The Jensen Huang Story

Table of Contents, Preface, 13 Chapters, Epilogue

Kim Kyung-jin

English PDF edition

The Jensen Huang Story

Kim Kyung-jin

The Jensen Huang Story is an online AI Library book by Kim Kyung-jin. It covers Jensen Huang, NVIDIA, GPUs, AI chips, and the AI industry.

Table of Contents

Preface

See every chapter

Next Chapter →

The reason I wrote this book is straightforward. One day in 2024, I saw a man in a black leather jacket on a YouTube screen. He was walking across the stage, explaining something with fierce intensity. Like a philosopher unraveling the nature of the world for his students, or a general laying out strategy for his troops. That man was Jensen Huang.

I have lived the life of a legal professional. I once served as a member of the National Assembly. Then I witnessed a massive wave called artificial intelligence rolling toward us. As I watched this new technology through a lawyer's eyes, I realized something. Artificial intelligence is not just a technological innovation. It is a fundamental turning point for human civilization.

And Jensen Huang stood right at the center of that turning point.

When you study Roman history, you discover something interesting. The people who built great empires were almost always born on the margins. Julius Caesar was a Roman aristocrat, but he came from a family pushed to the fringes of political power. Augustus was an ordinary young man born in a provincial town. Because they stood outside the existing order, they were the ones who could overturn it.

The same is true of Jensen Huang. He was born in the small city of Tainan, Taiwan. At nine years old, he crossed the Pacific to the United States and ended up at a boarding school in Kentucky, living alongside teenagers covered in tattoos. He learned to survive by scrubbing toilets and teaching his roommate how to read. He paid his way through school by washing dishes and waiting tables at Denny's, the family restaurant chain.

These experiences made him who he is. The immigrant boy from the periphery built an empire in the heart of Silicon Valley. In June 2024, NVIDIA became the world's most valuable company by market capitalization. It surpassed Apple. It overtook Microsoft. This happened 31 years after three young men founded a small company together on one of their thirtieth birthdays. By 2025, the market cap broke through $5 trillion, a figure exceeding the gross domestic product of Germany, the world's third-largest economy.

But this book is not meant to tell the success story of NVIDIA. I am not here to list how much the stock price climbed or how fast revenue grew. I wanted to look into the life of a human being. What trials he faced, how he overcame them, and what kept pushing him forward.

Nanami Shiono brought the history of ancient Rome vividly to life through her series The Story of the Romans. She portrayed Caesar not as a mere conqueror but as a human being with dreams and inner struggles. I wanted to do the same. I wanted to portray Jensen Huang not as a mythical figure of Silicon Valley, but as a person who opened his eyes every morning and pressed on despite anxiety

and fear.

Jensen Huang says that in 32 years of leading the same company, he has never felt at ease for a single day. He always thinks the company is on the verge of bankruptcy. Thirty days, forever. That is his secret. The statement is paradoxical. Even now, as one of the wealthiest people in the world, he is anxious. Yet that anxiety is what keeps him moving. He knows in his bones that the moment you settle, you fall behind.

While writing this book, I reviewed a great deal of material. I read Jensen Huang's interviews and speeches, along with articles and research about him. I studied the evolution of artificial intelligence and traced the history of the semiconductor industry. In the process, I came to understand something. This story is not just the biography of a single entrepreneur. It is the story of our era, a record of what is happening on the front lines as artificial intelligence reshapes the world.

I tried to write this book so that even a sixth grader could understand it. I unpacked complex technical terms and replaced difficult business concepts with simple analogies. At the same time, I tried not to lose a sense of dignity. Easy writing does not have to be lightweight writing. Anyone can read The Story of the Romans, but the insights within it are anything but shallow.

I hope this book inspires its readers. If there is something to learn from Jensen Huang's life, it is not genius or luck. It is the persistence to never give up and the willingness to keep learning. The boy who scrubbed toilets in Kentucky became the CEO of the world's most valuable company. His journey was filled with failures and setbacks. But every time, he got back up.

Greatness is not born from talent. It comes from the character to endure suffering. Jensen Huang said this at the Stanford University commencement. I believe this single line captures the heart of this book.

Now, let us begin. The story of a boy born under the Taiwanese sun in 1963.

January 1, 2026. Kim Kyung-jin

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 1. The Wandering Boy

Part 1. Building an Empire

← Previous Chapter

See every chapter

Next Chapter →

History is not the record of winners. It is the record of those who rose after defeat. Jensen Huang's life proves this.

Under the Tainan Sun

When Augustus, the first emperor of the Roman Empire, was born, no one imagined this frail child would one day command the entire Mediterranean. Jensen Huang's birth was much the same. On February 17, 1963, a baby boy let out his first cry in Taipei, the capital of Taiwan. His Chinese name was Huang Jen-Hsun (□□□). The boy who the world would come to know as Jensen Huang was the second son of an ordinary middle-class family.

His father, Huang Shing-Tai, was a chemical engineer working at an oil refinery. His mother, Luo Tsai-Hsiu, was a schoolteacher at a local elementary school. Both spoke Minnan, the traditional language of Taiwan, as their mother tongue. They built a household typical of Taiwanese families, one anchored in hard work and education. Shortly after the boy was born, the family moved south to Tainan, the ancient capital of Taiwan.

Tainan was an old city. Temples of red brick stood on every corner, and the scent of incense seeped through every alley. Street vendors called out their wares, bicycles streamed past in long lines, and people's lives were modest yet full of energy. To young Jensen, Tainan was the entire world, warm with sunshine and his family's love.

But the calm did not last. The year Jensen turned five, his father's work forced the family to leave Taiwan. Their destination was Bangkok, the capital of Thailand. His father had joined an oil refinery construction project there. Jensen and his older brother followed their father into unfamiliar tropical land.

Life in Bangkok was nothing like Tainan. The air was humid and scorching. The language, the food, the culture, all of it was foreign. Jensen attended Ruamrudee International School, where he slowly adjusted to English and his new surroundings. The family spent roughly four years in Bangkok. But the wheel of history would not leave this small family in peace.

In the late 1960s, the shadow of the Vietnam War stretched across all of Southeast Asia. The fighting crept closer and closer to the Thai border. American bombers flew over Cambodia and Laos, and inside Thailand itself, political turmoil and military coups followed one after another. Soldiers patrolled the streets. Anxiety hung in the air.

Jensen's father sensed it. This was no place to raise children. A few years earlier, he had traveled to the United States for professional training at Carrier, the air-conditioning manufacturer, at their New York plant. The America he saw then was a land of possibility. A vast continent, an atmosphere of freedom, and opportunity open to anyone. He made a quiet vow to himself: someday, I will send my children there.

His mother, too, was preparing for her sons' future. She could not speak a word of English. Yet she knew instinctively that her children would need the language to survive in America. She opened a thick English dictionary. Every day, she picked ten words at random and made her boys read them, memorize them, and explain their meanings.

She herself did not know the correct pronunciation of those words. She probably did not fully understand the definitions either. But she watched her children study, and she encouraged them never to give up. It was less a language lesson than an act of love. Years later, Jensen recalled those days.

"My mother didn't know any English at all. And yet she taught me English."

How was that even possible?

Jensen says the secret was his mother's persistence and devotion. She told her young son, over and over, "You can do it," and "You are going to do well on the test." Those words were not baseless encouragement. They were seeds of confidence planted in a child's heart. When Jensen Huang later faced countless impossible challenges, repeating to himself, "How hard can it be?", it was because of the belief his mother instilled during those years.

Nine Years Old, Crossing the Pacific

As the flames of war closed in on the Thai border, Jensen's parents made a momentous decision. The children, at least, had to be sent somewhere safe. But the family was not yet in a position to move together. They lacked the money, the visas, and the foundation to resettle.

The parents made a wrenching choice. They would send their two sons to America first. Jensen was nine; his brother was ten. Two young boys who could barely speak English had to cross the Pacific Ocean without a guardian.

In 1973, Jensen and his brother boarded a plane. As the boy stared out the window at the endless expanse of clouds and ocean, what could he have been thinking? He must have been afraid. He must have felt lost. And yet, perhaps somewhere in his chest, the seed of hope his parents had planted was already beginning to sprout.

Jensen's uncle lived in Tacoma, Washington. His parents had contacted this relative, a man the boys had never met, and asked him to take care of them. The uncle and his wife were recent immigrants, still finding their footing. Their means were limited, and they were hardly in a position to look after two young nephews.

The uncle searched for a school the boys could attend. He needed a boarding school that would accept foreign students. The tuition the parents could afford was tight. They had sold nearly everything they owned to send their children to America.

One school caught his eye. Oneida Baptist Institute in Kentucky. The photographs and descriptions in the pamphlet looked respectable enough. The uncle assumed it was a prestigious boarding school, something comparable to the well-known prep schools on the East Coast.

The reality was nothing of the sort. Oneida Baptist Institute was not a prep school. It was a place that took in boys who had dropped out of the regular American education system, troubled teenagers, young men with nowhere else to go. In today's terms, it sat somewhere between a juvenile detention center and an alternative school. Located in Clay County, Kentucky, one of the poorest areas in the country at the time.

When Jensen arrived, what greeted him was a rough dormitory. The rooms had no doors and no locks. Many had no dressers at all. Most of the students smoked. Each one carried a knife in his pocket.

Jensen was assigned to the dormitory with his brother. There was a problem: Jensen was too young. He was not old enough for the school's regular classes. So instead of attending the main campus, he had to go to a nearby elementary school, Oneida Elementary. His small frame, his long hair, his heavily accented English, and his Asian appearance. All of it made him stand out, and all of it made him an easy target for bullying.

Lessons from a Kentucky Boarding School

On his first night at Oneida, Jensen met his roommate. A seventeen-year-old white teenager. The age gap between them was eight years. The older boy lay on his bunk, his body covered in tattoos and lined with scars from knife wounds. He rolled up his shirt and showed Jensen every one of them.

To a nine-year-old boy, the sight must have been pure terror. It might have been the same dread Roman soldiers felt when they first came face to face with the barbarian tribes of Gaul. But Jensen had nowhere to run. His parents were on the other side of the Pacific, and international phone calls cost more than a child could imagine. He could not cry and beg them to come get him.

Oneida Baptist Institute had rules. Every student was required to perform manual labor each day. The school employed no custodians. Cleaning, maintenance, and farmwork all fell to the students. Jensen's brother was assigned to work at a nearby tobacco farm. And the younger Jensen? His job was cleaning the toilets.

Every day, Jensen had to clean the bathrooms in the boys' dormitory. He scrubbed toilets used by over a hundred rough boys, swept the floors, and cleared away foul-smelling waste. For a nine-year-old, it could have been a humiliating experience. But Jensen didn't complain. He did the work assigned to him without a word. Years later, he looked back on it this way.

"I probably cleaned more toilets than just about anyone."

Those years left Jensen with a lesson that stuck. No task is beneath you. If you give your best to even the lowliest job, that itself becomes a way of keeping your dignity. When he later became the head of a $3 trillion company, he never lost that attitude. He ate in the cafeteria alongside his employees, and if a restroom was dirty, he didn't hesitate to clean it himself.

Life in the dormitory was pure survival. Jensen was bullied often and sometimes beaten. The kids hurled racial slurs, and a small, slight Asian boy made an easy target. In a 2012 interview with NPR, Jensen described those days.

"The kids were really rough. Everybody carried a knife in their pocket, and when fights broke out, things got bad. Kids got hurt."

But Jensen didn't break. He learned to adapt. During his first week, he tried smoking to fit in with the other kids. Before long, though, he made a calculation that was perfectly sensible for a nine-year-old: the money spent on cigarettes was better spent on snacks and ice cream at the school store. So he quit.

Jensen found his own way to survive. It started with his roommate, a boy covered in tattoos. Jensen discovered something startling. The big, intimidating older boy couldn't read. A seventeen-year-old who was illiterate. Jensen made him an offer.

"I'll teach you how to read."

What did Jensen want in return? Protection. The older boy agreed to shield Jensen from the kids who picked on him. It wasn't a simple transaction. It was a partnership where each filled a gap the other had. Every night, Jensen taught his roommate the alphabet and read to him. In return, the older boy showed Jensen how to bench press and watched over him in the dormitory.

Through this experience, Jensen came to understand something. Even people with a rough exterior carry a desire to learn inside them. If you approach someone sincerely and offer what they need, you can become friends with a person from a completely different background. This became a real asset later, when he ran a company and had to work alongside all kinds of people.

Another treasure Jensen found at Oneida was sports. He picked up table tennis and joined the swim team. He was small, but at the ping-pong table, physical size didn't matter. Table tennis wasn't a game the strongest player won. It was a game for whoever could think fast and move with precision.

Jensen threw himself into table tennis. He learned to read his opponent's movements, calculate the spin of the ball, and control the momentum of a match. He stood out in swimming, too. Every night he did push-ups and sit-ups to build his strength. By the time he was fourteen, he was good enough to be featured in Sports Illustrated.

Communication with his parents happened through cassette tapes. International phone calls were too expensive. His parents bought cheap cassette recorders and sent them to the boys. Jensen and his

brother recorded everything that had happened over the past month onto a tape

and mailed it to Bangkok. Their parents recorded their own voices over the same tape and sent it back to Kentucky. A single cassette traveled back and forth across the Pacific, keeping the family connected.

One day, Jensen's school took the students to a nearby restaurant. It was a McDonald's. To a nine-year-old's eyes, McDonald's was a world of wonder. Gleaming lights, food served in boxes, and a delicious hamburger. Jensen recorded a message for his parents on the tape.

"Mom, Dad, I went to the most amazing restaurant today. Everything was bright and shining. It was like the future. The food came in boxes, and the hamburger was so good."

Years later, Jensen laughed when he recalled that moment. To the boy he was then, McDonald's was not just a fast-food restaurant. It was the "taste of the future," a symbol of all the abundance and possibility that America represented.

Those two years at Oneida formed the backbone of the person Jensen Huang became. In a 2002 interview, he made this confession.

"I remember my time in Kentucky more vividly than any other period of my life."

Many of those memories were painful. But the pain made him tough. Hitting rock bottom can destroy a person, yet some people push off that bottom and spring back up. Jensen was someone who rose from the bathroom floor.

Two years later, his parents finally immigrated to the United States. The family settled in Beaverton, Oregon, and brought the brothers back from Kentucky. The family was whole again. Jensen and his brother left Oneida and enrolled at Aloha High School in Oregon.

Life in Oregon wasn't easy either. An immigrant family's means were limited. His father scoured newspaper ads looking for work, and his mother took whatever job she could find. Jensen never forgot the look of helpless dismay on his mother's face when a piece of rented furniture broke.

But the family refused to give up. Jensen threw himself into his studies, skipped two grades, and graduated from high school at sixteen. In table tennis, he climbed high enough to earn a national ranking. He joined the math, science, and computer clubs, honing his talents along the way.

The harsh winter winds of Kentucky and the cold air of the dormitory had forged a nine-year-old boy into something hard and resilient. He was no longer a wandering child. He had become a small warrior who understood that he would have to carve out his own fate. The world is unfair and rough, but he had learned this: for those who refuse to quit, who keep grinding and sharpening their skills, opportunity will come.

Years later, Jensen Huang stood before students at Stanford University and told them this.

"I wish upon you enough suffering and hardship. Greatness doesn't come from talent. It comes from enduring pain."

Those were not empty words. The boy who scrubbed toilets in Kentucky, who taught his tattooed roommate to read, who was moved by a McDonald's hamburger. All of those experiences, layered one on top of another, built the Jensen Huang who stands today.

In 2019, Jensen Huang donated $2 million to the Oneida Baptist Institute, the school he had once attended. The money funded a building named Jensen Huang Hall, a dormitory and classroom facility for girls. The boy who once cleaned toilets had become a benefactor, giving other children a chance to learn.

Augustus, the first emperor of Rome, faced countless trials in his youth. Power struggles, assassination threats, failing health. Yet he rose above every hardship to build an empire that brought peace across the Mediterranean world. Jensen Huang's story follows a similar arc. From Tainan to Bangkok, Bangkok to Kentucky, Kentucky to Oregon. The wandering boy was thrown into new surroundings each time, and each time he adapted, learned, and grew stronger.

The wandering was over. But the lessons it taught him became the compass for a grand journey just beginning. The boy knew he could survive wherever the world dropped him. And beyond surviving, he knew he could build a kingdom of his own.

From Tainan to Kentucky, what the boy crossed was not the Pacific Ocean. It was the boundary of his own limits.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 2. The Waiter at Denny's

Part 1. Building an Empire

← Previous Chapter

See every chapter

Next Chapter →

Scrubbing Dishes, Washing Away Dreams

Some people believe history begins in grand places. Real history is rarely that glamorous. Alexander the Great grew up on the fringes of Macedonia, living in his father's shadow.

Jensen Huang's story is no different. His empire didn't start in a gleaming Silicon Valley lab. It started in an ordinary family restaurant, the kind you'd find in any American city, where the air hung thick with grease and the smell of pancakes.

1978. A Denny's in Portland, Oregon. The restaurant stayed open twenty-four hours a day, a place where late-shift workers and broke college students nursed cheap coffee refills to kill time. Fifteen-year-old Jensen Huang was there. He had only recently left the rough boarding school in Kentucky and reunited with his parents. The boy held a rag in his hands instead of a textbook, and a hazy dream of the future flickered somewhere in his mind.

His first title was Dishwasher.

A corner of the kitchen, the sink shrouded in steam, was his battleground. Plates from finished meals piled up without end. Plates smeared with ketchup, cups sticky with syrup, pans coated in grease. The boy scrubbed them one by one. His hands turned red in the scalding water, but he didn't complain.

Jensen Huang would later recall: "I was the greatest dishwasher." There is no bluster in that statement.

He set his own rules for washing dishes. He never walked empty-handed. When he carried plates to be put away, he always found another task to bring back with him, and the same was true on the return trip to the sink.

Wasting time was something he simply could not tolerate. Efficiency. Optimization. These words would become the core philosophy behind the semiconductor chips he would one day design. But the seed had already sprouted in the kitchen at Denny's.

Having proved himself as a dishwasher, he was soon promoted to busboy, clearing tables after customers left and setting them up for the next guests. Before long, he had risen to the position of waiter. This was

more than a promotion. For the introverted, shy immigrant boy, waiting tables was a challenge.

Approaching strangers. Explaining the menu. Apologizing and fixing things when an order went wrong. Every one of these felt as daunting as climbing a mountain. His English was still awkward, and his Asian features made him stand out from the people around him. But he didn't run from it. This was a boy who had fought his way through fistfights and cigarette smoke in Kentucky. He was already learning, in his bones, how to walk straight into discomfort.

His time at Denny's stretched from 1978 to 1983, nearly five years. During that span, Jensen Huang finished high school and entered college. By day he sat in lecture halls studying electrical engineering; by night he took orders from customers at the restaurant. When the graveyard shift ended, dawn was breaking. He carved time from sleep to do homework, then headed back to class. Sleep was scarce. His body was exhausted. But Jensen never called it suffering. It was just his life.

Years later, at a Stanford University commencement address, he would say: "My greatest superpower is that my expectations are very low. I wish upon you ample doses of pain and suffering." That was not a polished piece of rhetoric tossed out to sound inspiring.

In the Denny's kitchen, in the bone-deep fatigue of early mornings, he learned how to endure hardship. A temperament that doesn't crack under pressure. The strength to get back up after failure. These things don't come from intelligence. They come from character forged over long years.

Denny's taught Jensen Huang something else, too: the philosophy that no task is beneath you.

He didn't just wash dishes. He cleaned toilets without hesitation. In a later interview, he put it this way: "I cleaned a lot of toilets. More toilets than all of you combined, I guarantee." There is no self-pity in his voice, no bitterness. If anything, you can hear a quiet pride.

That philosophy runs through Nvidia's corporate culture to this day. Even now, as CEO, Jensen Huang digs into his employees' problems personally. Rather than scolding someone for being late to a meeting, he jumps in and helps solve the problem alongside them.

He has a phrase he uses often: "Have a good past, and it will follow you." His good past was not a sparkling resume. It was the days when he wore a grease-stained apron and quietly scrubbed dishes. That is his past. And that past made him who he is now.

There is a twist of fate worth noting. In 1993, Jensen Huang, now thirty years old, decided to start a new company with two colleagues. The place they met was none other than a Denny's. A corner booth at a Denny's restaurant in San Jose, California. They drank bottomless coffee and debated the future of graphics chips. "The place had endless coffee, and nobody kicked us out," Jensen recalled.

In the same restaurant chain where the boy had once scrubbed dishes, the grown man founded a company that would change the world. Denny's was his first workplace and his final launchpad. He sweated there, he nurtured his dreams there, and he sketched the blueprint of an empire there. In 2023, after Nvidia became a trillion-dollar company, Denny's corporate headquarters installed a

commemorative plaque at that historic booth. It reads: "The booth that launched a trillion-dollar company."

While scrubbing dishes, Jensen Huang was also washing away his old self. The worn-out things, the fears, the feelings of inferiority. They swirled down the drain with the soapy water. What remained was a will as clear and hard as polished glass.

An Afternoon in the Lab Where He Met Lori

The electrical engineering lab at Oregon State University. Wires tangled everywhere, oscilloscope lights blinking in the dim room. In this space, Jensen Huang's fate took another turn. This time it wasn't a circuit board but a person's heart that bent the arc of his trajectory.

In 1980, Jensen Huang entered college at seventeen. Having skipped two grades, he was far younger than his classmates. Sitting in the lecture hall, he stood out. A boyish face, a small frame. As he would later recall, "I was the kid in the classroom who looked like a child."

The electrical engineering department had 250 students. Only three were women. One of them was Lori Mills. She was nineteen, two years older than Jensen. She was also majoring in electrical engineering, and as it happened, she had been assigned to the same lab section. The two became lab partners.

Jensen was captivated from the moment he first saw Lori. He did not take the usual approach. Instead of asking her to a movie or suggesting dinner, he devised a strategy that was thoroughly, unmistakably that of an engineering student.

One day, Jensen walked up to Lori and asked: "Do you want to see my homework?"

There was a careful calculation behind it all. Jensen had already completed his homework perfectly in advance. When he compared answers with Lori, his plan was to let her see how smart he was, as naturally as possible. He figured his young appearance would actually work in his favor. He wanted to leave the impression: 'This kid is incredibly bright for his age.'

As time passed, Jensen took it a step further.

He made Lori a proposal. 'If you do homework with me every Sunday, I'll make sure you get A's in every subject.'

At first glance, it looked less like asking someone on a date and more like striking a deal.

But the offer brimmed with confidence. Lori looked at him with curious eyes. And she accepted.

That was how their Sunday dates began.

Calling them dates might be generous on the romance front. They huddled in a corner of the library, sketching circuit diagrams. Solving equations. Writing lab reports. But for Jensen, it was enough. Every Sunday, all of Lori's time belonged to him. He stretched the homework out all day long,

maximizing every minute with her.

Lori gradually opened her heart to Jensen's brilliance. What she saw was not just a student with good grades. She saw a fire in his eyes. A hunger to achieve something, a will to push forward into the world. It was not the recklessness of a young boy. It was the kind of seriousness you could stake a life on.

Around the age of twenty, Jensen made Lori a promise.

'By the time I'm thirty, I'll be the CEO of a company.'

At the time, this must have sounded absurd. A college student barely into his twenties, talking about becoming a CEO. But Jensen's gaze did not waver.

Years later, he would say: 'I had absolutely no idea what I was talking about back then.' And yet he said it. It was a vow to prove his worth in front of the person he loved. And he kept that promise. In 1993, the year he turned exactly thirty, he became the CEO of NVIDIA.

What Lori gave Jensen was not just love. She gave him emotional stability. As an immigrant carrying anxiety and fear about the future, he finally had someone to share all of it with.

Lori never laughed at Jensen's dreams. She never dismissed his ambition as fantasy. She watched the diligence he learned washing dishes at Denny's, the focus he showed designing circuits in the lab. And she believed in him.

The two married in 1984, five years after they first met, the same year Jensen graduated from college.

At the time, Jensen had no prestigious job, no comfortable savings. But Lori didn't care. What she chose was not a man's present circumstances but his potential to race toward the future.

Lori was an engineer herself. After graduating, she worked on microchip design at Hewlett-Packard. The two shared a common language: voltage and current, circuits and semiconductors. It was not just academic knowledge. It was the thread that bound them together.

Years later, when Jensen Huang founded NVIDIA and faced near-bankruptcy multiple times, Lori stood by his side without flinching. In 1996, when the company was in a desperate situation where they might have to close the doors within thirty days, she was still there. Lori remembered that look in Jensen's eyes from the lab. She believed he would not break.

Jensen and Lori have two children.

A son, Spencer, and a daughter, Madison. Both currently work at NVIDIA.

Spencer is a product manager, and Madison is a marketing director. Some might criticize this as a family business, but for Jensen, family is simply another word for trust.

In 2022, Jensen and Lori donated fifty million dollars to their alma mater, Oregon State University. The money is funding the construction of the Jen-Hsun and Lori Huang Collaborative Innovation Complex, a facility for research in artificial intelligence and robotics.

That afternoon in the lab where they first met is now extending into a new generation's laboratory.

The love that began on a rainy Oregon campus continues today, more than forty years later. Jensen Huang often mentions his wife in interviews.

His signature black leather jacket, he says, was picked out by his wife and daughter. 'My wife and daughter dress me. That makes me happy.' A simple, honest confession from a man worth hundreds of billions of dollars.

From Oregon to Stanford

In 1984, Jensen Huang graduated from Oregon State University. Bachelor of Science in Electrical Engineering. With highest honors.

He held a diploma in his hands, and in his heart burned a desire aimed at Silicon Valley. But he knew. A single college degree was not enough to build the world he dreamed of. He needed deeper knowledge, a sharper weapon.

After graduation, his first destination was AMD (Advanced Micro Devices).

He interviewed at Texas Instruments, LSI Logic, and several other companies, but chose AMD. Its California location and its position as a challenger competing against Intel in the semiconductor industry drew him in.

His time at AMD lasted just over a year. Brief as it was, he built his foundation in microprocessor design. How CPUs work, how instructions are processed, how the microscopic circuits etched onto a chip cooperate to perform calculations. He learned all of this on the job.

Jensen's curiosity didn't stop there. He moved to another company, LSI Logic. LSI Logic was a leader in the application-specific integrated circuit (ASIC) field.

At LSI Logic, Jensen didn't remain a mere engineer. He expanded his experience beyond technical roles into marketing and management. Eventually he rose to the position of director of a business unit.

During the day, he worked intensely at the company. But his day didn't end when he clocked out. At night, he became a student again.

He enrolled in Stanford University's evening master's program.

Stanford. It was more than a university. Situated in the heart of Silicon Valley, it was a crucible where brilliant minds from around the world came to design the future. Hewlett-Packard, Sun Microsystems, Cisco. Countless technology companies were born on Stanford's campus. Jensen Huang threw

himself into that heat.

Working by day and studying by night was grueling. During the day he ran projects at LSI Logic; at night he attended lectures at Stanford. On weekends he caught up on assignments and looked after his two children.

He was also a married husband and father. He carved time from sleep, compressed every hour, and kept pushing forward.

In 1992, he finally received his master's degree in electrical engineering from Stanford.

Eight years had passed since joining AMD, seven since starting at LSI Logic. The young immigrant who had graduated from Oregon's state university now held a degree from Stanford, the breeding ground of Silicon Valley's elite.

His time at Stanford gave him something beyond a diploma.

There he began thinking not just about 'what' to build, but 'why' it needed to be built. A way of thinking that digs into the essence of a problem, an approach that starts from fundamental principles (First Principles Thinking). This would become the philosophical foundation driving Nvidia's innovations in years to come.

At LSI Logic, Jensen met two important people.

Chris Malachowsky and Curtis Priem. Both were veteran engineers who had come from Sun Microsystems.

LSI Logic had a contract with Sun Microsystems at the time to develop a graphics accelerator card, and Jensen ended up working alongside Chris and Curtis on that project.

The three hit it off.

They met frequently after work to talk about the future of technology. At the time, PCs could barely handle text and 2D graphics. Rendering 3D graphics required workstations costing tens of thousands of dollars. But the three of them saw a different possibility.

"What if ordinary PCs could render 3D graphics?" "What if games could move as vividly as movies?" "What if, instead of a general-purpose CPU, there were a separate processor specialized for graphics?"

These questions circled through their minds. Gradually they grew certain. The old ways could not open the future of computing. A new architecture, a new philosophy was needed.

In 1992, the three gathered often at Denny's. A Denny's restaurant on the east side of San Jose. The very same chain where Jensen had once washed dishes. They sat in a corner booth, refilling their coffee while sketching out a business plan. "The coffee was unlimited, and nobody kicked us out," Jensen recalled.

Jensen went to a bookstore to learn about starting a company. He picked up a book called How to Write a Business Plan. Four hundred fifty pages. He flipped through a few and put it down. "By the time I finish this book, I'll be bankrupt." His way was never to learn from a book. It was to jump in headfirst.

In April 1993, the three founded Nvidia.

The name traces back to the Latin word 'invidia' (envy). It carried the ambition of making everyone envy their chips.

Jensen was thirty years old. The promise he had made to Lori, "I'll be a CEO by thirty," had become reality.

Their startup capital was $40,000. It was everything the three of them had saved.

But Jensen had credibility built up at LSI Logic. The company's founder, Wilfred Corrigan, thought highly of Jensen.

He introduced Jensen to Don Valentine, a venture capitalist. Valentine, who headed Sequoia Capital, decided to invest in NVIDIA.

Twenty million dollars in initial funding was secured.

Jensen Huang became CEO from day one.

He was the youngest of the three, but Chris and Curtis handed him the leadership. "Jensen, you run the company. The things we don't know, that's your job." Those were Curtis Priem's words.

They knew Jensen was not just an engineer. He was a rare talent who understood both technology and business.

The journey that began on a rainy campus in Oregon passed through Stanford under the hot California sun, and finally came together at a corner booth in Denny's.

The boy who once washed dishes was now a thirty-year-old CEO, taking his first step toward the world.

But this was only the beginning. Failure, setbacks, and the brink of bankruptcy lay ahead. History does not make heroes easily. Before Jensen Huang could become a true empire builder, there were many more mountains to climb. He had crossed the Rubicon, but the road to Rome was long and treacherous.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 3. Crossing the Rubicon

Part 1. Building an Empire

← Previous Chapter

See every chapter

Next Chapter →

1993: The Decision on His Thirtieth Birthday

In January of 49 BC, the Roman general Julius Caesar was marching south with his legions after finishing his campaign in Gaul. Before him flowed a river called the Rubicon. It was a small river. On horseback, you could cross it in moments. But this river carried the weight of law.

Roman law forbade any general leading an army from crossing this river into Rome. Cross it, and you become a traitor. Don't cross it, and your political enemies take everything from you.

Caesar stood at the riverbank and thought for a long time. Then he spoke.

"The die is cast."

He crossed the river. He chose the path from which there was no return.

Every great figure in history who made a decisive choice had their own Rubicon to cross.

In 1993, in California's Silicon Valley, such a river appeared before a young man approaching his thirtieth birthday. His name was Jensen Huang.

February 17, 1993. Jensen Huang turned thirty. The ancient Chinese sage Confucius called the age of thirty 'erli,' meaning the age when one stands on one's own. The age when one establishes an unwavering purpose. But for Jensen Huang, the number thirty carried a meaning more concrete than that. It was the hour of a promise kept.

Let us turn the story back a few years. The electrical engineering lab at Oregon State University. Seventeen-year-old Jensen Huang met a young woman there. Her name was Lori. She was a student in the same department. The two became lab partners, then friends, then lovers. One day Jensen told Lori:

"By the time I'm thirty, I'll be a CEO."

It was the bold declaration of a seventeen-year-old. It could have been the kind of bluster any young man utters in front of the person he loves. Did Lori believe him? Did she not?

What matters is that Jensen Huang himself never forgot those words. Years passed, the two married, and the age of thirty was approaching.

Let us look at where Jensen Huang stood in 1993. He was working at a semiconductor company called LSI Logic. His title was Director. Before that, he had designed microprocessors at AMD, the well-known chip company. He held a master's degree in electrical engineering from Stanford University. His salary was high, and he had earned the respect of those around him. He was, by any measure, a successful engineer.

Most people would have been satisfied with this. A stable job, a comfortable income, a respected position. For a young immigrant from Taiwan who had come to America, achieving this level of success was remarkable. His mother thought so too. She told her son:

"Why don't you look for a good job?"

In his mother's eyes, her son was doing well enough as things were. There was no need to take unnecessary risks. But a different idea had been growing inside Jensen Huang's mind.

He was reading the shifts taking place in the computer industry.

1993 was a time when personal computers were spreading rapidly. As Microsoft's Windows operating system reached the masses, computers were no longer machines reserved for specialists. But Jensen Huang saw something beyond that.

The future of computers lies not in text, but in images.

At the time, most computer screens were filled with text. Making spreadsheets in Excel and writing documents in word processors were the primary uses of a computer.

As for graphics, simple two-dimensional images were all there was. Three-dimensional, or 3D, graphics were a technology possible only on specialized workstations costing tens of thousands of dollars. The idea of playing 3D games on a home PC sounded like something out of a science fiction movie.

But Jensen Huang saw it differently.

He was certain that before long, 3D graphics would become possible on PCs. And he resolved that he would be the one to build the chip that made it happen.

As his thirtieth birthday drew closer, the urgency in his heart grew.

Would he stay on the comfortable bank of a secure job, or leap into the rough current of starting a company? He had to choose one of two paths. He remembered the promise he had made to his wife Lori. The words about becoming a CEO by thirty.

Jensen Huang made his decision. He would leave LSI Logic.

It meant burning the bridge behind him. If he failed, he could lose everything. The stable job, the high salary, the respected position could all vanish. But he chose to trust his own instinct. He knew where the future of technology was headed. At least, he believed he did.

When Caesar crossed the Rubicon, he cried out, "The die is cast." Jensen Huang threw his own die. It was 1993. He was thirty years old. He crossed the river.

Three young men in a corner booth at Denny's

Every legendary Silicon Valley company has its birthplace.

Apple started in the garage of Steve Jobs's parents' house. Hewlett-Packard was born in a garage in Palo Alto. There's even a saying in Silicon Valley: "Great companies start in garages."

Nvidia was different.

This company wasn't born in a garage. It was born in a family restaurant. The name was Denny's, an ordinary 24-hour chain with locations across the United States.

Berryessa Road in San Jose, California. The Denny's there held a special meaning for Jensen Huang. It was where he first worked at age fifteen. He washed dishes, cleared tables, carried food to customers. It was also the place where a shy immigrant boy learned how to talk to strangers.

In 1993, Jensen Huang, now thirty, returned to Denny's. This time he hadn't come to wash dishes. He had come to start a company.

Two people came with him.

Chris Malachowsky. He was an engineer at Sun Microsystems, a well-known computer company.

Curtis Priem. Also from Sun Microsystems, he was a veteran technologist who had worked at IBM as well.

The three of them settled into a corner booth at Denny's. A booth is a seat with high backs that creates a semi-private space away from other diners. They sat down and ordered coffee.

Why Denny's, of all places? Jensen Huang later explained it this way.

"It was quieter than home, and the coffee was cheap."

Denny's offered free refills on coffee. Pay for one cup and you could drink as many as you wanted. The three young men took full advantage. They would meet in the morning and stay until evening, sitting in the same spot, talking for hours. Refilling their coffee ten cups at a time.

Jensen Huang later looked back on those days.

"The waitstaff probably wanted to kick us out."

Customers who barely ordered food but occupied a table for hours on end. The restaurant couldn't have been happy about it. But the Denny's staff never chased them away. Thanks to that, the three young men could talk as long as they pleased.

What were they talking about?

3D graphics.

At the time, the personal computer screen was a two-dimensional world.

But change was happening in the world of games.

Arcade game centers already had games running 3D graphics. Hollywood films were using more and more computer graphics. The movie Jurassic Park, released in 1993, stunned audiences around the world by making computer-generated dinosaurs look as if they were alive and moving.

The three young men had an idea.

Soon, home PCs would be able to run 3D graphics like these. And they would be the ones to build the chip that made it possible.

The CPU, the central processing unit, serves as a computer's brain. It is the core component that handles all calculations.

But rendering 3D graphics requires a staggering number of calculations. For each of the millions of pixels on screen, the system must compute where the light comes from, how colors shift, and how objects move.

A CPU alone couldn't handle that workload.

The three young men had an idea.

Build a dedicated chip to help the CPU. Create a new brain devoted solely to graphics processing.

This was the beginning of what would later be called the GPU, the graphics processing unit.

A napkin lay spread on the table. They sketched circuit diagrams on it and scribbled down ideas.

Inside that restaurant, where the smell of pancake syrup hung in the air, the three young men's eyes were shining.

They divided up the roles.

Chris Malachowsky would handle the hardware design.

Curtis Priem would take charge of chip architecture and algorithms.

Jensen Huang would be responsible for business strategy and management.

Curtis Priem asked Jensen Huang a question.

"You're going to be CEO, right?"

Jensen Huang nodded. The promise he had made to Lori at seventeen was about to become reality. Choosing a company name turned out to be no simple task either. At first they considered 'NVision,' a name meant to suggest 'new vision.' But a company with that exact name already existed. It was apparently a toilet paper manufacturer.

They needed a name. They flipped through a Latin dictionary.

That was when they found the word 'invidia.'

It meant 'envy.' Curtis Priem said:

"Let's make something that turns our competitors green with envy."

And so the company had its name. NVIDIA. Taken from the Latin word for envy.

April 5, 1993. Jensen Huang went to see a lawyer to officially register the company. The lawyer's name was James Gaither. He asked Jensen Huang for the capital required to incorporate. Jensen Huang dug through his pockets. He had 200 dollars. The lawyer said that would be enough.

Jensen Huang called Chris and Curtis.

"Can you each put in 200 dollars?"

Both of them did.

NVIDIA's initial capital came to 600 dollars. Converted to Korean won, that is roughly 700,000 won. The fact that a company now worth trillions of dollars started with 700,000 won is, by any measure, astonishing.

That corner booth at Denny's. Nobody knew then that the drawings three young men sketched on a napkin would one day change the world.

A declaration to create a market where none existed

NVIDIA had been founded. But there was a problem. They had a product to sell, yet no one to sell it to.

In 1993, who would buy a 3D graphics chip?

PC users were perfectly content running Excel and word processors.

Some people loved games, but most thought 2D graphics were good enough. 3D gaming itself barely existed yet.

Jensen Huang went around meeting investors. The responses were cold.

"We can see your technology is impressive. But who's going to buy it? Where's the market?"

The investors' questions were reasonable.

To start a business, you need a market. You need customers who will buy your product. But the market for what NVIDIA wanted to build did not exist. It was a zero-dollar market.

Jensen Huang called it the "Zero Billion Dollar Market." A market worth zero dollars at present. A market with nothing in it. And yet that was exactly where he aimed.

An ordinary executive would have looked for a market that already existed.

Grabbing even a small share of a large market is the safe strategy. Business textbooks say the same thing. But Jensen Huang thought differently. "We're not going to fight over market share. We're going to create the market."

It was a bold declaration. He was saying he would build a market where none existed. To anyone asking whether the chicken or the egg came first, he was answering that he would lay the egg himself.

His logic went like this.

If we build an affordable, powerful 3D graphics chip, game developers will use that chip to create 3D games.

Once 3D games appear, gamers will buy our chip to play those games.

Once the chips sell, we can build even better chips.

Better chips will produce better games. A virtuous cycle begins.

Jensen Huang focused on video games. Why games, of all things? He explained it this way.

"Gaming is the only field that poses the hardest technical challenge while also commanding the highest sales volume."

Building a 3D game requires enormous computing power, because the screen must be rendered in real time. A movie can show pre-drawn images. Games are different. The scene must change instantly depending on where the player looks. That is an extremely difficult technical problem.

At the same time, games sell in massive quantities.

You watch a movie once and you're done, but people play games over and over. A hit title sells millions, even tens of millions of copies. High sales volume means more money to invest in research and development. That was the opening Jensen Huang was targeting.

Reality, though, was unforgiving.

NVIDIA's first product, the NV1, was a failure.

The performance was solid, but it was too expensive. On top of that, it did not conform to the standards Microsoft had set. Insisting on proprietary technology turned out to be the fatal mistake. The product did not sell, and the company was driven to the brink of bankruptcy.

Jensen Huang later recalled the situation like this.

"We always worked as if we were 30 days away from going out of business."

Thirty days. One month. That was how close they were to shutting down. There was no money to pay employees. Funding for continued R&D was running out. In the five years after founding the company, they faced bankruptcy three separate times.

Still, Jensen Huang refused to give up. He cut his own salary to one dollar. The CEO would take just one dollar. It was a message to his employees: I believe in this company, and I've staked my future on it. I'm asking you to believe, too.

Then he changed direction.

Instead of insisting on proprietary technology, he decided to follow Microsoft's standards. The team developed a new product called the NV3. This time, the market responded well. The company survived.

In 1999, NVIDIA finally released something the world had never seen: the GeForce 256. Alongside this product, Jensen Huang coined a new term. GPU. Graphics Processing Unit. Before that, no product had carried such a name. He was the one who gave it that name.

It was the moment when the dream sketched on a napkin at Denny's in 1993 became real.

Jensen Huang later said:

"If I had known how hard it would be to start a company, I probably never would have started one."

His ignorance became his courage. Because he didn't know, he wasn't afraid. Because he wasn't afraid, he could take the leap.

When Caesar crossed the Rubicon, there was no guarantee he would win. When Jensen Huang founded NVIDIA, there was no guarantee he would succeed. But both men crossed the river. And both changed history.

In September 2023, Jensen Huang returned to that Denny's. The same restaurant where he had discussed his startup thirty years earlier. Denny's CEO Kelli Valade greeted him. She presented him with a plaque. A nameplate was affixed to the corner booth where he had sat.

"The booth where a trillion-dollar company was born."

Jensen Huang stood at that spot and said:

"Get your first job at a restaurant. Working in a restaurant teaches you what it means to work hard. You learn humility. You learn kindness."

A fifteen-year-old boy washing dishes. A thirty-year-old man crossing the Rubicon. A CEO past sixty leading one of the world's most valuable companies. All of it is one person's story.

His declaration that he would create a market where none existed came true. The 3D graphics chip market has grown to hundreds of billions of dollars. And the GPU did not remain a chip just for gaming. As a new world called artificial intelligence opened up, the GPU became the heart that powers it. We will return to this story later in the book.

What we should remember now is this: the decision three young men made that day in 1993, in a corner booth at Denny's, opened the door to the age of artificial intelligence.

The die was cast. And it landed on numbers that would change the world.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 4. Thirty Days of Survival

Part 1. Building an Empire

← Previous Chapter

See every chapter

Next Chapter →

They say history is written by the victors. But true history is written in the moment of defeat. When Rome suffered a catastrophic loss to Hannibal at the Battle of Cannae, senators screamed that they should abandon the city and flee. A young man named Scipio drew his sword and declared: "Anyone who swears to abandon Rome is an enemy of Rome." Great empires are forged not in victory but in the moment they rise from defeat. The story of a company called NVIDIA is no different.

When three young men shared their dreams at a Denny's restaurant in 1993, they had no idea they would soon be pushed to the brink of death. Jensen Huang founded his company on his thirtieth birthday, and within just two years he faced the crisis of losing everything. This is not a story about technology. It is a story about how one person stands back up from the depths of despair.

The Disastrous Failure of the First Chip

NVIDIA's first creation was a graphics chip called NV1.

In May 1995, Jensen Huang and his colleagues unveiled the chip to the world. Their hearts were full of pride. NV1 was not just a graphics card. It was an ambitious chip that could handle 3D graphics, 2D graphics, and audio all at once. Plug it into a computer, and you didn't need to buy a separate graphics card and sound card. It could even connect to a controller for the Sega Saturn, a game console from the Japanese gaming company.

The problem lay in the technology they had chosen.

At the time, there were two main methods for drawing three-dimensional images on a computer screen.

One method stitched together countless tiny triangles to represent objects, building cars and people by assembling triangles much like stacking LEGO bricks. The other used quadrilaterals to render smooth curves.

Jensen Huang and his colleagues chose the quadrilateral approach. They believed quadrilaterals could draw curves more beautifully than triangles.

Technically speaking, they weren't wrong. When rendering a round ball or a human face, quadrilaterals could produce smoother results. But the world wanted convenience over beauty.

In 1995, Microsoft, the reigning emperor of the computing world, released Windows 95. Alongside it came something called DirectX. DirectX was a toolkit for game developers. The problem was that Microsoft built this toolkit around the triangle method. It did not support the quadrilateral approach.

Overnight, NVIDIA became a stranded island.

Game developers all rallied under Microsoft's banner. Games ran smoothly on Windows when built with triangles. Meanwhile, the number of games that supported NVIDIA's NV1 could be counted on one hand. Sega's Panzer Dragoon and Virtua Fighter were about all there was. It didn't matter how good your graphics card was if there were no games to run on it.

The chips didn't sell. Unsold inventory piled up in the warehouse. The company's coffers began running dry. NVIDIA had to let go of more than half its employees. A staff that had numbered in the dozens shrank to roughly thirty. Every morning when Jensen Huang opened the office door, one thought ran through his mind: Will I still be able to open this door tomorrow? Silicon Valley was not a place that readily gave second chances to those who had failed.

Do you know what happened to a Roman consul who lost a war? He was dragged before the Senate and humiliated, or he had to take his own life. Jensen Huang in 1995 must have felt something similar. The thirty-two-year-old CEO had to accept that the technology he had believed in so deeply was wrong.

Honesty Saved the Company

At the very end of despair, a single ray of light appeared.

Sega, the Japanese gaming company, extended a hand to NVIDIA. Sega was preparing a new game console called the Dreamcast. It needed a weapon to compete against Sony's PlayStation. Sega asked NVIDIA to build the graphics chip for the Dreamcast. The project was called NV2. A contract worth five million dollars was promised as an advance payment.

For NVIDIA, it was a lifeline.

But the goddess of fate tested Jensen Huang with terrible cruelty.

While developing NV2, Jensen Huang came to a dreadful realization. The approach Sega wanted was also based on quadrilateral technology. But the world had already shifted entirely to triangles. With Microsoft's DirectX becoming the standard, a game console built on the quadrilateral method would be incompatible with Windows games.

Jensen Huang saw the truth.

If he finished NV2 as contracted, Sega would release a game console out of step with the Microsoft era. And NVIDIA itself would be branded as a company clinging to an isolated technology, doomed to vanish for good.

He was trapped.

If he built the chip as agreed, the company would receive its payment right away.

But that chip would fail due to compatibility issues with Microsoft, and the company would collapse in the end.

If he stopped development, on the other hand, it would be a breach of contract. Without the money, the company would go bankrupt immediately. No matter which path he chose, death waited at the end.

Jensen Huang spent several sleepless nights wrestling with the decision. Then he made his choice.

He boarded a plane to Japan.

He was going to meet Shoichiro Irimajiri, the CEO of Sega.

Irimajiri was a legendary engineer who had originally worked at Honda Motor Company. He was the man who put Honda's motorcycles on the world stage. Now he had moved to Sega and was leading its gaming business.

The air in the conference room was heavy. Jensen Huang spoke in a voice that trembled but remained firm.

"The technology we promised you is wrong."

Irimajiri's eyebrows twitched. Jensen Huang continued.

"Microsoft's DirectX has become the standard. The chip we're building right now, even if we finish it, won't be compatible with that standard. If Sega builds a game console around this chip, that console will fail. I would recommend that Sega look for another company."

What would a typical executive have done? Hidden the problem, stalled for time, strung together excuses. But Jensen Huang chose honesty. He admitted his own failure and put his customer's interests first.

Up to this point, anyone might have said the same thing.

But then Jensen Huang added something almost unbelievable.

"However, Nvidia has no money. If we're going to redesign the chip with the right technology and succeed, we need the remaining $5 million you committed under our contract. Without that money, we go bankrupt. Give us that money, and we will start over in the right direction."

By any normal standard, the request made no sense.

Here was a subcontractor who had come to disclose problems with the deliverables, fully aware the contract might be terminated, and yet was asking for the full remaining payment. Most executives would have thrown him out in a rage. Jensen Huang himself expected to be refused.

In a later interview, he recalled the moment this way: "If we hadn't gotten that money, Nvidia would have evaporated on the spot. It would have vanished in an instant."

Irimajiri was silent for a long time. His eyes studied Jensen Huang's face. What did Jensen Huang see in that gaze?

Perhaps Irimajiri saw his own younger self. The days at Honda when he threw himself into motorcycle racing, when he refused to fear failure.

Irimajiri opened his mouth.

"We will do that."

Jensen Huang could not believe his own ears.

Irimajiri instructed his team to pay Nvidia the remaining contract balance, even though Nvidia would no longer be making the chip for the Dreamcast.

Sega had already decided to use another company's chip, and yet it helped Nvidia.

Years later, Jensen Huang recalled this moment on Joe Rogan's podcast: "I'm grateful for that man's kindness. And I surprised even myself."

Why did Irimajiri give money to a company that had failed?

It went beyond business logic. He saw honesty in the young man named Jensen Huang. He saw someone who admitted his own mistakes, who put his customer's interests first, and who still burned with a refusal to give up.

Irimajiri invested in the person, not the technology.

That $5 million became the last drop of blood that brought Nvidia back from the dead.

A Gamble at the Edge of the Cliff

With the money from Sega, Nvidia was back at the starting line. This time there was no room to repeat past mistakes. Jensen Huang set aside every last scrap of pride. He abandoned the quadrilateral method he had so stubbornly championed and fully embraced the triangle-based approach that Microsoft had established as the standard.

But there was a problem. They didn't have enough money.

When you make a semiconductor chip, you have to go through a process called tape-out. That means physically producing your designed chip and testing it. A single tape-out costs hundreds of thousands of dollars. If the chip has problems, you redesign it and produce it again. Each time, more money goes out the door. Normally, finishing one chip requires multiple tape-outs.

NVIDIA couldn't afford that. With $5 million, they could only produce the chip once. One failure and it was over.

Jensen Huang made a radical decision.

He spent half the remaining money on a piece of equipment called an emulator. An emulator is a device that lets you test a chip on a computer before actually manufacturing it. At the time, it was extremely expensive and unfamiliar technology.

The plan was to simulate the chip perfectly in software, then go straight into mass production without any physical test runs. This was unprecedented in the semiconductor industry. If the simulation missed any errors, they would be left with hundreds of thousands of useless chips.

Jensen Huang called TSMC, the foundry company in Taiwan. He commissioned chip production from TSMC's founder Morris Chang. They were launching mass production without being certain the chip would actually work.

NVIDIA's employees worked around the clock. They cut the design period from one year to seven months. Everyone worked until dawn. Jensen Huang stood before his team and said:

"We only have 30 days left."

In 1997, a new chip called the RIVA 128 finally came into the world.

It was a massive success.

The RIVA 128 was faster and cheaper than 3dfx's Voodoo chip, which dominated the market at the time. Most importantly, it fully supported Windows DirectX. Game companies cheered.

The chip sold like wildfire. One million units shipped within four months of launch.

A company that had been at death's door came back to life with a single, all-or-nothing bet.

When Sega later sold its NVIDIA shares, the payout came to about $15 million. Not a bad investment for Sega. But what if they had held those shares until today? As of 2024, NVIDIA's market capitalization exceeds $3 trillion. That $5 million investment would have been worth over $1 trillion.

"We are always 30 days from going out of business"

This experience left an indelible mark on Jensen Huang's mind.

He realized that a company's life is as fragile as a glass cup. No matter how brilliant your technology, a single wrong choice can erase everything. The moment you grow complacent from success, death comes knocking.

From that point on, one sentence was etched deep into Jensen Huang's mind.

"We are always 30 days from going out of business."

What does 30 days mean?

It means that once the company pays its employees one month's salary, the vault is empty. That's how urgent it is. You have to sell something before the money runs out. If you can't sell, the company shuts down.

This phrase became NVIDIA's unofficial motto. Jensen Huang never forgets it, even now that the company generates trillions of won in profit and the whole world watches its every move.

In 2023, appearing on Joe Rogan's podcast, Jensen Huang said: "I've been using this phrase for 33 years. But the feeling never changes. The feeling of vulnerability, uncertainty, anxiety. It doesn't go away."

He added: "Every morning I wake up with this feeling. 'We might be going under soon.' It's no different from what I felt waking up this morning."

This is not exaggeration or self-pity.

Jensen Huang wakes up at 4 a.m. every day to check his email. He reads thousands of emails a day. Seven days a week, on Thanksgiving, on Christmas. His two children, Madison and Spencer, also work at NVIDIA, and like their father, they work every day.

"It's exhausting work. You're always in a state of anxiety."

Why does he push himself this hard? He already owns the most valuable company on the planet.

It goes back to 1996. The memory of choosing the wrong technology and being driven to the edge of a cliff. The memory of bowing his head before Sega. The memory of nearly losing everything because of a single mistake. That memory became the whip that keeps him running forever.

There is a story about Roman generals during their triumphal processions. They would have a slave stand behind them and whisper in their ear.

"Memento mori. Remember that you will die." Even amid the cheers of victory, never forget that death comes for everyone. For Jensen Huang, the phrase "thirty days" is his memento mori. Even at the peak of success, he remembers what bankruptcy felt like.

During a 2024 lecture at Stanford University, Jensen Huang told the students something unexpected.

"I wish upon all of you enough pain and suffering."

The audience was stunned. Wishing pain on people in a graduation speech? But Jensen Huang was serious.

"People with high expectations have low resilience.

Low resilience makes you prone to unhappiness.

Unhappiness makes success difficult. People who have been through pain learn to lower their expectations. When your expectations are low, you become grateful for small things. Gratitude makes you happy. And happiness lets you succeed."

This was the lesson Jensen Huang learned at rock bottom in 1996.

The Weapon of Honesty

The Sega episode became a Silicon Valley legend.

An ordinary executive would have hidden the problem. He would have stalled for time, found excuses, dodged responsibility. But Jensen Huang chose honesty. He admitted his own ignorance and asked for help.

That honesty saved him.

NVIDIA has a core value called "intellectual honesty." When you know something is going wrong, you don't hide it to save face or protect your ego. You admit it immediately and correct course. Jensen Huang tells his employees:

"We survived as a company because we admitted we were wrong and asked for help."

At NVIDIA, no one hides failure. They talk about it openly. In meetings, employees present their own mistakes. They share why they failed and what they learned. It's embarrassing at first. But Jensen Huang believes it is the only path to growth.

One longtime employee put it this way: "When someone starts getting defensive, I know they won't last long here."

Admitting failure is not easy.

It stings your pride, it's humiliating, it's frightening. But Jensen Huang says it takes far more courage than dreaming of success. And he believes only those with that courage can truly succeed.

The Foundation of an Empire

The empire called NVIDIA was not built on glorious victories. It was built on devastating defeats.

The failure of the NV1 taught them humility. They learned that no matter how brilliant a technology might be, it cannot survive if it goes against the current of the world. The negotiation with Sega taught them the value of trust. They learned that honesty can sometimes be the key to survival. And the terror of being "thirty days from bankruptcy" turned them into wild horses that never stop running.

Now they had moved beyond survival. They were ready to change the world.

The RIVA 128 in 1997 was only the beginning.

Two years later, in 1999, NVIDIA coined a word the world had never heard before.

GPU.

Graphics Processing Unit. A device that processes images. Just as a computer has a CPU serving as its brain, the concept was that there could be another brain, one dedicated to handling graphics.

At the time, nobody truly understood what the term meant. Most people saw it as a fancy label slapped on a graphics card. But Jensen Huang knew. He knew the GPU would one day change the world. Not just by making game visuals look better, but by becoming the key that would reshape the future of humanity.

That story continues in the next chapter.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 5. GPU as a Weapon

Part 1. Building an Empire

← Previous Chapter

See every chapter

Next Chapter →

When a Roman legionnaire drew his sword on the battlefield, it was no mere piece of metal. The gladius, short but razor-sharp, was the weapon that let Rome conquer the Mediterranean. A good weapon changes how battles are fought, and a changed style of warfare redirects the course of history.

Silicon Valley in 1999 looked a lot like the Mediterranean coast in the third century BC.

Just as countless city-states had clashed over supremacy in that ancient era, graphics chip companies were crowding the field and strangling one another. 3dfx, ATI, Matrox, S3, and NVIDIA. Amid this tangle of names, Jensen Huang was quietly sharpening his blade.

His blade had a name. GPU. Those three letters did not exist anywhere in the world before 1999.

1999: Inventing a Word That Never Existed

Caesar was great not just because he conquered vast lands. He planted Roman law and language in the lands he took. To name a new concept is to become its owner. Jensen Huang understood this truth by instinct.

On August 31, 1999, NVIDIA announced a new chip and introduced an unfamiliar term. The chip, called GeForce 256, was presented as the world's first GPU. GPU: short for Graphics Processing Unit.

Why was this word, which seems obvious today, so revolutionary? To understand that, you need to know how the computer world worked in 1999.

In those days, the computer had only one king. The CPU, the Central Processing Unit. This chip, made by Intel, governed everything that happened inside a computer. Writing a document, running a calculation, playing music, running a game; everything had to pass through the CPU.

The CPU was the computer's brain and heart.

So what was a graphics card? It was little more than a servant that splashed the CPU's processed results onto the monitor. People called it a graphics accelerator. As the word "accelerator" suggests, it meant a helper device that sped up the CPU's work a little. The chip's name made its limitations plain from the start.

Jensen Huang questioned this arrangement.

Was rendering three-dimensional graphics on screen really a simple task that needed only a little assist from the CPU?

His answer was an emphatic no.

Drawing a three-dimensional object on screen requires complex math.

You have to calculate the position of each object, calculate where light is coming from, and calculate where shadows fall. These calculations must happen millions of times per second for the image to move smoothly. The CPU had to handle all of this on its own while also computing enemy movements inside a game, so the burden was enormous.

Jensen Huang thought: Isn't this a fundamentally different kind of computation from what the CPU does? If the CPU is a single genius mathematician, wouldn't graphics calculation be better served by thousands of ordinary clerks working at the same time?

He decided to give the graphics chip a new name. He declared it would no longer be called an accelerator. He defined it as an independent processing unit that served as a second brain for the computer, standing on equal footing with the CPU.

NVIDIA defined the GPU this way:

"A single-chip processor that integrates transform, lighting, triangle setup and clipping, and a rendering engine, capable of processing a minimum of 10 million polygons per second." One phrase in this definition deserves attention: Processing Unit.

Until that moment, "Processing Unit" had been a sacred title reserved for the CPU alone. Jensen Huang dared to attach that title to his own chip. This was not just a marketing term. It was a declaration of independence, a challenge thrown at the CPU's kingdom.

People in the industry scoffed.

A graphics card calling itself a processor? But language has power. The moment people started using the term GPU, the chip's standing began to shift inside their heads. It was no longer seen as a helper device but as an independent brain.

Whoever controls the name controls the concept, and whoever controls the concept controls the market. Jensen Huang grasped this truth exactly.

GeForce 256: The Opening Act of Parallel Computing

Having staked out a new name, he had to prove the substance behind it. On October 11, 1999, NVIDIA officially launched the GeForce 256. Carrying the title of the world's first GPU, the chip lived up to expectations.

To explain what set the GeForce 256 apart from earlier graphics chips, you first need a basic idea of how three-dimensional graphics are created.

The 3D characters and buildings you see on screen are made up of countless small triangles called polygons. Every time a character moves, the new position of each of those thousands of triangles must be calculated. This is called coordinate transformation, or in technical terms, "Transform."

Next comes calculating the effects of light.

You have to determine where the sun is shining from, where shadows fall, which parts of an object are bright and which are dark. This is called lighting processing, or simply Lighting. Combined with the first word, Transformation, the two are known as T&L.

Before 1999, T&L calculations fell entirely on the CPU. The graphics card only stepped in to paint colors after the CPU had finished its math.

Think of it this way. On a construction site, the architect draws the blueprints, calculates coordinates, and decides the direction of light before the painter even picks up a brush. That was how it worked.

The GeForce 256 flipped this structure on its head. It built a T&L engine right into the chip. Now the CPU only needed to issue rough commands. "Put a pillar over there."

The GPU would then calculate the pillar's coordinates on its own, analyze the direction of light, and render the shadows. The CPU was freed from this grueling work.

So what happened when the CPU had less to do?

The CPU could now focus on computing enemy AI and physics inside the game. Characters got smarter. Explosions and falling objects looked more realistic. The entire gaming experience jumped up a level.

The GeForce 256 proved itself in raw numbers too.

The chip could process over 10 million polygons per second. That was more than 50 percent faster than competing products at the time, such as 3dfx's Voodoo3 or NVIDIA's own previous chip, the RIVA TNT2.

Gamers cheered. The monsters on screen no longer looked like boxy crates. They felt like living creatures with bulging muscles and glistening skin.

But Jensen Huang was looking somewhere far beyond the game screen. He was focused on the essence of the T&L engine: the concept of parallel computing.

What is parallel computing? Here's a simple way to picture it.

Imagine a classroom where thousands of exam papers, each containing 20 questions, need to be graded. The CPU approach is like having one brilliant teacher pick up each paper and grade all 20

questions in order, one sheet at a time. It's not slow, but it has to go through them one by one.

The GPU approach is like assigning each teacher to grade only one specific question across all the papers. If there are 20 questions, 20 teachers each handle their own. The overall speed is far faster this way.

3D graphics is exactly this kind of situation.

When millions of pixels on the screen need to be calculated at the same time, thousands of workers beat one genius. The GeForce 256 was the first product to implement this philosophy of parallel processing in hardware.

Jensen Huang was quietly looking ahead at this moment.

What if this parallel processing power could be used for something other than games? What if it could be applied to scientific calculations or data analysis? The idea hadn't yet taken concrete shape.

But it would later bloom under the name CUDA.

And that CUDA would ride the massive wave of artificial intelligence and change the world.

The GeForce 256 of 1999 was a dazzling toy for gamers.

But in the history of computing, this chip carried a far greater meaning. It marked the moment when parallel computing appeared as a new way of thinking in a world that had been dominated by serial processing.

Just as the steam engine appeared in a world ruled by horse-drawn carts, the very method of computing itself began to change.

The Fall of 3dfx and the Rise of a New King

History has no permanent champion. The Roman Empire couldn't last a thousand years. Genghis Khan's empire fractured within a generation. The history of technology is no different.

In the mid-1990s, the undisputed ruler of the 3D graphics market was a company called 3dfx. Their graphics card, the Voodoo, was both legend and religion among gamers. Owning a Voodoo card was every teenager's dream in those days. At its peak, 3dfx commanded a market share of roughly 85 percent.

When NVIDIA's first chip, the NV1, failed miserably, it was 3dfx that ruled the market. In most people's eyes, NVIDIA was just another chaser in a crowded pack. A dwarf standing before a giant. But the throne of number one is a poisoned chalice. Those drunk on victory grow arrogant, and the arrogant fail to read change.

3dfx's first mistake was misjudging technical standards.

They insisted on Glide, their own proprietary programming language.

Glide was a closed technology that only worked on 3dfx cards. If a game developer built a game in Glide, that game only ran properly on 3dfx hardware.

At first, this looked like an advantage. 3dfx cards dominated the market, after all. But when Microsoft started pushing DirectX, an open standard, the picture changed. Games built on DirectX ran on any graphics card. Developers began tilting toward DirectX. They could sell their games to a far larger audience.

NVIDIA read the shift quickly.

Instead of clinging to proprietary technology, Jensen Huang chose to fully support open standards like DirectX and OpenGL. He called it an "ecosystem" strategy.

The idea was simple: don't try to do everything alone. Work with others and grow together.

3dfx's second mistake was restructuring its business model.

Originally, 3dfx only designed chips. Multiple manufacturers built and sold the actual graphics cards. Then in 1998, 3dfx acquired STB, a graphics card manufacturer. The ambition was to control everything from chip design to card manufacturing to retail sales. On paper, it looked like a smart way to capture all the profit.

Reality told a different story.

The moment 3dfx started selling cards directly, the other manufacturers who had been building cards with 3dfx chips turned hostile. They lined up behind NVIDIA instead. 3dfx found itself fighting alone, while NVIDIA commanded an alliance.

Jensen Huang took the opposite path.

He focused solely on chip design. The actual fabrication went to a Taiwanese company called TSMC. Building finished cards from those chips was left to dozens of partners: ASUS, MSI, Gigabyte, and many others. When NVIDIA made a good chip, scores of companies sold cards carrying that chip under their own brands, all around the world.

The way Rome folded conquered nobles into alliances rather than making enemies of them, Jensen Huang never turned his partners into adversaries. This ecosystem strategy became the backbone of the NVIDIA empire.

3dfx's third mistake was misreading the technology trend.

They dismissed 32-bit color. 32-bit color could render far richer, more natural visuals than 16-bit. But 3dfx argued that 32-bit color would drag down frame rates, and stuck with 16-bit. "Gamers care about speed, not color depth," was their verdict.

NVIDIA saw it differently.

Jensen Huang decided that combining 32-bit color with hardware T&L could deliver both speed and image quality. The result was the GeForce 256. The moment people saw a GeForce screen, the Voodoo started to feel like yesterday's hardware.

The most decisive difference was speed, specifically the speed of product launches.

Jensen Huang set a fearsome target: a new chip every six months. Moore's Law, the prediction by Intel co-founder Gordon Moore, held that chip performance doubled every 18 to 24 months. Jensen Huang was chasing a pace far faster than that.

For NVIDIA's employees, it was a death march.

But that punishing tempo let NVIDIA release the RIVA 128, RIVA TNT, RIVA TNT2, and GeForce 256 in rapid succession. While competitors paused to catch their breath, NVIDIA already had the next product on shelves.

While 3dfx struggled to develop its next-generation chip, the GeForce 256 hit the market. Armed with hardware T&L, it made 3dfx's latest products look like relics. Gamers were ruthless. They said goodbye to yesterday's hero and cheered for the new king.

On November 16, 2000, 3dfx announced it would stop manufacturing cards. It was a declaration to go back to selling chips only. But it was already too late. A month later, on December 15, 3dfx's creditors initiated bankruptcy proceedings, and the company agreed to sell its assets to NVIDIA.

The acquisition price was $70 million in cash and $1 million worth of NVIDIA stock. A shabby sum for an empire that once held 85 percent market share. Jensen Huang picked up the fallen rival's sword and strapped it to his belt. 3dfx's engineers and intellectual property were absorbed into NVIDIA.

With that, the long first act of the graphics war ended in NVIDIA's victory. A formidable competitor named ATI still remained, but NVIDIA had become the undisputed ruler of the market.

The fall of 3dfx left an important lesson in the history of technology.

Technical prowess alone isn't enough to win. You need the vision to read where standards are heading, a strategy for building an ecosystem, and the speed to adapt when things change. Without all three, survival is not guaranteed.

3dfx was no slouch in technology, but the company stumbled on the other three fronts.

Jensen Huang did not let victory go to his head. Watching 3dfx collapse, he must have felt the old terror of being "30 days from bankruptcy" all over again. Today's winner can become tomorrow's loser. The moment you get comfortable, you fall apart. That was the lesson 3dfx taught him.

The success of the GeForce 256 brought NVIDIA a massive influx of cash.

The question was where to spend it. A typical executive would have paid dividends to shareholders or poured the money into expanding existing operations. Jensen Huang was not a typical executive.

He decided to invest it in a place nobody else was willing to look at.

The power of parallel computing he had witnessed through the GeForce 256 felt too valuable to confine to gaming alone.

What if thousands of tiny calculators working simultaneously could be applied to scientific research or data analysis? What if they could replace supercomputers?

This idea would eventually reach the world under the name CUDA. But it would prove to be a lonely investment, generating zero revenue for more than a decade.

Wall Street laughed. Analysts called it a waste. Jensen Huang did not stop.

In the next chapter, we will talk about that lonely decade of investment, and how a gamble that everyone dismissed as insanity rode a massive wave called artificial intelligence and changed the world.

Three letters Jensen Huang coined in 1999: GPU. It was not just a marketing term. It was the invention of a concept that altered the history of computing. It opened a door from the serial era, where a single CPU handled everything alone, to a parallel era, where thousands of small brains worked together. Beyond that door, a new world called artificial intelligence was waiting.

Jensen Huang did not yet know the specific shape of that world.

But he sensed that was the direction to go. "We started this company to solve problems that general-purpose computers cannot solve." Those words carry a consistent philosophy spanning from the 1993 founding to the 1999 invention of the GPU to the arrival of the AI era that followed.

Just as Roman legionaries armed with the gladius, a short sword, conquered the Mediterranean, NVIDIA, armed with the GPU, began building a new empire in the digital world.

But a true empire is not built overnight. It is completed only at the end of a solitary journey, staying up through countless nights, walking a road that no one else believes in.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 6. Ten Years of Lonely Investment

Part 1. Building an Empire

← Previous Chapter

See every chapter

Next Chapter →

Every great turning point in history began with a lonely decision. When Caesar crossed the Rubicon, even his own legionnaires didn't fully grasp what that choice meant. All they could see was the cold current and the darkness on the far bank.

For Jensen Huang, the year 2006 was that river. Not long after he had claimed the throne of graphics cards to the cheers of gamers everywhere, he steered his company toward a direction nobody was asking for.

After the success of the GeForce 256, a wide-open road stretched out before NVIDIA.

Flashier graphics, faster frame rates, more gamers. That alone could have kept the company growing. But Jensen Huang's gaze reached somewhere beyond the flames pouring from dragons on the screen. What he saw was an uncharted territory called computation.

CUDA: The Gamble Wall Street Laughed At

In late 2006, Jensen Huang made a decision that would go down in Silicon Valley history. It was the launch of a software platform that would transform the GPU, a graphics processing unit, from a gaming device into a general-purpose computing engine.

Its name was CUDA, short for Compute Unified Device Architecture.

To understand what CUDA was, you first need to look at the computing world of that era. Back then, a computer had two kinds of brains. One was the CPU, the central processing unit, which was like the class president who did reasonably well in every subject. Writing, math, drawing. But it could only do one thing at a time.

The other was the GPU, the graphics processing unit, which was like that kid in art class who could draw incredibly fast. Inside a GPU were hundreds, even thousands, of tiny workers that could splash countless dots, lines, and colors onto a screen all at once.

Here's what Jensen Huang was thinking: what if those thousands of workers, assembled for art, tackled math problems simultaneously? Instead of the class president solving problems one by one, wouldn't thousands of workers charging in together finish far more quickly?

There was a problem, though.

To get thousands of art workers to do math, you had to give them orders in a language they could understand. CUDA was precisely that language.

It was an interpreter that let programmers use C, a language they already knew, to assign tasks to the GPU's thousands of workers.

The trouble was that this decision would hurt the company right away.

Jensen Huang decided to put CUDA capability into every graphics card NVIDIA made, whether it was built for gamers or for professionals. It was an extraordinarily strange decision.

When you manufacture a semiconductor chip, adding one more feature makes the chip bigger.

A bigger chip raises manufacturing costs, generates more heat, and draws more power. For gamers, scientific computing features were completely useless. They wanted monsters to look more lifelike; they had no interest in calculating molecular structures.

Yet Jensen Huang crammed this seemingly pointless feature even into the budget graphics cards that gamers bought.

Wall Street's reaction went beyond cold. It bordered on fury.

Securities analysts pored over NVIDIA's financial statements and shook their heads.

From 2006 to 2017, roughly $12 billion, more than 15 trillion Korean won, was poured into research and development.

In 2014, over 30 percent of revenue flowed into R&D. A large share of that went toward building the software ecosystem around CUDA.

Investors kept pressing Jensen Huang with the same question: "Who on earth is going to use this technology? Where is the market?" In their eyes, CUDA was a monster that devoured money without producing any revenue.

The stock price nosedived. When the 2008 global financial crisis hit on top of everything else, NVIDIA's share price plunged more than 80 percent from its peak.

Discontent erupted inside the boardroom, too. "Why don't you shut this division down? We're doing perfectly fine with gaming alone."

Analysts wrote in their reports: "CUDA is nothing but a cost center with unproven returns. NVIDIA should focus on its core business of gaming graphics."

From 2009 to 2015, NVIDIA's market capitalization bobbed between $8.5 billion and $10 billion. Growth appeared to have stalled. Wall Street valued CUDA's market worth at zero dollars. That meant it was considered worthless.

But Jensen Huang refused to back down.

He repeated this in meeting rooms. "The future of computing is accelerated computing. CPUs alone can't keep up with this speed. If we don't go down this road, someone else will eventually find their way to it."

Jensen Huang was looking at a different timeline. CPU performance improvements were hitting physical limits. There were limits to how small semiconductor circuits could be made. Moore's Law, the rule that computer performance doubles every year, was gradually slowing down.

Jensen Huang was convinced that the only way to break through those limits was parallel computing, a method where thousands of small processors work at the same time.

It was the cry of a lonely prophet. Nobody believed him. But just as the Romans laid roads across every corner of their empire, he was building the highways that future data would travel on. There were no carriages to ride on them yet.

But he believed that someday all the world's information would be racing down these roads.

Handing Supercomputers to Scientists

While Wall Street was laughing at NVIDIA, Jensen Huang was out meeting an entirely different group of people. They

weren't the wealthy guests invited to glamorous parties.

They were the underfunded scientists pulling all-nighters in the corners of university labs, the graduate students sighing over research budgets, the young researchers stomping their feet because they had theories to prove but no tools to prove them with.

To understand their situation, you have to look at the world of scientific computing at the time.

Imagine a scientist studying weather prediction. To figure out whether it will rain tomorrow, you need to run incredibly complex calculations. Temperature, humidity, air pressure, wind direction and speed; all of it has to be plugged into mathematical formulas. These calculations are so complex that an ordinary computer can't finish them even after running for days.

That's why scientists use supercomputers. A supercomputer links thousands of computers together to calculate at tremendous speed. The problem is that supercomputers are extremely expensive. They cost tens of millions of dollars, sometimes hundreds of millions. Only governments and large corporations could afford them.

So what did cash-strapped researchers do?

They submitted applications to national supercomputer centers. Then they waited for months. When their turn came around, they could borrow time on a supercomputer for a short window. And that was only if they were lucky. If a researcher's topic was outside the mainstream, or if their institution's

backing was weak, their number might never come up at all.

Jensen Huang saw an opportunity here.

What if a single GPU he built could match the computing power of a supercomputer costing hundreds of millions of won? What if a graphics card you could buy for a few hundred thousand won at the Yongsan electronics market could do the same work that scientists waited in line at a national supercomputer center to perform?

It was the democratization of computing.

NVIDIA began donating GPUs to universities and research labs around the world, or supplying them at low cost. And it distributed CUDA, its software platform, for free. It was like Rome building aqueducts to deliver water to the lands it had conquered.

Once the water flowed, shoots began to sprout in what had been barren ground.

In graduate research labs around the world, the whir of graphics card fans running through the night became a familiar sound.

Chemists studying molecular structures, biologists analyzing gene sequences, geologists interpreting seismic waves, physicists exploring the origins of the universe; they all started buying NVIDIA's chips.

Calculations that used to take a week were finishing in a single day. Simulations that required months of waiting could now run on a desk in someone's own room.

Reflecting on this period, Jensen Huang once said: "We didn't hand scientists a microscope. We turned them into telescopes."

Inside NVIDIA, people described CUDA this way:

"CUDA is the switch that turns the gaming graphics card plugged in under your desk into a budget supercomputer."

Those words carried real weight.

Researchers no longer had to wait for government approval to do their work. They no longer had to persuade a panel of grant reviewers.

With a single graphics card and CUDA, they could start today what they had considered impossible yesterday.

Jensen Huang poured resources into this community that generated no revenue. CUDA training programs were established at more than 200 universities worldwide. He dispatched engineers to help developers solve their problems. He kept adding libraries for researchers, pre-built bundles of program code, one by one.

Even Nvidia employees were puzzled.

"Why should we be solving problems for physicists?"

Jensen Huang's answer was firm. "Because they are building the future."

He wasn't selling hardware. He was planting an ecosystem.

Scientists who had grown fluent in CUDA couldn't easily switch to another company's chips.

It would mean more than inconvenience. It would mean a rupture in their research.

This ecosystem would later become Nvidia's powerful moat. A moat is the trench of water dug outside a castle wall, a defensive line that keeps invaders out. While competitors scrambled to boost their chips' raw performance, Nvidia had already built an ecosystem that hundreds of thousands of scientists used every day.

A quantum chemist once said this to Jensen Huang: "Jensen, thanks to your technology, I can finish my life's work within my lifetime." Huang said he was deeply moved. He later recalled: "If that doesn't move you, you're basically a corpse."

It was a confirmation of purpose, that technology could expand the boundaries of human knowledge. Moments like these were part of what kept him going through Wall Street's ridicule.

Betting Everything on a Zero-Dollar Market

Business textbooks say this:

Make the product the market wants. Survey demand, then supply goods that match it. That is the foundation of business.

Jensen Huang made a decision that ran directly against this principle.

When CUDA first launched in 2006, the market for general-purpose computing on GPUs was zero.

Nobody was buying it. Nobody thought they needed it. Jensen Huang staked his company's fate on precisely this zero-value territory.

He called it "the zero-billion-dollar market." And he said:

"We decided to build something where there were no customers. Because where there are no customers, there are no competitors either."

The logic is deeply paradoxical.

Most businesspeople charge into fiercely contested markets and fight. They try to poach each other's customers. They battle over market share. But Jensen Huang said: "We don't fight for market share. We create markets."

The pain of this strategy was beyond description.

In the early 2010s, the smartphone era had arrived and the mobile chip market was exploding. Companies like Intel and Qualcomm were earning astronomical sums in mobile.

People advised Nvidia to make mobile chips too. Nvidia did, in fact, produce a mobile chip called Tegra. But Jensen Huang eventually decided to pull out of that market.

"We don't exist to do what everyone else does, only better. We exist to do what no one else can."

Instead of the mobile market, already a blood-red ocean, he chose AI computing, a wasteland no one bothered to look at. It must have felt like being a farmer quietly tilling his field while a glamorous party raged next door.

His strategy can be summed up in three parts.

First, he evolved hardware and software together.

Every time a new generation of GPUs came out, CUDA was upgraded to match.

Tesla, Fermi, Kepler, Maxwell, Pascal... With each new GPU architecture, CUDA expanded to take advantage of more cores, more complex memory structures, and higher bandwidth.

Second, he built an ecosystem.

NVIDIA didn't just offer a programming language. The company created tools that scientists could use right away and distributed them for free. A library for linear algebra, a library for deep learning, a library for image processing. A forest of software grew on top of CUDA.

Third, he invested in education.

NVIDIA taught CUDA at universities, research labs, and developer conferences. Through developer forums and sample code, the company opened a path for anyone to learn on their own.

By the numbers alone, all of this looked like a losing bet. But what Jensen Huang saw was potential energy.

Problems requiring parallel computation grew every year. The volume of data flowing onto the internet was increasing exponentially. Simulation, analysis, and a field that didn't even have a proper name yet, artificial intelligence, were demanding more and more computing power.

CUDA simply hadn't caught fire yet. Huang was digging the channels in advance so that when that massive wave of demand arrived, it would have somewhere to flow.

And finally, the moment came.

In 2012, something remarkable happened at the University of Toronto in Canada.

Alex Krizhevsky and Ilya Sutskever, students of Professor Geoffrey Hinton, won the ImageNet artificial intelligence competition by an overwhelming margin. The name of their AI was AlexNet.

AlexNet was an AI that identified objects in photographs.

Until then, the gap between first and second place had been one or two percentage points at most. AlexNet left the runner-up behind by more than ten percentage points. The error rate dropped from 26 percent to 15 percent in one leap.

What matters here is how AlexNet was trained.

Krizhevsky trained AlexNet in a bedroom at his parents' house using two NVIDIA GTX 580 gaming graphics cards. A calculation that would have required thousands of CPUs was done by two GPUs. The tool he used was CUDA.

Yann LeCun, a leading researcher in computer vision, said of this event: "It was an obvious turning point in the history of computer vision."

When Jensen Huang heard the news, he knew instantly. "This is it. The moment we've been waiting for has arrived." He immediately redirected every resource in the company toward artificial intelligence. "We are no longer a graphics company. We are an artificial intelligence company."

Suddenly, IT companies around the world turned their attention to NVIDIA. Google, Facebook, and Microsoft began buying up GPUs for AI research. What had been a toy for scientists was becoming the beating heart of the Fourth Industrial Revolution.

In 2016, Jensen Huang unveiled another groundbreaking system: the DGX-1, the world's first AI supercomputer. It linked eight GPUs together to deliver 170 teraflops of computing power, a figure almost unimaginable at the time. The price was $129,000, roughly 170 million Korean won.

When Huang first presented the machine on stage at GTC, NVIDIA's developer conference, the audience's reaction was cold. "There was no applause, no interest, just 100 percent confusion." People asked, "Does this thing run Windows?"

Only one place, a small startup in San Francisco, recognized what this machine was worth.

Its name was OpenAI. Led by Elon Musk and Sam Altman at the time, OpenAI was nothing more than a nonprofit research lab. But they desperately wanted a DGX-1.

Jensen Huang loaded the first supercomputer into his car and drove it to OpenAI's office himself. Joking that he was a "DoorDash delivery driver," he handed the machine directly to Elon Musk and Sam Altman.

This delivery was not just a hardware drop-off. It was the act of handing over a seed from which a future AI empire would grow. OpenAI used this DGX-1 to train its early language models. That became the decisive stepping stone leading to the birth of ChatGPT.

What if Jensen Huang had given up on CUDA ten years earlier, unable to withstand Wall Street's criticism? What if he had chased short-term profits and focused solely on improving gaming performance? The AI revolution we know today might never have happened. Or it might have started much later, in someone else's hands.

Jensen Huang's ten years were not just a period of holding on.

It was the time a seed spent putting down roots in dark soil. On the surface, nothing seemed to change. But underground, a vast network was forming. NVIDIA's legion of scientists, researchers, and students was growing in every corner of the world.

Just as Rome was not built in a day, NVIDIA's empire was not built on a single brilliant chip. It was raised over ten long years on a foundation of a software ecosystem built while enduring ridicule, and an unshakable faith in the future.

Jensen Huang puts it this way.

"We have been preparing for this moment for 30 years. It looks like an overnight success, but it actually took 30 years."

His lonely investment has come full circle. It now beats at the heart of a new industrial revolution: the AI factory. Without NVIDIA's GPUs, the miracles of artificial intelligence we enjoy today might never have existed.

This is a triumph of technology. But it is, above all, a triumph of character, the kind that endures pain.

As Jensen Huang himself said, "Greatness does not come from talent. It comes from character, the kind that can endure suffering." His reckless gamble on a space worth zero in market value came back as a revolution that changed the world.

The preparation is over. The world stands ready to be pulled, swirling and roaring, into the very future Jensen Huang foresaw.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 7. The Big Bang Moment

Part 1. Building an Empire

← Previous Chapter

See every chapter

Next Chapter →

The moments when history changes course usually arrive quietly.

No grand fanfare, no advance notice. In some small corner of a lab, or during an idle afternoon conversation, the spark that will alter the fate of humanity begins to catch. That is what happened in 2012.

That year, people were still cheering over the latest smartphones. Gamers were thrilled by ever more dazzling graphics. But in a lab at the University of Toronto in Canada, an event that would later be called the Big Bang of artificial intelligence was quietly taking shape.

2012: The Day AlexNet Changed the World

The story begins with a competition called the ImageNet Challenge.

ImageNet was a contest in the field of computer visual recognition.

The task was to show a computer a photograph and have it tell whether the image was a cat or a dog. For humans, this judgment is as natural as breathing. For a machine built on zeros and ones, it was an insurmountable wall.

For decades, brilliant minds had tried and failed. Accuracy barely budged.

Computer vision scientists at the time were struggling. They tried to write algorithms by hand, defining what made a cat a cat and a dog a dog. Pointy ears meant cat; tongue sticking out meant dog. That sort of thing. But cats and dogs in the real world come in endless variety. No fixed set of rules could cover every case.

Then Professor Geoffrey Hinton at the University of Toronto and his students turned their attention to a different approach.

Alex Krizhevsky and Ilya Sutskever, two young researchers working under Hinton's guidance, attempted something radical. They dusted off an old theory called deep learning, a deep neural network modeled after the neural pathways in the human brain.

Deep learning as a concept had existed for a long time. Its core idea was to build a computer system where countless neurons connect and learn on their own, just like the human brain. The

theory was elegant, but there was a practical problem. Processing all those calculations required enormous computing power. With conventional CPUs, training could take months or even years. That is why the field had been pushed to the margins of academia.

Here, a fateful choice was made.

The research team decided not to use a supercomputer. Instead, they turned to a gaming graphics card that anyone could buy at a Best Buy or a local electronics shop. It was NVIDIA's GeForce GTX 580, priced at just over 500 dollars. Nobody imagined that a chip designed to render flashy game visuals would become the key to artificial intelligence.

Alex connected two GTX 580 cards in parallel. He used CUDA, the programming language Jensen Huang had released for free six years earlier, to train the neural network. The training took place in a bedroom at Alex's parents' house. Not hundreds of servers, but a single desktop computer on a desk, running an experiment that would change the world.

Calculations that would have taken months on a CPU finished in just a few days. Thousands of GPU cores worked simultaneously, processing matrix operations in parallel. The core computations of the neural network turned out to be an astonishingly perfect match for the GPU's parallel architecture.

On September 30, 2012, the ImageNet competition results were announced.

Their creation, AlexNet, slashed the image recognition error rate from 26% to roughly 15% in one leap. The gap over second place was more than 10 percentage points. This was not incremental progress. In an era when computer vision experts had been fighting for decades to shave off 0.1%, three researchers with gaming graphics cards had rewritten history.

The venue was stunned.

Yann LeCun of New York University declared the moment a clear turning point in the history of computer vision. People asked how AlexNet had pulled off such a miracle. The answer was straightforward: deep neural networks and GPUs. A deep learning theory that had been dormant for decades had finally opened its eyes, powered by NVIDIA's graphics card as its engine.

When Jensen Huang heard the news, a jolt ran through him.

He did not take it as just another contest result. What he saw was a fundamental shift in computing. He sensed that AlexNet was not merely an image recognition tool. It was something that, given input and output data, could learn the complex relationship between them on its own, without a single line of human-written code.

A universal learning machine had been born, one that could learn any pattern in the world as long as it had enough data.

He reached his conclusion immediately. This is the future.

On Friday evening, he sent an email to every employee. We are no longer a graphics company. We are now an artificial intelligence company.

By Monday morning, NVIDIA had been reborn as an artificial intelligence company.

Only the Prepared Can Seize the Moment

They say fortune smiles only on those who are prepared.

AlexNet's success may look like a stroke of luck, but behind it lay nearly a decade of Jensen Huang's sweat and tears. As described in earlier chapters, Huang bet the company's future when he unveiled CUDA in 2006.

CUDA was a programming tool that let developers use the GPU, a graphics chip, for general-purpose computing. Huang believed GPUs could do far more than push pixels on a screen; they could power scientific simulations and complex mathematical calculations. At the time, nobody agreed.

Wall Street investors scoffed.

Criticism poured in. Why was a gaming company dumping enormous sums into developing a supercomputing language that only scientists would ever touch?

CUDA was free, yet almost nobody used it.

Even inside NVIDIA, people called it the "CUDA tax," treating it as a money pit that ate into the company's profits. The stock price nosedived, plunging 70% in 2007, the year after CUDA's launch.

Every time Huang met with investors, he faced the same question: How on earth will you make money from this? His answer: I don't know yet. But I believe this will change the future of computing.

Still, Huang did not waver.

He made a lonely decision as a leader. He believed the company had to lay the groundwork for the future of computing, even if it meant sacrificing short-term profits. He donated NVIDIA GPUs to universities and research labs around the world and taught them CUDA. With the patience of a farmer scattering seeds, he waited for someone, someday, to make something bloom on that platform.

He often said: If we don't build it, they can never come.

He chose to build where there were no customers, because where there are no customers, there are no competitors either. He held on for ten years, sustained by a single conviction: that one day, this technology would solve problems everyone else considered impossible.

In his own words, it was a decade of pain and suffering.

Huang had observed that Moore's Law was hitting its limits. Transistor counts kept climbing, but gains in performance and power efficiency were slowing down. He concluded that the old CPU-centric

approach could not keep up with the computing demands of the future. The only answer was accelerated computing, processing tasks in parallel.

So it was no accident that the AlexNet team was able to pick up a GTX 580. Huang had already seeded hundreds of millions of CUDA-capable GPUs across the globe. Professor Hinton's students could not afford to wait in line for time on a national supercomputer. Instead, they rewrote the history of artificial intelligence with graphics cards bought at a local store.

What if the CUDA ecosystem had never existed? Researchers would still have been wrestling with arcane hardware languages, unable to even attempt implementing their breakthrough algorithms. Documentation existed. Sample code existed. A developer community existed. That is why they were able to learn CUDA.

Huang saw this moment as a victory for those who had prepared.

He often said he builds markets where no market exists. This was not bravado. Long before the enormous market called artificial intelligence existed, he had been building and refining the tools that market would need. The Big Bang of 2012 was not a sudden explosion. It was compressed energy that had been building for years, finally crossing a critical threshold and bursting into the open.

Think of the Roman general Caesar crossing the Rubicon. He never looked back.

Huang was the same. The moment he confirmed the potential of artificial intelligence, he made a bold decision to redirect all of the company's resources toward AI computing.

Beyond the existing graphics card lineup, he launched development of data center chips designed exclusively for AI workloads.

It was a declaration: NVIDIA would shed its identity as a gaming company and be reborn as an AI computing company.

The employees were stunned, but Huang aligned the entire organization around a single goal. He knew. If they failed to seize this opportunity now, NVIDIA would remain a component supplier forever.

Years later, he recalled: It took almost twenty years to build the CUDA platform and transform it into the computing platform it is today. Five years, ten years, twenty years; those timescales don't feel that long to me.

A decade of solitary investment and preparation collided with the AI Big Bang of 2012, and NVIDIA launched onto a trajectory of explosive growth. The reasoned hope Huang had carried for so long had finally become reality.

The Man Who Hand-Delivered the DGX-1 to OpenAI

The year was 2016.

AI research had moved beyond university labs and become a battlefield where Silicon Valley's giants competed for dominance. Google, Facebook, and Amazon were vying for supremacy in artificial intelligence. At the center of it all stood OpenAI, the nonprofit research lab founded by Elon Musk and Sam Altman.

OpenAI had a grand ambition: to build safe artificial intelligence that would benefit humanity.

But its resources were meager. It was a tiny startup, a handful of researchers working out of a shabby office in San Francisco. They needed enormous computing power, yet Google and Facebook had locked up thousands of GPUs, leaving OpenAI desperately short of the resources it needed to conduct research.

Jensen Huang saw exactly what they needed.

NVIDIA was developing what could be called the world's first AI supercomputer, the DGX-1. The machine linked eight cutting-edge GPUs at high speed to deliver 170 teraflops of computing power. It weighed 32 kilograms and contained more than 35,000 components. It was the crystallization of NVIDIA's technology, with over $2 billion in development costs poured into it.

When he first unveiled the DGX-1 on stage at the GTC conference, the response was cold. Jensen Huang recalls the moment: "There was no applause, no interest, just 100 percent confusion. People just asked, 'What is this? Does it run Windows?'"

The $129,000 price tag, its massive size, and its unfamiliar purpose meant nobody placed an order. Then one hand went up.

It was Elon Musk. He told Jensen Huang, "We can use that." Jensen Huang laughs as he remembers: "My first customer was a nonprofit. When I heard that, I felt the color drain from my face."

But Jensen Huang didn't hesitate.

He loaded the first DGX-1 off the production line into his own car and drove to OpenAI's office in San Francisco. He didn't send a delivery driver. He wanted to hand it over personally, not as a courier, but as the creator of the technology, delivering the tool directly to the innovators who would put it to its best use.

He compared himself to a DoorDash delivery driver. "I picked up a box and drove all the way to San Francisco. Then I handed it to Elon myself."

August 2016, OpenAI's office.

Elon Musk, Sam Altman, Greg Brockman, and Ilya Sutskever were waiting.

Ilya Sutskever was the same young man who had built AlexNet four years earlier and triggered the Big Bang of artificial intelligence. Now he had become a core researcher at OpenAI.

Jensen Huang showed up in his leather jacket. In his hand was a silver marker. He wrote on the case of the DGX-1:

"To Elon and the OpenAI team. For the future of computing and humanity. I present the world's first DGX-1."

The photo from that moment became an iconic image in the history of technology. Jensen Huang in his leather jacket, signing the machine. Elon Musk standing with his arms crossed, looking on with satisfaction. The researchers gathered around them. Everyone looked as happy as children opening Christmas presents.

This delivery was not just a shipment of equipment.

If computing in the past was a calculator that faithfully executed commands typed in by humans, then from this moment on, computing would serve as the midwife for the birth of intelligent beings that could learn and reason on their own. Jensen Huang didn't just sell a machine. He gave a gift of time to the people researching the future.

Thanks to the DGX-1, the pace of research at OpenAI accelerated dramatically.

According to Ilya Sutskever, research experiments moved weeks ahead of schedule. Experiments that had been impossible before became possible. Training times shrank from months to days.

The researchers used the machine to train early language models. A researcher named Andrej Karpathy said in a video at the time: "We plan to use the DGX for large-scale language models. Someday we will be able to talk to computers the way we talk to people."

His prediction came true. Six years later, in 2022, OpenAI released ChatGPT to the world. The era of generative AI had arrived, sweeping the globe. It was technology conceived in the very supercomputer Jensen Huang had delivered in 2016.

Jensen Huang remembers that shabby office.

"A few geniuses were gathered in a small room. They were dreaming of the future, and I wanted to give them the tools they needed."

He didn't see his customers as mere consumers buying a product. He saw them as allies building a new industry together, and as pioneers proving just how far NVIDIA's technology could reach.

In a later interview, he spoke about the thrill of watching the tools they had built end up in the hands of geniuses who would change the world.

The DGX-1 was an expensive piece of equipment at $129,000. But Jensen Huang was looking through it at a future ten or twenty years away. He wasn't just selling hardware. He was planting the concept of an AI factory in the minds of the world.

In his own words: "We wanted to shrink a supercomputer down to a single chip. And now we are making an entire data center operate like one chip."

The story continues.

In October 2025, Jensen Huang personally delivered another machine to Elon Musk. This time, the destination was Starbase, the rocket launch site in Texas. Standing next to the largest rocket in the world, he handed over the DGX Spark, the smallest supercomputer in the world.

He cracked a joke. Delivering the smallest supercomputer right next to the biggest rocket.

Nine years separate the 2016 delivery from the 2025 delivery.

2016 was an investment in the future. 2025 was the portrait of a reigning empire. NVIDIA's market capitalization had surpassed $4 trillion, and every major company on the planet was lining up to get its hands on NVIDIA chips.

But the same philosophy connects those two moments. If a developer reaches out and says they need GPUs, the answer must always be yes.

The discovery of AlexNet in 2012 and the partnership with OpenAI in 2016 deserve to be called NVIDIA's Big Bang moments.

They are a perfect example of the explosive growth that becomes possible when preparation meets opportunity.

Jensen Huang never stopped investing in the future, even as Wall Street laughed at him. Those years of patience eventually carried him from the fringes of Silicon Valley to the center of the AI revolution.

He was no longer the head of a company that made graphics cards for gaming PCs. He had become the architect laying the foundation for the coming age of artificial intelligence, the supplier of the engine powering a new industrial revolution.

He puts it this way: "We spent 30 years preparing for this moment. We believed that accelerated computing would one day change everything, and that single belief carried us through years of pain. Now, that future has arrived."

History was being rewritten at the fingertips of one man who had been ready all along.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 8. The Black Leather Jacket

Part 1. Building an Empire

← Previous Chapter

See every chapter

Next Chapter →

A Monk in the Closet, a Commander on the Battlefield

When people see Jensen Huang, they think of the black leather jacket.

Under Silicon Valley's scorching sun, on stages flooded with bright lights, in the humid night markets of Taiwan, he always looks the same. Like a Roman legionnaire strapping on armor before heading into battle, he puts on that jacket every morning and walks out the door.

Steve Jobs had his black turtleneck. Jensen Huang has his black leather jacket.

If Jobs's turtleneck was an aesthetic of simplicity, Huang's leather jacket is something closer to combat gear. It's an unspoken declaration that he's ready to rush to the front lines at any moment, a sign of resolve that he won't step back from the cutting edge of a fast-moving technology war.

Some people ask: why does he always wear the same thing? The answer is straightforward. For him, time spent choosing clothes is time wasted. He doesn't wear a wristwatch either. The only time that matters to him is right now, this moment. What concerns him isn't what he wears but what he decides.

We tend to picture the emperors who built great empires sitting on high thrones in lavish robes, issuing commands from above.

Jensen Huang is different. He doesn't even have a fixed office. He roams NVIDIA headquarters and sets up wherever he's needed. He'll sit in the corner of a conference room huddling with engineers, or stop a new hire in the hallway and fire off a question.

Inside that black jacket lies a distinctive and intense governing philosophy that drives the vast empire called NVIDIA. It's nothing like what traditional management textbooks teach. It is a set of survival rules, both wild and deeply rational.

Sixty Direct Reports and the Beauty of Flatness

Management professors say a single manager can effectively oversee somewhere between six and ten people. Beyond that, communication breaks down and oversight becomes difficult. This has been a tested principle for a long time.

Jensen Huang, as if mocking that rule, has roughly sixty direct reports. Vice presidents, key engineers, all of them report straight to him. Management scholars call this impossible. One person

listening deeply to sixty people and evaluating their work is, physically, extremely hard to do.

But for Jensen Huang, this isn't a choice. It's a necessity.

He despises information getting distorted as it passes through layers. Just as a messenger crossing mountains and rivers subtly changes the battlefield report along the way, multiple tiers inside a company dilute the raw truth.

He puts it this way: "A report is already filtered information. It carries bias; interpretation has been added. I want the raw, unprocessed information straight from the field."

The sixty-person direct reporting structure connects the massive organism of NVIDIA like a single neural network. The Roman Empire built roads so that news from every corner could reach the capital quickly. Jensen Huang eliminated organizational layers so information could flow without distortion.

In a traditional pyramid organization, a story from the field passes through several levels before it reaches the CEO. An engineer's concern goes to a team lead. The team lead reports it to a department head. The department head passes it to a division chief, who then briefs the CEO. Along the way, the story changes. Some parts get exaggerated, others get minimized. Bad news gets wrapped in soft language. Good news gets inflated.

Jensen Huang skips that entire process. He talks directly with the commanders in the field.

A natural question follows: how on earth do you manage sixty people? If he held a weekly one-on-one with each of them, he'd have no time left for anything else.

The answer is straightforward. He doesn't do one-on-ones.

Jensen Huang shares all information in open meetings. He says: "Almost everything I say, I say to everyone at the same time." Giving exclusive information to any single person breeds internal politics and creates power imbalances.

This means he doesn't treat information as power. In some organizations, there are secrets only the CEO knows, meetings only senior executives attend. The people who hold that information gain an upper hand. Jensen Huang hates that kind of structure. He believes information is not power; it is raw material we need to get the work done.

That's why meetings at NVIDIA look unusual.

Dozens of people attend the sessions where major decisions are made. A new hire might be sitting right next to a vice president. Everyone watches how Jensen Huang works through a problem. They all see how he thinks, what logic leads him to a conclusion, and they all learn from it.

He also has a distinctive communication method. NVIDIA employees write down the five most important things they're working on and send it by email. Every morning, the moment Jensen Huang opens his eyes, he reads dozens, sometimes hundreds, of these emails. He isn't reading polished

reports. He's reading his people's thoughts and concerns firsthand.

He says: "I want to know what 30,000 employees are thinking, what they're prioritizing."

The advantage of this flat structure is agility.

Decisions happen fast. There's no room for middle managers to distort or delay information. Because all sixty leaders share the same context, different departments rarely drift in different directions.

There are downsides, of course. Managing sixty people at once is mentally exhausting. Tracking each person's performance in detail becomes nearly impossible. But Jensen Huang overcomes those drawbacks with transparency and trust.

He doesn't do career coaching with his direct reports. They're already experts in their respective fields. What they need isn't advice; it's information and context about where the company is headed. Jensen Huang defines the CEO's role not as a manager but as a provider of information, a strategic guide.

The board once made a suggestion: "How about bringing in a chief operating officer to reduce the administrative burden?" Jensen Huang refused outright. "Direct communication," he replied, "is the best way to make sure everyone knows what's going on."

At NVIDIA, nobody uses the words "division" or "business unit." Jensen Huang doesn't want the company carved into pieces. He runs NVIDIA like one enormous team. The way billions of neurons in a human brain are interconnected yet produce a single thought, tens of thousands of NVIDIA employees move toward a single goal.

Jensen Huang trusts young talent.

He'll hand important projects to new hires. Fresh perspective matters more to him than years of experience. The philosophy he absorbed while scrubbing dishes at Denny's, that no task is beneath anyone, lives in that decision.

His sixty direct reports are not courtiers serving an authoritarian emperor.

They are sixty antennas. Sensory organs that detect shifts in the world and relay those signals to Jensen Huang. And Jensen Huang is the helmsman who synthesizes those signals and sets the ship's course.

Share the Failures, Split the Successes

Look at Roman history. Defeated generals were sometimes executed as punishment. But in NVIDIA's court, failure itself is not a crime. The only crime is hiding it.

Jensen Huang understands, better than anyone, that an organization afraid of failure can never innovate. He was a survivor of those thirty days in NVIDIA's early years when the company teetered on the edge of bankruptcy. When the deal with Sega nearly fell apart,

when the NV1 failed, when the company's bank balance held just one month's payroll, he learned the taste of failure down to the bone.

That's why he tells his employees often:

"You failed? Then you need to tell everyone what you learned from that failure."

This happened in the early 2000s. NVIDIA released a graphics card with a problem. The fan was unbearably loud. Most companies would have disciplined the people responsible. Jensen Huang took a different approach.

He had the project team give a presentation. In front of hundreds of employees.

The team members explained every decision they had made. What choices they went with, why they made those choices, and where exactly things went wrong. This wasn't meant to humiliate them. It was meant to let the entire organization learn from that failure.

NVIDIA even produced a parody video starring the failed graphics card. The premise: using the noisy card as a leaf blower. The team members appeared in it themselves. The video was distributed to the press. It was a declaration that failure would not be hidden.

Jensen Huang dislikes one-on-one meetings. Reprimands or praise delivered behind closed doors breed information asymmetry. He discusses problems in open settings instead. If someone made a mistake, he frames it not as personal incompetence but as evidence that the challenge the team took on was that difficult.

He says:

"Feedback is learning. Why should you be the only one who gets this valuable lesson? All of us should learn from that mistake."

He believes the best learning comes from other people's mistakes. Why should you only learn from your own errors? Why should you only learn from your own embarrassment? You can watch someone else's mistake and learn from it.

This puts enormous pressure on employees, too. Having your mistakes discussed publicly is never a pleasant experience. But at the same time, a culture takes root where mistakes are not grounds for discipline but opportunities for learning.

Jensen Huang says: "I don't fire anyone."

He hates giving up on people. He puts it this way: "I cleaned bathrooms, and now I'm the CEO of a company. Anyone can learn. Given the chance."

He continues: "I'm always watching sixty brilliant people do their work. They probably don't realize it. But I'm constantly learning from every one of them. That's why I hate giving up on people. I believe they can grow."

He says it like a joke: "People know. I'd rather torture them into greatness."

That is the moment when one person's failure becomes intellectual capital for the entire organization. It gives employees a profound sense of psychological safety. If I stumble while taking on a challenge, the boss won't abandon me. He'll study why I fell and help make sure my colleagues don't fall the same way.

The fruits of success, on the other hand, he shares thoroughly. When NVIDIA's stock price soared, Jensen Huang was wary of letting the gains concentrate in his own hands. NVIDIA employees can connect the company's growth to their own wealth through stock compensation.

Every year, he personally reviews the compensation of all 42,000 employees. In almost every case, he increases the company's spending to take care of his people. Reports have emerged that more than half of NVIDIA's employees are millionaires.

Jensen Huang says, 'I've created more billionaires among our executive team than any other CEO in the world.' This isn't a boast. It's a philosophy. He believes employees can only unleash genuine creative passion, rather than working just to survive, when they've achieved financial freedom. Take care of people, and the rest takes care of itself.

Even when the company faced hard times, he tried to keep layoffs to a minimum.

His view is that when a ship hits a storm, you don't throw the rowers overboard. Pain is something a leader shoulders first; glory is something a leader passes to the team first. That is the essence of how he leads tens of thousands of brilliant minds.

Intellectual Honesty: The Ultimate Weapon

One of the values Jensen Huang holds dear is intellectual honesty. This goes beyond the moral dimension of not telling lies. It means clearly distinguishing what you know from what you don't, and immediately acknowledging and correcting course when circumstances change or your judgment turns out to be wrong.

He puts it this way: 'Without intellectual honesty, you can't have a culture that tolerates failure. Because people become too attached to their ideas. Even when an idea is bad or doesn't work, if they feel their reputation is tied to it, they can't admit failure.'

Many leaders cling to bad decisions to protect their authority, or because they can't stomach the money already spent. Economists call this the sunk cost fallacy: continuing to push forward with a wrong decision because of costs already incurred.

Jensen Huang, though, is ruthless about this.

The most dramatic example of intellectual honesty dates back to NVIDIA's early days and the failed project with Sega. NVIDIA had chosen the wrong architecture, and the company was on the brink of bankruptcy. Jensen Huang went to see Sega's CEO.

He said, 'We were wrong. If we continue down this path, we will fail, and we'll waste your money. Cancel the contract and find another partner.'

What came next was the remarkable part. He added, 'But we need the remaining payment you promised us to avoid going bankrupt.'

It was an almost brazen confession of honesty. Admitting he was wrong while simultaneously asking for help. Most people would have been too embarrassed to say it. Jensen Huang said it anyway.

His disarming frankness moved Sega's CEO. Sega canceled the contract but paid the remaining balance. NVIDIA used that money to pull off a miraculous recovery.

Jensen Huang took away a lifelong lesson from that moment: confronting your failure head-on and admitting it honestly is the first step toward overcoming a crisis.

During the explosive growth of the smartphone market, NVIDIA jumped into the mobile chip business too. They developed a processor called Tegra. But they struggled against competitors like Qualcomm and Samsung. An ordinary company would have poured massive marketing dollars into the fight, if only out of pride.

Jensen Huang, though, stood up one day and told his employees, 'We can't win here. This is not our battlefield. We're pulling out.'

He withdrew from the mobile business without hesitation. It was a project they'd poured countless hours and dollars into, but he felt no attachment. Instead, he steered the ship toward territories nobody was paying attention to at the time: artificial intelligence and autonomous driving.

He told his employees plainly, 'We can't win there. But there's somewhere else where our technology is truly needed.'

If he had let pride keep him slugging it out in the mobile market's mud fight, the AI empire that NVIDIA is today would never have existed.

He says, 'A lot of people say a CEO should always be right and never change their mind. That makes absolutely no sense to me.'

Jeff Bezos of Amazon has said something similar: 'People who are right a lot change their minds a lot.' In a fast-moving industry, clinging to a bad idea is worse than having no idea at all.

Jensen Huang doesn't hesitate to admit when his decisions as CEO have been wrong. Even a strategy he announced in front of 50,000 employees, he'll reverse immediately if new information or a shift in the environment tells him it was a mistake.

He says, 'A leader's job is not to be right. A leader's job is to help us succeed.'

He came to realize that when he says 'I was wrong,' employees don't lose faith. They feel relief that he sees the truth, and they trust him even more.

This intellectual honesty springs from first principles thinking. He rejects the logic of 'everyone else is doing it, so we should too.' Instead, he digs into the essence of things the way a physicist would.

What is the fundamental principle behind this technology? Are we the only company that can solve this problem? Is this something the future of computing absolutely requires?

He asks questions relentlessly. If the answer is no, he walks away without hesitation. If the answer is yes, he will wait ten or twenty years to create a market, even when none exists.

The patience that carried him through a decade of Wall Street mockery while developing the CUDA platform came from this intellectual honesty. He knew the logic he had built was mathematically and physically correct, so he never wavered.

Intellectual honesty cuts the cost of communication inside an organization. At NVIDIA, reports written to flatter a boss or excuses crafted to hide failure do not work. Only data and logic hold currency.

Even a new hire can challenge the CEO's opinion with sound reasoning. Jensen Huang is pleased when someone proves him wrong. It means the company has just corrected its course.

He knows.

In the world of technology, the truly dangerous enemy is not a competitor but the arrogance of lying to yourself. That is why, every morning, he looks in the mirror and then faces his sixty direct reports and asks himself: Are we being honest right now? Are we seeing the truth?

Approaching seventy, he still says, "I am learning." He reads new papers, poses questions to ChatGPT, and stretches the boundaries of his own cognition. For him, intellectual honesty is not merely refraining from lies; it is the attitude of admitting what you do not know and never stopping to learn.

The man in the black leather jacket wakes every morning carrying the anxiety that his company could collapse. The thought that bankruptcy might be thirty days away. Yet that anxiety does not shrink him; it keeps him honest. He refuses to look away from harsh reality, drunk on success. He stares at things exactly as they are.

Communicating with sixty direct reports, listening to voices from the front lines, sharing failures rather than hiding them, and above all admitting his own errors. This intellectual honesty is the force that drives Jensen Huang forward into the future at the helm of a trillion-dollar company.

His black leather jacket is not just clothing. It is a declaration that he rejects pretense and focuses on substance, a warrior's symbol of wearing pain like armor and charging at the impossible. It is also a shield that guards him against the temptation to settle or compromise with himself.

The world may call him an emperor, but the man in the leather jacket is still hungry, still brimming with curiosity, still a challenger who has never lost the humility and urgency of the days he washed dishes at Denny's.

This philosopher in the black leather jacket still roams the offices of Silicon Valley, asking his people: "What do you think?"

His question is not a command but an invitation. A courteous yet powerful invitation to move together into the unknown, armed with nothing but truth and unafraid of failure.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 9. The Age of AI Factories

Part 2. Architect of a New Age

← Previous Chapter

See every chapter

Next Chapter →

Turning points in history always arrive quietly.

It was true when the steam engine remade the world, and again when electricity banished the night. People never grasp, at first, how large the change really is. They see a clever invention, nothing more. Then one day they glance over their shoulders and find that the world has already become something else entirely.

What Jensen Huang is doing right now is exactly that kind of change.

He is redefining what the word "computer" means, from the ground up. The computer we knew was a box sitting on a desk. You tapped the keyboard and letters appeared on the screen; you moved the mouse and an arrow followed it around. That kind of machine.

But that is not what a computer is to Jensen Huang. To him, a computer is a building the size of a city block. It is a factory that runs without stopping, venting heat instead of smoke from its stacks.

The Data Center Is No Longer a Warehouse

Have you heard of a place called a data center? Chances are you have never visited one. They tend to sit where ordinary people never look.

On a sprawling lot at the edge of a city, or in the middle of a desert, a massive building stands with no sign on its wall. There are almost no windows. Inside, thousands, tens of thousands of computers line up in rows, humming.

For the past thirty years, data centers did one thing. They stored information.

Photos you took on your smartphone, messages you sent to friends, videos you watched online: all of it piled up neatly somewhere inside a data center. Think of it as an enormous library. When someone asked for a piece of information, the center pulled it out and showed it; when new information arrived, it slotted into an empty space. That was the entire job of a data center.

Jensen Huang calls this retrieval-based computing. It means a system that finds information someone already created. It works like borrowing a book from a library. The book was written by somebody else long ago, and all you do is locate it on the shelf and carry it home.

Then, in May 2025, at an event in Taiwan called Computex, Jensen Huang made a declaration to the world.

"The data center is no longer a warehouse. It is a factory."

To understand what that single sentence means, you have to think about the difference between a warehouse and a factory.

A warehouse is a place that holds goods. Items come in, get sorted and stacked, and wait until someone needs them. Nothing new is created inside a warehouse. What goes in comes back out unchanged.

A factory is different. Raw materials go into a factory. Iron ore goes in. Plastic pellets go in. All sorts of components go in. Then something happens inside. Machines spin, sparks fly, workers move. And out the other end comes something completely different. Cars come out. Refrigerators come out. Smartphones come out. What enters and what exits bear no resemblance to each other. A factory is a place that makes things.

This is why Jensen Huang called the data center a factory. Data centers no longer pull stored information off a shelf and hand it over. They produce new information.

When you ask ChatGPT a question, the answer was not written down somewhere in advance. It is generated fresh at that very moment. It is like a chef who takes your order and cooks the dish right there on the spot.

Let me use an analogy to show how enormous this shift is.

Imagine walking into a library and asking the librarian, "Tell me about King Sejong of the Joseon Dynasty." The librarian would walk between the shelves, find the right books, and bring them to you. Those books would contain stories about King Sejong that someone had already compiled. The librarian hands them over. That is what data centers used to do.

Now it works differently.

When you ask an AI the same question, the AI does not go looking for a book. Instead, it writes a new account of King Sejong right there on the spot. The account draws on knowledge the AI absorbed during training. But every sentence is being composed for the first time in that instant, shaped to fit your question, arranged in a way you can follow easily.

Jensen Huang calls this generative computing. Instead of retrieving what already exists, the system creates something new.

This shift has changed even the physical appearance of data centers. The old data center was a quiet place. Cool air drifted through the aisles. Blinking lights were the only signs of activity. Storing information does not take much force.

An AI factory is nothing like that.

Tens of thousands of GPUs run without rest. They throw off tremendous heat. Air alone cannot cool them, so chilled liquid flows through pipes to draw the heat away from the chips, much the way blood circulates through a human body to carry heat where it needs to go. If the old data center was a hushed library, the AI factory is a blast furnace burning white-hot.

The chip Jensen Huang developed called Blackwell, and the one coming after it called Rubin, are the engines that power this factory. These chips do not work alone. Thousands, tens of thousands of them link together and operate as if they were a single enormous brain. Jensen Huang puts it this way:

"The entire data center is now one computer factory."

Put Electricity In, Get Intelligence Out

You've probably heard the term "industrial revolution." It was a massive transformation that began in England roughly two hundred and fifty years ago. Before that, humans relied on their own strength or the strength of animals to get work done. Oxen pulled plows through fields. Horses hauled carts loaded with goods. What a person's arms and legs could accomplish had clear limits.

Then came the steam engine. Burn coal, boil water, and steam comes out. That steam pushes a piston, and the piston drives a machine. A force far stronger than human muscle appeared, one that never tired. Spinning machines began turning in factories. Steamships crossed oceans without waiting for wind. The world changed.

After that came electricity. A machine called a generator converted the power of water or coal into electrical current. Electricity is invisible. You can't touch it. But it travels along wires to anywhere you need it. Factories, homes, every corner of every city. Night no longer meant darkness, and machines ran with new precision.

Jensen Huang calls what is happening right now a new industrial revolution. He explains its core like this.

"In the past, we boiled water to make steam, and used a machine called a dynamo to make electricity. It was invisible, but enormously valuable. Now we have a new dynamo. You put electricity in, and intelligence comes out." This isn't a metaphor. It is what is actually happening.

Electricity flows into an AI factory. That electricity powers tens of thousands of GPUs. Inside those GPUs, trillions of calculations take place. And on the other side of the factory, something called tokens come out.

What is a token? Think of it as the smallest piece of information that an AI produces. When writing text, a single word is a token.

When generating an image, a single pixel is a token. Gather enough tokens and they become sentences, then paragraphs, then entire books. Or they become images, videos, the movements of a robot.

This is what an AI factory looks like through Jensen Huang's eyes. Electricity goes in one side; intelligence comes out the other. Just as the steam engine turned coal into mechanical power, just as a power plant turned the force of water into electricity, an AI factory turns electricity into intelligence.

Consider why this matters.

Until now, intelligence came from only one place: the human mind.

A skilled doctor diagnosed illness. A talented lawyer solved legal problems. A creative artist produced works of art.

That kind of intelligence required years of study and experience. It was scarce. It was expensive. Not everyone could have it.

But what if intelligence could be manufactured in a factory?

Imagine it rolling off the line the way cars roll out of an automobile plant. If that happened, the price of intelligence would drop. Anyone could afford to use it.

Much like electricity itself. When it first appeared, it was rare and costly. Today you just plug something into a wall outlet and it's there for everyone.

That is the world Jensen Huang dreams of.

A world where intelligence is as common and as ordinary as water from a faucet or light from a switch.

Running these factories, of course, takes an enormous amount of electricity.

AI factories are hungry for power. A single facility can consume as much electricity as a small city. Jensen Huang knows this. That is why he is partnering with energy companies. Nuclear plants, solar farms, wind farms. He is working with every source of clean energy he can find.

In December 2025, Jensen Huang issued a warning at a research institute in Washington, D.C.

While the United States takes three years to build a data center, he said, China can put up a hospital in a single week. Winning the AI race takes more than good chips. You need factories to run those chips, and you need to build them fast. You also need power plants to supply those factories with electricity.

The world as he sees it looks like a five-layer cake. At the bottom sits energy. Above that, chips. Above that, data centers. Above that, AI models. And at the very top, applications. All five layers must be in place for an AI factory to run properly. Jensen Huang is drawing the blueprint that ties every

layer together.

The Engineer Driving a $100 Trillion Industrial Revolution

Every time Jensen Huang steps onto a stage, he throws out staggering numbers.

"This is a $100 trillion opportunity."

One hundred trillion dollars. It's hard to picture how large that is. Add up everything every country on Earth earns in a single year and you get roughly $100 trillion. That's the size of the entire global economy. What Jensen Huang is saying is that AI will reshape all of it.

The IT industry, up to this point, occupied only a fraction of the global economy. It was a market worth roughly three trillion dollars. Building computers, selling software, providing internet services. That was what the IT industry did. A large industry, sure, but compared to manufacturing, agriculture, healthcare, and finance, it was just one piece of the picture.

Jensen Huang sees AI breaking past the boundaries of IT. Because AI is not a tool. AI is the work itself.

Let me explain what that means.

There is a program called Excel. It calculates numbers and creates tables. It is a tool. A person opens Excel, types in numbers, enters formulas, and gets results. The person does the work; Excel helps.

AI is different.

Tell an AI, "Analyze this data and write me a report," and the AI writes the report itself. No one needs to stand over it giving step-by-step instructions. The AI reads the data on its own, finds the meaning, and produces sentences. The AI is doing the work. The person just reviews the output.

Jensen Huang calls this a digital workforce.

A non-human intelligence performing labor. This digital workforce can be deployed in manufacturing, in healthcare, in finance, in transportation, in every industry. That is why the market for AI is not the three trillion dollars of IT but the one hundred trillion dollars of the entire world economy.

Jensen Huang's gaze reaches beyond the computer screen. What he sees is the physical world. Robots turning on factory floors, cars running on roads, machines hauling goods through warehouses. All of these can become the body of AI.

Until now, AI stayed inside the screen.

Writing text, drawing images, carrying on conversations. All of that happened within the digital world. But what Jensen Huang is preparing for is AI stepping out of the screen and into the real world. He calls this physical AI.

To prepare for physical AI, Jensen Huang built something called Omniverse. Omniverse is a virtual replica of the real world. You can build a factory virtually. You can move a robot virtually. You can drive a car on a virtual road.

Why is this necessary? A robot learning to walk has to fall down countless times. A car learning to drive itself has to log hundreds of millions of kilometers. Practicing that much in the real world takes too long and is too dangerous.

So Jensen Huang built a virtual world. Inside Omniverse, a robot falls millions of times and learns to walk. A car drives billions of kilometers on virtual roads and learns to steer. Only after enough training does it come out into the real world.

Throughout 2025, Jensen Huang traveled the globe. Taiwan, Japan, South Korea, India, France, Saudi Arabia, the United States. Everywhere he went, he delivered the same message: every country needs its own AI factory. He calls this sovereign AI.

Each country has its own language. Its own culture. Its own history. An AI that understands these things must be trained on that country's data and built in that country's factories. Depending on AI made by a foreign company could become as risky as depending on another country for your food supply.

Jensen Huang came to South Korea, too. In November 2025, at the APEC summit held in Gyeongju, he announced a plan with the Korean government and its major corporations to install more than 260,000 GPUs in the country. Samsung, SK, Hyundai, Naver. Korea's biggest

companies are all on board. The goal is to build AI that understands the Korean language and Korean culture, right on Korean soil.

Jensen Huang has a clear reason for doing all of this. He controls the infrastructure that powers the global economy. Think of the Romans laying roads across their empire. Roman roads were the corridors through which armies marched, merchants traveled, and information flowed. By controlling the roads, Rome could rule the Mediterranean world.

Jensen Huang is laying the roads of the twenty-first century. Those roads are chips made of GPUs, buildings called data centers, and software called CUDA. On these roads, the wagons of artificial intelligence roll forward. AI for healthcare, AI for finance, AI for manufacturing. Every AI runs on the roads Jensen Huang has paved.

He is not just someone who makes and sells chips. He is an architect drafting the blueprint for a new era. If there is a massive train carrying the hundred-trillion-dollar world economy into transformation, Jensen Huang is sitting in the engineer's seat.

Yet he does not let himself grow complacent. There is a phrase he repeats like a reflex.

"We are always thirty days from going out of business."

The world of technology changes at a terrifying pace. Yesterday's winner can become tomorrow's loser. Jensen Huang knows this better than anyone. So he never stops. Past sixty, he still puts on the black leather jacket and bounds across the stage. He talks about the future with more intensity than people half his age.

No one knows exactly where the intelligence produced in AI factories will take humanity. When the steam engine first appeared, people had no idea how it would reshape the world. The same was true when electricity arrived. New technology always transforms the world in directions no one predicted.

One thing, though, is certain. Jensen Huang stands in the front row of that change. The AI factories he designed are humming at this very moment. New AI factories are going up in every corner of the globe. Electricity goes in; intelligence comes out.

Jensen Huang asks us a question.

"Are you ready for this new age of intelligence?"

His question carries weight. We stand now before a massive door he has opened. The door is already open. Beyond it lies a world none of us have experienced before. History has always favored those who move forward. It is unforgiving to those who hesitate behind. Jensen Huang has already crossed that threshold and is running far ahead.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 10. The Integration of OS, Programs, and Intelligence

Part 2. Architect of a New Age

← Previous Chapter

See every chapter

Next Chapter →

When Rome built its roads, people thought of them as paths. They were not paths. Rome's roads were the nervous system of an empire. Legions marched on them, merchants hauled goods across them, and news traveled through them. Without those roads, the Roman Empire could not have existed.

The same is true of artificial intelligence as Jensen Huang describes it. It is not a technology. It is the road system of a new era.

June 2025, on stage at London Tech Week.

Jensen Huang stood next to the British Prime Minister. He wore his black leather jacket, as always.

He turned to the audience and said:

"There is a new programming language. The name of that language is 'human.'"

Human.

He was not joking.

When computers first entered the world, operating them required specialized knowledge. There was something called machine language. A language made of zeros and ones. Computers understood nothing else.

Over time, programming languages appeared. Fortran, COBOL, C, Python. They were closer to human speech than machine code, but they were still hard to learn. You had to study for years in college. The world was divided between people who could program and people who could not. That boundary looked like a tall wall.

Jensen Huang says that wall is coming down.

Think about it.

If you want a computer to do something for you, what do you have to do now? In the past, you had to write code.

Declare variables, define functions, run loops. If you didn't know these things, you couldn't make a computer do anything.

Now it's different.

You just talk to it.

"Organize this table for me." "Summarize this document." "Schedule my meeting for tomorrow."

The AI understands and executes.

In Jensen Huang's words, the way you ask a computer to do something has become the same as the way you ask a person.

"Just ask politely."

Not many people know C++. Fewer know Python. But everyone knows how to speak. Korean, English, Spanish, it doesn't matter. The language you already use has become a programming language.

But we need to go one step further. The future Jensen Huang envisions is not just "a computer you talk to." What he sees runs deeper.

There is something called an operating system. Windows, macOS, Linux. It is the first program that runs when you turn on a computer. Every other program runs on top of it. Word, Excel, games, all of them. Without an operating system, a computer is an empty shell.

Jensen Huang says artificial intelligence will become the new operating system.

Let me explain what that means.

Right now, we "launch" programs. If you want to use Word, you click the Word icon. If you want to use Excel, you click the Excel icon. Each program exists separately. We move back and forth between them as we work.

In the future, that won't be the case. The very concept of a program will start to blur.

When you say "Write me a report," the AI will pull together whatever capabilities the task requires in that moment. Text editing, table creation, chart drawing. It will combine them and produce the report. You won't need to think, "I should open Word first." The AI will handle that on its own.

When you say "send an email," you don't need to open the email app. When you say "schedule a meeting," you don't need to open the calendar app. Artificial intelligence acts as a bridge between everything.

In June 2024, at Computex in Taiwan, Jensen Huang declared:

"The era of the software factory is over. The era of the AI factory begins."

What is a software factory? It's the way things have worked until now. Programmers gather and write code. They spend months, years building a program. They sell the finished program to people. That's how Microsoft makes Word, how Adobe makes Photoshop.

The AI factory is different. It doesn't sell pre-built programs. Artificial intelligence creates what's needed on the spot. It doesn't search. It generates.

I mentioned the napkin story in the preface. It's an analogy Jensen Huang likes to use.

A napkin is something you use once and throw away. You blow your nose, wipe your hands, jot a note. When it's served its purpose, it goes in the trash. Nobody treasures a napkin.

That's what software will become in the future.

When you ask artificial intelligence something, it creates a small program in that moment. A program built to answer your question. After showing you the answer, the program vanishes. Like a napkin.

Until now, programs had to be "reusable." You built one and a million people had to be able to use it. That was the only way to recoup development costs. Nobody made a program meant to be used once and discarded. It would have been wasteful.

Artificial intelligence changed that equation. For AI, creating a program is like breathing. It costs almost nothing. It's fine to build a program for just one person, just one time. Use it and toss it like a napkin, no loss.

There's an important distinction here.

Jensen Huang emphasizes the difference between a "tool" and "work."

Excel is a tool. It's like a hammer. A hammer doesn't drive nails on its own. A person picks up the hammer and drives the nail. The same way, Excel doesn't analyze numbers on its own. A person has to open Excel and enter formulas.

Artificial intelligence is different. AI is not a tool. AI is the "work" itself.

When you say "analyze this data," AI chooses its own tools. It decides which analytical method to use. It performs the analysis. It organizes the results and presents them. The person just receives the output.

In Jensen Huang's words, artificial intelligence is "a worker that knows how to use tools."

January 2026, CES in Las Vegas.

Jensen Huang took the stage again. The black leather jacket was the same as always. He told the audience:

"Computing has been fundamentally reshaped. You no longer write software. You train software. You no longer run CPUs. You run GPUs."

Not writing but training. That phrase deserves a closer look.

The old way of making a program went like this. A programmer sat down and wrote code line by line. "If A, then do B. Otherwise, do C." Thousands of lines, tens of thousands of lines of such instructions. The programmer had to think through every possible case in advance.

The way we build artificial intelligence today is different. Instead of writing code, you show it data. "This is a photo of a cat, and this is a photo of a dog. Learn the difference." The AI looks at millions of images and discovers the rules on its own.

The programmer doesn't hand over the rules. The AI discovers them itself. Much like a child learning about the world.

Jensen Huang calls this "training." You don't "write" software; you "train" it. So what does it mean for artificial intelligence to become the operating system?

Think about what an operating system does. You press a key and a letter appears on screen. You move the mouse and the cursor follows. You save a file and it's recorded on the hard drive. The operating system handles all of this. We don't notice the process. We just feel that the computer "works."

In the future, artificial intelligence will play a similar role. When we want something, AI understands it and executes. Which program to use, which data to pull, what order to process things in. AI decides all of it. We won't notice the process. We'll just feel that "things get done."

NVIDIA is already moving in this direction.

In November 2025, Jensen Huang said this during an earnings call.

"NVIDIA's architecture and platform are the only framework that runs every AI model in the world. We are in every cloud."

Everywhere. That was how Jensen Huang put it.

NVIDIA chips sit inside Google's data centers. They sit inside Amazon's servers. They sit inside Microsoft's cloud. Most of the computers running artificial intelligence around the world operate on top of NVIDIA's technology. This is not a coincidence.

Think back to the CUDA story covered in Chapter 6.

In 2006, Jensen Huang introduced CUDA. It was software that let GPUs be used for tasks beyond graphics processing. Wall Street laughed. Who would ever use something like that?

Jensen Huang waited ten years. He grew the CUDA ecosystem slowly. Scientists started using CUDA. Researchers cited CUDA in their papers. Universities began teaching it.

Then in 2012, the age of artificial intelligence opened. An AI called AlexNet delivered a dominant performance in an image recognition competition. That AI had been trained on NVIDIA GPUs and CUDA.

Everything changed from that point. If you wanted to do AI research, you needed NVIDIA. There was no other option. The ecosystem built over ten years had become a barrier to entry.

Jensen Huang is running the same playbook again. This time, it is an AI operating system.

NVIDIA is not a company that sells chips alone. NVIDIA sells systems. It sells platforms where chips, software, and tools are bundled into one package.

There is something called NIM. It stands for NVIDIA Inference Microservices. It is a tool that makes it easy to deploy AI models. There is also something called Omniverse. It is a platform for building and simulating virtual worlds. Robots practice in this virtual world before they go out into the real one.

All of this runs on NVIDIA's chips. It is developed with NVIDIA's software. It moves within NVIDIA's ecosystem.

The partnership with Siemens is a good example of where this is heading.

In 2025, NVIDIA and Siemens announced they would build an "industrial AI operating system." It is artificial intelligence that runs factories. From product design to manufacturing to quality control, AI coordinates every step.

Jensen Huang said this.

"Generative AI and accelerated computing have sparked a new industrial revolution. Digital twins are no longer passive simulations. They become active intelligence for the physical world."

A digital twin means creating a replica of the real world inside a computer. You build an exact copy of an entire factory inside a computer. You run experiments in this virtual factory first. If there are no problems, you apply the results to the real factory.

But now these digital twins think for themselves. They find problems, suggest solutions, and carry out optimizations. Artificial intelligence becomes the brain of the factory.

This is what Jensen Huang calls "the integration of OS, programs, and intelligence."

The operating system, programs, and intelligence merge into one. The boundaries disappear.

It used to work like this. There was an operating system, programs sat on top of it, and people operated those programs. The three were separate.

In the future, it works like this. Artificial intelligence plays the role of the operating system. It creates whatever functions are needed on the fly. People give instructions to the AI by speaking. The three melt into one.

There are concerns, of course.

Can we trust programs that AI creates? What if there are errors? Can we use them for anything important?

Those worries are fair. But looking back through history, new technologies have always started out imperfect.

The first automobile was slower than a horse. The first airplane stayed airborne for only a few seconds. The first computer filled an entire room and performed worse than a modern calculator.

Every technology starts that way. It is born imperfect, improves little by little, and at some point changes the world.

Jensen Huang said this at the 2024 World Government Summit.

"For the last 10 to 15 years, everyone has been saying that children should learn computer science. That they should learn programming. Now it's the opposite. Our job is to create computing technology so that nobody has to program. So that the programming language becomes human language. Everyone in the world is now a programmer. This is the miracle of artificial intelligence."

Everyone is a programmer. Consider the weight of that statement.

Until now, technology belonged to the few. People who learned to code, people who majored in engineering, specialists working inside companies. Only they could truly wield the power of computers.

Artificial intelligence tears down that wall.

Anyone who can speak can command a computer. Farmers, fishermen, teachers, students. If you can say what you want, artificial intelligence will make it happen.

Think again of Rome's roads.

Rome's roads bound the empire together. A farmer on the frontier could reach the markets of Rome. A merchant in the provinces could sell goods in the capital. The roads opened opportunity to everyone.

The AI operating system Jensen Huang envisions works the same way. The benefits of technology don't flow only to a handful of experts. They open to everyone. Anyone can bring their ideas to life. With words.

Of course, new questions arise.

If everything runs on artificial intelligence, who controls that intelligence? What happens when the AI makes a wrong judgment? Doesn't the company building the AI end up holding too much power?

Jensen Huang doesn't dodge these questions.

He talks about the concept of "sovereign AI," covered in Chapter 11. Every nation should have its own artificial intelligence. One country's AI must not be dependent on another's.

Whether that alone is enough remains unclear. New technology demands new rules. Writing those rules is not the job of technologists alone. It belongs to all of us.

Jensen Huang left this piece of advice at London Tech Week.

"You will not lose your job to artificial intelligence. But you may lose your job to someone who uses artificial intelligence."

Artificial intelligence is a tool. Or, to follow Jensen Huang's own words, artificial intelligence is "work." But deciding how to use that work still falls to people. Setting the direction for what AI should do is a human responsibility.

So Jensen Huang says: run. Don't walk.

Think back to the napkin from 1993.

When thirty-year-old Jensen Huang scribbled his idea on that napkin, did he know what the world would look like thirty years later? Probably not.

But he started. With an incomplete idea, on a flimsy piece of paper. That became a $5 trillion company. A company that makes the heart of the world's artificial intelligence beat.

We don't yet fully know what the operating system of the future will look like. Whether AI will truly change everything, how fast it will change things, is hard to predict. The change has already begun. And at the center of that change stands a Taiwanese immigrant in a black leather jacket. A man who started with a single napkin. A man now trying to rewrite the very way the world computes.

Jensen Huang.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 11. Data Is Territory

Part 2. Architect of a New Age

← Previous Chapter

See every chapter

Next Chapter →

The Romans laid roads on every land they conquered.

The Appian Way was one. The Aurelian Way was another. They understood something: without roads, legions couldn't move, and if legions couldn't move, the land didn't belong to Rome.

Roads were not just rows of stones. They were the veins through which Rome's will flowed, the channels that carried the empire's law and culture to its farthest reaches.

Two thousand years later, a man named Jensen Huang is meeting with leaders around the world, telling a similar story. Only the roads he talks about aren't made of asphalt.

They are fiber-optic cables that carry data, and AI factories that process it. And the territory he speaks of isn't physical land drawn by borders. It is a new kind of territory, invisible but undeniably real, called data.

Every Nation Must Have Its Own AI

In February 2024, the World Government Summit was held in Dubai. More than 4,000 delegates from 150 countries gathered there. Jensen Huang was invited as a speaker. Omar Al Olama, the UAE's Minister of State for Artificial Intelligence, sat down with him for a conversation. At that event, Jensen Huang introduced a concept to the world.

"Sovereign AI."

What does sovereignty mean? It is the right of a nation to govern itself without outside interference. The concept has been a foundation of the modern nation-state since the Treaty of Westphalia. Jensen Huang took this old word and brought it into the world of artificial intelligence.

His argument was straightforward. Every nation must have its own AI. It cannot simply borrow someone else's. Why not?

A nation's data contains everything about that nation. What its people think, what they buy and sell, what they say to each other, what diseases they get. This is not just a collection of numbers. It is a people's collective memory, the DNA of their culture, the footprints of their history.

Jensen Huang posed this question:

"Would you send your country's data to another country for processing, and then buy the resulting intelligence back at a premium?"

The question carries a sharp historical insight. Think of the 19th century. Colonized nations shipped raw materials to their colonial masters at rock-bottom prices. Rubber, cotton, crude oil. The imperial powers then processed those materials in their factories, turning them into tires, textiles, and gasoline. The finished goods were sold back to the colonies. At steep markups. Wealth flowed in one direction only.

Jensen Huang warns that the same thing could happen with data in the 21st century. Data is the new oil. But it matters far more than oil. Oil runs out; data keeps being generated. Oil is physical; data is the mind of an entire people.

What happens if a nation fails to build its own AI infrastructure? All of that nation's data flows into servers in Silicon Valley. Some company in California trains an AI model on it. Then that company sells the model back to the world. On a subscription basis.

This is a new form of dependency. A conquest carried out without guns or swords. Jensen Huang argues that every country needs to build its own 'AI factory' to prevent this. Just as power plants were once national infrastructure, AI factories that take in data and produce intelligence must now become national infrastructure.

In a BG2 interview, he put it this way:

"Nobody wants nuclear bombs. But everyone needs AI."

Nuclear weapons are arms you possess but can never use. They exist only as deterrence. AI is different. It gets used every day. In government administration, in healthcare, in education, in defense. The day is coming soon when a nation cannot function without AI. If that AI depends on another country, can we really call that nation sovereign?

In November 2024, Jensen Huang attended the APEC summit in Gyeongju, South Korea. There he announced large-scale partnerships with the Korean government and its major corporations. Samsung, SK, Hyundai, Naver. These companies were set to receive 260,000 NVIDIA GPUs. South Korea's Ministry of Science and ICT also committed to securing 50,000 GPUs to build sovereign AI infrastructure.

Jensen Huang said:

"Korea needs to build not just AI infrastructure but an entire ecosystem."

The word 'ecosystem' matters here. Buying chips and stacking them up doesn't complete sovereign AI. You need software running on those chips. You need talent to develop that software. You need companies where that talent can work. And all of it has to be connected organically.

Korea is well positioned in this regard. It has semiconductor manufacturing capability. It has software development capacity. It has homegrown tech companies like Naver and Kakao. It has a world-class manufacturing base. That is why Jensen Huang called Korea 'a key testbed for the physical AI era.'

Digital Sovereignty: Protecting Language and Culture

The importance of sovereign AI isn't only economic. There are cultural reasons too. These may be the more fundamental ones.

Think about the large language models that dominate the world today. GPT, Gemini, Claude. What data were they trained on? Mostly English. American English, specifically. Billions of documents, books, articles, and conversations floating around the internet. The overwhelming majority of this data comes from the Western world.

These models are smart. But that smartness carries built-in biases. Western values, Western historical perspectives, Western cultural nuances. All of it is baked into the model.

Say a Korean child asks an AI, "What is filial piety?" How would an AI trained on American data answer? It would explain the parent-child relationship from a Western lens, through the frame of independence and individualism. But could it convey the depth that the concept of filial piety holds in Korea, its historical weight, its connection to Confucian culture?

Jensen Huang took the stage at GTC Paris in 2025 alongside French President Emmanuel Macron and Mistral AI CEO Arthur Mensch. There, he said this.

"Intelligence is an inalienable national resource. You can outsource many things, but outsourcing your entire intelligence doesn't make sense."

He added.

"Intelligence comes from your country's data. That data belongs to your country. It is the knowledge of your people, their culture, their common sense. Most of the data isn't even on the internet. It's in libraries, in corporations. That data is yours."

France has a company called Mistral AI. It builds AI models tailored to the French language and European culture. NVIDIA is working with Mistral to build an AI cloud equipped with 18,000 GPUs. A 1.4-gigawatt AI campus is planned for construction near Paris.

Germany is doing the same. Deutsche Telekom has partnered with NVIDIA to build an "Industrial AI Cloud." Ten thousand Blackwell GPUs are being installed in a Munich data center. This is Germany's effort to maintain its standing as a manufacturing powerhouse in the AI era.

Look at India.

It is a country with 22 official languages and thousands of dialects. Can Silicon Valley's AI understand that diversity? The Indian government is developing AI that understands India's languages through

the "Bhashini" project. Major conglomerates including the Tata Group are partnering with NVIDIA to build an Indian sovereign AI ecosystem.

What about Saudi Arabia? This country once refined crude oil to accumulate wealth. Now they want to refine data to produce intelligence. Under the HUMAIN initiative, they are developing an Arabic AI model called ALLAM. They have also drawn up plans to provide AI training to one million citizens. AI factories are rising where oil refineries once stood.

Ukraine's case is striking. It is a country at war. Yet under its Minister of Digital Transformation, Ukraine launched a project called "Diia.AI LLM," an AI model tailored to the Ukrainian language, Ukrainian law, and public services. It is a declaration that the country will not surrender its digital sovereignty, not even in the middle of a war.

South Korea is moving too. In late 2025, the Ministry of Science and ICT and the National IT Industry Promotion Agency held a launch event at COEX in Seoul for the "Independent AI Foundation Project." Naver Cloud, SK Telecom, LG AI Research, Upstage, NC AI. Five companies are each developing Korean AI models in their own way.

Their strategies differ. Some pursue ultra-large-scale models. Some focus on efficiency. Others prioritize real-world industrial applications in manufacturing, defense, or gaming. But they share a common goal: building an AI that fully understands Korean, and that can produce answers suited to Korean law and culture.

There is still a gap. In global AI performance benchmarks, Google's Gemini and OpenAI's GPT score in the 70s. Korean models score in the 40s. But this is not a matter of short-term ranking. It is a matter of long-term survival. Securing data sovereignty, building AI that can be applied to industry, and achieving economic self-reliance through it. That is the essence of sovereign AI.

A New Form of National Power Competition

Looking back at history, the measure of national power has changed with each era. In the agrarian age, the size of your land was your power. After the Industrial Revolution, steel production was power. In the 20th century, oil reserves and nuclear arsenals were the yardsticks.

What will measure national power in the 21st century? Jensen Huang says it is computing power. How many high-performance chips a country possesses, how fast those chips can process data, how efficiently the resulting intelligence can be produced. This is the new measure of national strength.

He compares AI to a five-layer cake. At the bottom is energy. Running an AI factory requires enormous amounts of electricity. The second layer is semiconductor chips, high-performance computing devices like GPUs. The third layer is infrastructure, data centers and networks. The fourth layer holds AI models, things like GPT or Gemini. The top layer is applications, the services we actually use. Only countries that are competitive across all five layers can become true AI powers.

Start with energy. AI data centers consume staggering amounts of electricity. Jensen Huang put it this way: "China produces twice as much energy as the United States."

This is a warning that the United States urgently needs to expand its energy infrastructure. The best chips in the world are useless without the electricity to run them. This is why Nordic countries like Norway, Finland, and Sweden hold an advantageous position in the AI infrastructure race. They can produce abundant, cheap electricity from hydropower and wind. Then there is talent. According to Jensen Huang, a significant share of the world's AI researchers are in China. In sheer numbers, China's talent pool is massive. Some statistics show that China holds 70% of all AI patents. If the United States wants to maintain its technological edge, export controls alone won't be enough. It needs to create an overwhelming technology gap. Huang has expressed concern about American strategy on this front. If the U.S. tries to shut China out through export controls, China will build its own independent technology ecosystem. In the long run, that could weaken America's grip on global technology standards. He said this.

"We are not losing market share. We could be losing the future standard."

This is a complex geopolitical dilemma. Shutting out a competitor may look favorable in the short term, but if that competitor goes its own way, the world fractures. Jensen Huang's strategy is different. He wants to help every country build its own sovereign AI on top of NVIDIA's technology stack. American technology becomes the global standard, while each nation's data sovereignty is respected.

This is why Jensen Huang travels the world so relentlessly. Taiwan, South Korea, Japan, India, France, Germany, the United Kingdom, Saudi Arabia, the United Arab Emirates. He meets heads of state, meets business leaders, signs partnerships. His title says CEO of NVIDIA, but his role goes beyond that. He serves as the hub of a Western technology alliance. You could call him the technology diplomat of the AI era.

In June 2025, NVIDIA announced it would deploy 3,000 exaflops of Blackwell systems to major European countries. France, Germany, Italy, the United Kingdom. AI infrastructure is being built in all of them. At his GTC keynote in Paris, Jensen Huang said this.

"Europe has now realized the importance of AI infrastructure. Within two years, we will increase Europe's AI computing capacity tenfold."

He reminded his European audience of their own history. Europe was the birthplace of the Industrial Revolution. The steam engine first appeared here. Electricity was first adopted on a mass scale here. What is happening now, he said, is a revolution of the same magnitude. AI is the electricity of this era. AI factories are the power plants of this era.

Jensen Huang divides national security into two kinds. Written with a lowercase n and s, national security means military defense. Written with an uppercase N and S, National Security encompasses the dynamism of an economy, the productivity of industry, and cultural influence all together.

Sovereign AI is the cornerstone of this macroscopic national security framework. Without AI infrastructure, a country's scientists cannot conduct research. Its companies cannot create innovative services. Its manufacturing sector falls behind in automation and efficiency. Healthcare, education, public administration, every field stagnates.

A country that successfully builds sovereign AI stands in a different position. It can rearm its industries with AI. It can raise productivity by leaps and bounds. Jensen Huang puts it this way.

"Every industry will become a producer of intelligence."

An automaker is not just building cars. It is building the intelligence of autonomous driving. A pharmaceutical company is not just manufacturing drugs. It is building biological intelligence for new drug discovery. At the center of all these processes sits each nation's sovereign AI.

What we are witnessing is a new industrial revolution. Jensen Huang calls it "the hundred-trillion-dollar opportunity." If the old IT industry was a three-trillion-dollar tools market, AI affects the entire global economy because it replaces the work that humans used to perform.

In this massive transition, countries that fall behind become "intelligence consumers," forced to import intelligence produced by other nations at steep prices. Countries that move ahead become "intelligence producers," exporting intelligence and transforming their own industries.

Jensen Huang offers this advice to governments around the world.

"Start right now. The pace of technological progress will not wait for you. If you miss this moment, you may never catch up."

This is not an exaggeration. The speed of AI advancement has outpaced Moore's Law. NVIDIA has publicly committed to reducing the cost of token production by more than a millionfold within ten years. Every year brings a new architecture, and performance improves exponentially. Those who fail to board this train will stand on the platform forever, watching it pull away.

Data is territory. Chips are weapons. AI models are a nation's mind. Only countries that possess all three can remain sovereign states in the twenty-first century.

If the Treaty of Westphalia in 1648 established the concept of the modern sovereign state, then the sovereign AI that Jensen Huang champions is defining the concept of the "digital sovereign state." Just as a people who lose their territory have no future, a nation that surrenders its data and entrusts its intelligence to another country has no independent future.

Jensen Huang is not just a person who sells chips. He is spelling out the preconditions for nations to survive and prosper in the coming era. He hands each country the tools and helps them build their own digital walls.

If Caesar expanded the borders of Rome, Jensen Huang is teaching each nation how to build its own digital fortress. The heat pouring out of AI factories is not a mere byproduct of computation. It is the

hot breath proving that a nation still exists as a sovereign entity in the twenty-first century.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 12. The Arrival of Physical AI

Part 2. Architect of a New Age

← Previous Chapter

See every chapter

Next Chapter →

Intelligence Escapes the Screen and Gains a Body

It took the Romans a long time before they could call the Mediterranean 'our sea,' Mare Nostrum. Rome started as a small city-state, unified the Italian Peninsula, destroyed Carthage, and finally embraced Greece and Egypt. Only then could they declare the entire sea their own territory. Jensen Huang's path looks similar. He conquered the small pond of gaming first. Then he crossed the river of scientific computing. He sailed across the ocean of artificial intelligence. Now his gaze reaches beyond the sea, toward machines that move in the physical world.

For the past several years, we have been captivated by AI that lives inside screens. We watched in amazement as ChatGPT wrote poetry, Midjourney painted pictures, and Sora produced videos.

But what happens if you ask this brilliant intelligence to make you a cup of coffee? Nothing. It has no hands. No feet, no eyes, no skin. These AIs are ghosts trapped in a world of bits. No matter how smart a ghost is, it can't turn a doorknob.

Jensen Huang saw through this limitation early on. He thought: why should intelligence remain locked inside a screen? If the human brain had stayed confined within the skull, neither the Pyramids nor the Great Wall would exist. The brain's commands travel along nerves to reach muscles, and those muscles moved stones and stacked bricks. That is how civilization rose.

The same applies to artificial intelligence. It reaches completion only when it moves beyond giving answers on a screen and begins helping in the real world.

This is what Jensen Huang calls Physical AI. To understand the concept, you first need to know the limits of existing AI. Ask ChatGPT how to peel an apple, and you'll get a flawless text answer.

It explains the angle to hold the knife, the direction to peel the skin, and the order to remove the seeds, all in fine detail. But hand that AI an actual knife and an apple, and nothing happens. It knows only the probabilistic relationship between the word 'apple' and the verb 'peel.' It has no understanding of the slipperiness of a round apple surface or the resistance a blade meets when it cuts into the skin.

Physical AI is different.

It knows gravity. It understands friction. It grasps what inertia is and recognizes cause and effect.

When a ball rolls under a sofa and disappears from view, a conventional AI might conclude the ball has vanished.

Physical AI, however, recognizes that the ball still exists; it is simply out of sight. Psychologists call this object permanence. Even a baby a few months old picks up this concept, yet large language models with trillions of parameters failed to understand it.

Across several keynote speeches in 2024 and 2025, Jensen Huang defined this shift as AI's next wave. According to him, artificial intelligence has passed through three stages of evolution.

The first wave was perception. Computers learned to recognize images and understand speech.

The second wave was generation. They wrote text, drew pictures, and produced videos.

The third wave is physical action. AI begins to perceive its environment through sensors, make judgments with its brain, and intervene in the real world with arms and legs.

Consider why this shift matters. The digital revolution so far has mostly played out on office desks. What computers and the internet changed was the way we draft documents, send emails, and search for information.

Yet a large share of the global economy still runs on physical labor. Workers assemble parts in factories, carry boxes in warehouses, lay bricks on construction sites, and care for patients in hospitals. Jensen Huang believes a true industrial revolution is complete only when AI enters this physical domain.

Transplanting digital intelligence into physical machines, though, is far harder than it sounds. A badly written sentence can be erased, but a robot that moves incorrectly can injure someone. An autonomous vehicle that misjudges a situation can cost lives. That is why Jensen Huang says an entirely new computing architecture is required.

His solution is a three-computer system.

The first computer is a supercomputer that trains the AI. NVIDIA's Blackwell and DGX systems fall into this category. Here, the AI ingests vast amounts of data and forms its intelligence.

The second computer handles simulation. It creates a virtual world where robots will operate, and it tests and validates the AI inside that world. The third computer is the brain that goes inside the robot's body. It is a compact computer loaded with a trained AI model, enabling the robot to interact with the real world.

When these three computers are linked together, intelligence finally leaves the screen and takes on physical form. Intelligence born in a data center learns to walk in a virtual world, then is implanted into a body of steel and begins working in factories, on roads, and in homes. Jensen Huang transformed NVIDIA into the company that supplies the brain and the nervous system for this entire process.

Where earlier AI converted text into tokens for processing, Physical AI converts a robot's movements into tokens. In Jensen Huang's words, from the computer's perspective there is no difference between writing a poem, drawing a picture, or moving a robotic arm. They are all just numbers, tokens. Now, when we tell a robot to pick up a blue box and place it on a red shelf, the AI understands the instruction, calculates the joint angles of the robotic arm and the force of each motor, and carries out the action.

What Caesar demonstrated in the Gallic Wars was not mere military genius. He understood the importance of supply lines, systematized logistics, and made the Roman legions move as a single organism. Jensen Huang operates the same way. He went beyond making good chips and designed an entire ecosystem for AI to function in the real world. By connecting three axes, training, simulation, and execution, he opened a channel for digital intelligence to flow into the physical world.

This was something he had been preparing for eight years. Back when no one paid attention, he developed robot simulators, refined physics engines, and designed edge computers. So when the era of Physical AI arrived, NVIDIA already held every piece of the puzzle.

Omniverse: The School Where Robots Learn About the World

What made the Roman legions so powerful? Superior weapons played a part, but the real secret was training. Roman soldiers trained at wartime intensity even during peacetime. They marched dozens of kilometers a day carrying heavy loads. They formed and broke battle formations against imaginary enemies, over and over. So when an actual battle came, war was already a familiar routine. Jensen Huang believed robots needed a training ground like this, too.

But there was a problem. Training robots in the real world is too dangerous, too slow, and too expensive. Consider that a robot has to fall millions of times to learn how to walk. Every fall damages its parts. Repairs cost time and money. If a self-driving car causes an accident while learning to drive, the damage can't be undone. And real-world time moves too slowly. You can't wait ten years to teach a robot a skill that takes a human ten years to master.

This is where Jensen Huang's counterintuitive thinking comes in. If you can't learn in reality, build a virtual world identical to reality. That was the origin of Omniverse.

Omniverse is not just a 3D graphics tool. It is a complete physical world where gravity acts, friction exists, and light reflects off objects. Think of it as a digital Earth governed by Newton's laws of motion.

NVIDIA developed a physics engine called Newton in collaboration with Google DeepMind and Disney Research. This engine calculates not only rigid body dynamics but also soft body deformation, fluid flow, and haptic feedback, all in real time.

Robots are born, learn, and grow inside this virtual world. They fall tens of thousands of times as they figure out how to walk. In reality, a single fall means stopping for repairs. In Omniverse, a robot gets back up with the push of a button. Time can be fast-forwarded so that years of learning are compressed into just a few days. Jensen Huang called this a school where robots learn how to be

robots.

The teaching method at this school resembles how humans learn.

Think about a child learning to ride a bicycle. At first, the child falls. Through countless rounds of trial and error, the child figures out how hard to push the pedals, what angle to turn the handlebars, and how to keep balance. No parent instructs the child to 'lean three degrees to the right.' The child learns through the body. Robots inside Omniverse work the same way. No programmer dictates every movement through code. The robots go through trial and error on their own and discover the optimal way to move.

NVIDIA took it one step further.

They introduced a model called Cosmos. Cosmos is an AI that learns from video and data of the physical world to predict and generate the future. It is as if a robot were dreaming.

Humans learn through dreams too. Experiences from the day replay in dreams at night and solidify into memory. Cosmos gives robots this ability to dream. When a developer shows a robot a single demonstration of picking up an object, Cosmos imagines thousands, even tens of thousands of scenarios for how that action would play out in different environments. What happens if the lighting changes. What should the robot do if the object is in a different position. How should it plant its feet if the floor is slippery. Through this generated virtual data, a robot can acquire general-purpose skills from just a handful of demonstrations.

Jensen Huang described this as converting computing into data. The biggest obstacle in robot training had always been a shortage of data. Language models like ChatGPT could learn from the vast ocean of text on the internet. But data showing robots walking, picking up objects, and moving through complex environments simply did not exist online. NVIDIA solved this problem by generating unlimited data inside virtual worlds.

In his keynote addresses of 2024 and 2025, Jensen Huang demonstrated a small Disney robot learning to walk inside Omniverse. This robot, called BD-X, fell over and over in the virtual world until it learned to keep its balance. When it stepped into the real world for the first time, it walked as naturally as a seasoned veteran. Experts call this Sim-to-Real, the transfer from simulation to reality.

The uses of Omniverse go well beyond robot training. Entire factories can be built virtually. Giant manufacturers like BMW, Mercedes-Benz, and Foxconn construct perfect digital twins inside Omniverse before breaking ground on real facilities. Inside these virtual factories, they test how thousands of robots collaborate. They check in advance whether workflow paths get tangled, where bottlenecks emerge, and whether robotic arms collide with each other.

Jensen Huang says: every factory of the future will exist in two versions. One is a physical factory that produces physical goods. The other is a digital factory that produces the intelligence to run the physical one. Instead of building from a blueprint, you build after previewing the finished future.

This is a fundamental shift. In the past, you built first and fixed problems as they appeared. Now, you discover and solve every problem before construction begins. The cost of trial and error has moved from the real world to the virtual one.

On top of the Omniverse platform, NVIDIA placed an ambitious initiative called Project GR00T.

GR00T is a foundation model for humanoid robots. Just as we teach a child to take its first steps, GR00T teaches robots to understand the world and move through it. The model comprehends human language, observes and imitates human movements, and acquires skills through trial and error in virtual worlds.

Jensen Huang released this technology as open source. Robot developers around the world no longer need to build AI from scratch; they can simply adopt the brain NVIDIA has built. Leading robotics companies like Figure AI, Boston Dynamics, and Agility Robotics are already developing their next-generation robots on this platform.

When Rome paved its roads, it was not only armies that traveled on them. Merchants came and went, scholars journeyed, and culture spread. Omniverse works the same way. The virtual world NVIDIA built has grown beyond a mere technology platform into the infrastructure of the physical AI era.

Autonomy for Everything That Moves

Jensen Huang often says this: everything that moves will become autonomous. This is not simple automation. Automation is a machine repeating a fixed set of instructions. A washing machine running a laundry cycle or an elevator moving between floors is automation. Autonomy is different. A machine perceives, judges, and acts on its own. When conditions change, it finds its own response.

The first industry hit by this wave of autonomy is the automotive sector. Jensen Huang defines cars as robots on wheels. Where cars once competed on engine horsepower and exterior design, they now compete on the computing power and software capability inside the cabin.

NVIDIA's DRIVE platform has reshaped the automotive landscape. Mercedes-Benz, Jaguar Land Rover, Hyundai Motor, Volvo, and other major automakers install NVIDIA chips in their vehicles. Jensen Huang calls this the Software-Defined Vehicle. Just as a smartphone gains new features through software updates, a car has become a machine that keeps getting smarter after purchase.

The Alpamayo model, unveiled in 2025, shows the peak of this transformation. It is an end-to-end learning model that takes visual information from cameras and directly controls the steering wheel and brakes. What makes it remarkable is that the AI doesn't just drive; it can explain why it made a given decision. 'The car ahead has its brake lights on, and a vehicle in the next lane is trying to merge, so I'm slowing down.' It drives while thinking this way. This solved the black-box problem of earlier self-driving systems, the inability to know why the system made a particular judgment.

But Jensen Huang's ambitions don't stop on the road. His gaze turns to factories, warehouses, hospitals, and homes. Humanoid robots are the crowning achievement of physical AI. The reason Jensen Huang focuses on humanoids is straightforward: our world was designed for humans. The height of door handles, the rise of stair steps, the shape of tools, the size of chairs are all built around the human body. Only a robot with a human form can be deployed instantly throughout our society without modifying the environment.

As of 2026, humanoid robots are appearing across the globe. Boston Dynamics' Atlas has switched to electric actuation and independently performs part-handling tasks in factories. Tesla's Optimus targets a price range of 20,000 to 30,000 dollars and aims to be a general-purpose robot capable of everything from factory work to household assistance. Figure AI's humanoid assists with assembly and material transport at BMW factories.

The dexterity of these robots is advancing at a striking pace. The sensory resolution of a human fingertip is about 2.8 to 3.5 millimeters, while the latest tactile sensors boast precision at the 30 to 100 micrometer level. Hands with sensitivity far exceeding that of humans are beginning to be attached to robots.

Factories themselves are becoming giant robots.

Jensen Huang sees the factory of the future as one massive intelligent system. The building itself is a robot, and the thousands of robotic arms and autonomous mobile robots moving inside it operate like organs in a living body. Amazon has already deployed one million robots across its fulfillment centers. These robots use artificial intelligence to recognize and sort items, then travel along optimal routes.

Physical AI is bringing a revolution to medicine as well. AI-powered surgical robots are achieving 40 percent greater surgical precision, 15 percent shorter patient recovery times, and 20 percent higher surgical efficiency. The da Vinci robot learns from surgical footage to autonomously adjust camera angles and track a surgeon's instruments the way a human assistant would.

Behind all of these changes lies a demographic shift. The working-age population is shrinking worldwide. In developed nations, birth rates are falling and populations are aging. Jensen Huang is convinced that robots won't steal human jobs; instead, they'll become essential partners filling the gaps left by a shrinking labor force.

On dangerous construction sites, in semiconductor fabrication lines that demand extreme precision, and in medical settings that require round-the-clock patient care, physical AI will do the heavy lifting in place of humans. Robots take on the repetitive, hazardous labor that people avoid, freeing humans to focus on more creative and valuable work.

There are challenges, of course. Reproducing in the real world what a system learned in simulation is never straightforward. A gap always exists between virtual and physical environments. Lighting conditions differ. Surface textures differ. Unexpected situations arise. Experts call this the Sim-to-Real Gap.

NVIDIA and robotics companies are narrowing that gap through techniques such as domain adaptation, domain randomization, and transfer learning.

Safety and ethics present their own issues. Because physical AI systems exist in the physical world, they carry physical risks. Some reports indicate that workplace injuries have risen after AI-controlled robots were introduced. There have been cases of robots stepping on children's feet or operating erratically. Balancing innovation with safety and ethical standards remains a task ahead.

The impact on labor markets can't be ignored either. According to some studies, hundreds of millions of workers worldwide could be displaced by automation by 2030. Low-skilled workers who are replaced by automation struggle to find comparable or better jobs. High-skilled workers, on the other hand, can cushion the blow of automation with stronger career prospects. Making sure the benefits of technological progress are distributed fairly remains a challenge for society.

Jensen Huang is optimistic all the same. He defines robots not as tools but as partners. When robots take over the dangerous, repetitive labor that humans don't want to do, people can concentrate on more creative and meaningful work. He believes physical AI doesn't replace humans; it amplifies what humans can do.

NVIDIA's chips, Jetson and Thor, serve as the brain for everything that moves. From small delivery robots to massive construction equipment, an NVIDIA brain is implanted inside. Just as the emperors of Rome made all roads lead to Rome, Jensen Huang has made every moving machine think and act on top of NVIDIA's platform.

We now stand at the threshold of the future Jensen Huang designed. AI has stepped out from behind the screen and walks beside us, works alongside us, and sometimes helps us. For Jensen Huang, none of this was an accidental discovery. It's the inevitable result of years spent building robot chips and developing simulators, quietly stacking one layer on top of another since a time ten years ago when nobody gave any of it a second look.

He always says it: "We're a company that could be out of business in 30 days." That humble, urgent attitude kept him from ever settling, and it's what drove him to bring intelligence into the physical world. The arrival of physical AI is not a mere step forward in technology. It is a sweeping civilizational shift that fundamentally changes how humanity uses tools and coexists with them. And standing at the center of that shift is an engineer in a black leather jacket: Jensen Huang.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Chapter 13. Coexistence of Humans and Machines

Part 2. Architect of a New Age

← Previous Chapter

See every chapter

Next Chapter →

When the Roman Empire was at its height, all roads led to Rome. A network of roads stretching tens of thousands of kilometers connected every province into a single whole. Soldiers marched along these roads, merchants hauled their goods, and scholars spread knowledge. As physical distances shrank, civilization grew at an explosive pace.

What Jensen Huang is building is also a kind of road. A neural network made of silicon and code, an invisible highway where data flows. A new traveler called artificial intelligence walks along this road. And this traveler wants to become our colleague.

In the period crossing from 2024 into 2026, a single question echoed around the world. Will artificial intelligence really take our jobs? Painters wondered whether they should put down their brushes. Writers agonized over whether to break their pens. Even programmers trembled at the fear of being replaced by the very machines they had built.

Jensen Huang saw it differently.

AI Won't Take Your Job

Jensen Huang repeated the same message across speeches and interviews. "AI isn't going to take your job. Another person who knows how to use AI well is going to take your job."

The first time you hear this, it can sound cruel. But think about it for a moment and you realize it echoes a universal truth of history. The one holding the tool has always had the advantage over the one with bare hands. A hunter with a stone axe caught more prey than a hunter without one. A farmer pulling an iron plow harvested more grain than a farmer pulling a wooden plow.

Jensen Huang draws a clear line between a person's "job" and their "tasks." A job is a collection of roles a person performs. Tasks are the individual pieces of work required to carry out that job. What AI takes over is not the entire job but the tedious, repetitive tasks buried inside it.

He offered an interesting example. About ten years ago, Geoffrey Hinton, a leading authority in artificial intelligence, made a prediction: "Within five years, AI will read medical images better than radiologists. There will be no need to train radiologists anymore."

On a technical level, Hinton was right. Today, AI reads X-rays and MRIs faster and more accurately than humans. But something strange happened. Radiologists didn't disappear. Their numbers

actually grew.

Why?

One of a radiologist's "tasks" was reading images. But the reason the "job" existed, its purpose, was to diagnose disease and treat patients. Once AI took over the grueling, repetitive work of image reading, doctors could analyze more patient data in greater depth. Diagnostic accuracy went up. Hospitals could care for more patients. The quality of medical services improved. The result was that more doctors were needed.

Jensen Huang compares this to the spreadsheet. When Excel came along, did accountants vanish? They did not. The hours spent punching calculators and reconciling numbers simply shrank. With that freed-up time, accountants could craft financial strategies and advise executives, work of far greater value.

Jensen Huang points to NVIDIA's own example. Every software engineer and chip designer at NVIDIA works with the help of AI tools. Inside the company there is an AI coding assistant called Cursor. It suggests code, finds bugs, and summarizes documents.

So did NVIDIA lay people off? No. It hired even more. As productivity rose, new projects that had been impossible for lack of time suddenly became feasible. Human ideas are limitless. When tools get better, more of those ideas can become real.

Jensen Huang puts it this way: "If a tool can shrink a week's worth of work to an instant, will you end up doing less work, or more?" His answer is clear. "You're going to be busier."

This can sound like a paradox. If the tools get better, shouldn't life get easier? Why would you get busier? The reason is that humans have endless desires and curiosity. Once one problem is solved, they tackle the next. Once they cross one mountain, they go looking for a taller one.

Jensen Huang raises another critical point. The world is facing a severe labor shortage. In manufacturing, transportation, logistics, healthcare, and nearly every other field, there aren't enough people to do the work. The cause is aging populations and falling birth rates. Some projections say the global workforce will be short by roughly 50 million workers by 2030.

AI and robots fill that gap. They take on jobs that humans avoid or that are physically impossible for them. Dangerous factories, deep mines, sprawling warehouses. When machines work in places like these, humans can focus on safer, more creative work.

In January 2026, at CES, the technology trade show, Jensen Huang used a striking phrase. He called robots "AI immigrants." Just as immigrants come to a new land and help its economy grow, robots and AI would join our economy and grow alongside us.

He said: "When robots arrive, jobs arrive with them. Population decline is keeping us from achieving the economic growth we want. AI immigrants need to help us on the production floor. When the economy grows, you naturally end up hiring more people."

In the end, artificial intelligence is not humanity's enemy. It is a lever that amplifies human ability. Just as the tractor didn't replace the farmer but multiplied the farmer's strength, AI multiplies the strength of knowledge workers.

A Future of Working Alongside Digital Colleagues

The office of the future Jensen Huang envisions looks quite different from today's. In it, "biological employees" and "digital employees" work side by side. The biological employees are humans. The digital employees are AI.

The digital employees he describes go well beyond answering questions. They reason on their own, make plans, and use tools to solve problems. Jensen Huang calls this "agentic AI."

The word "agent" originally means a representative, someone who handles affairs on your behalf. Insurance companies have agents. Real estate firms have agents. AI agents are no different. They handle complex work in our place.

Imagine a situation like this. You say, "Plan a trip to Paris for next week." A basic chatbot would stop at listing tourist spots. An agentic AI is different. It searches the internet on its own

and compares flights. It checks hotel prices. It puts together an itinerary and adds it to your calendar. It even makes reservations. It handles a variety of tools with the same ease as a human.

Jensen Huang predicts that future companies will have billions of digital employees. In a marketing department, one AI agent handles market research. Another helps with copywriting. Another analyzes customer data. Human employees collaborate with these agents, distributing tasks and gathering results like a team leader.

He offered an amusing analogy while explaining this shift. "The IT department of the future will become the HR team for digital employees."

Think about what an HR team does. They hire new employees. They teach those employees the company's culture and philosophy. They evaluate job performance. They manage and train people when problems arise. According to Jensen Huang, the IT department of the future does exactly the same thing. It decides which AI models to adopt. That is hiring. It trains the AI on the company's knowledge and data. That is onboarding. It monitors and evaluates whether the AI is functioning correctly. That is performance management.

He says he told NVIDIA's CIO: "You will become the HR director for digital employees in the future. Digital employees will work alongside our biological employees. That is the future of our company."

This shift will fundamentally reshape corporate organizational structures. In the past, software was used as a 'tool.' You opened Excel, opened Word, opened Photoshop. Humans used tools to do work. But agentic AI is not a tool. It is a workforce that performs tasks on its own. AI becomes the entity that uses tools.

Jensen Huang points to NVIDIA's chip design process as an example. NVIDIA's latest chips contain billions of transistors. A transistor is a tiny switch that either allows or blocks electrical current. Billions of these switches come together to form a complex city. Designing and verifying this enormous maze by human effort alone is nearly impossible.

NVIDIA designs chips through AI. AI assists with the design, and those chips in turn create more powerful AI. This is a cyclical structure. Jensen Huang calls it the 'AI factory.' If factories of the past produced goods, factories of the future produce intelligence.

Humans are not marginalized in this process. They focus on higher-order problem-solving instead. How do we stop climate change? How do we develop new treatments? How do we explore space? Human intelligence alone is insufficient to solve these challenges. Digital colleagues called AI work around the clock, analyzing data and running simulations. Humans make decisions based on those results.

In the world Jensen Huang envisions, the human role is that of a 'conductor.' Just as an orchestra conductor gives direction to each musician, humans set the direction for AI agents. Defining what to do and why it matters remains a human responsibility. AI provides the answer to how.

Humans also become 'curators.' A curator selects artworks and plans exhibitions at a museum. Among the countless outputs AI generates, judging what holds real value, what is ethically sound, and what can touch human emotion requires human discernment.

Jensen Huang describes this discernment as 'taste.' AI can write excellent poetry. But feeling why that poem is beautiful is human. AI can present a flawless treatment plan. But holding a patient's hand and offering comfort is the doctor's role.

As technology advances, the value of what is distinctly human rises paradoxically. Empathy, creativity, ethical judgment. These are things machines cannot easily replicate. We are not competing with machines. We must learn to dance with machines and lead more human lives.

Jensen Huang says: "The era of programming languages is passing. Now the programming language is human language." In the past, you had to learn complex machine languages like C++ or Python to give commands to a computer. Only a handful of specialists enjoyed the benefits of technology. But now anyone can give commands by speaking. In English, in Korean, in Japanese. "Analyze this data for me," "Write code that implements this feature," "Draw me a picture like this." If you don't even know how to use AI, you can just ask it."

"How do I use you?"

This means an era has opened where everyone can be a programmer. A farmer can analyze agricultural data. A biologist can study gene sequences. An artist can use new creative tools. Countless people who couldn't realize their ideas because of technical barriers will now unleash their creativity with the help of digital colleagues.

Jensen Huang calls this the 'democratization of technology.' Democratization means power spreading from the few to the many. The democratization of technology means the power of technology spreading from a small number of specialists to everyone.

"Run, Don't Walk"

May 2023, the commencement ceremony at National Taiwan University. Thousands of graduates have gathered on the athletic field. A man in a black leather jacket stands at the podium. It is Jensen Huang. He has just received an honorary degree from the university and is delivering the commencement address.

His words resonated in the hearts of young people around the world.

"Run, don't walk. Either you are running for food, or you are running from being food. Sometimes it's hard to tell which. Either way, run."

Run, don't walk. Either you are running for food, or you are running from being food. This sentence is Jensen Huang's life philosophy. It is also the question he poses to all of us living in a rapidly changing age of AI.

In the natural world, walking means danger. You get caught by a predator, or you miss your prey and starve. The world of technology is the same. In an era where the speed of change is exponential, 'walking' is not maintaining the status quo. It is falling behind.

Jensen Huang has weathered countless crises in the more than 30 years since founding NVIDIA. The company came to the brink of bankruptcy more than once. When his first chip failed, he went to Sega and honestly admitted his mistake. When he had to withdraw from the mobile market, he boldly pivoted toward AI, a field no one was paying attention to.

Each time, he chose to run instead of walk. He felt in his bones that the moment he stood still, he would be left behind.

The pace of technological progress shoots upward along an exponential curve. Yesterday's innovation becomes today's common sense and tomorrow's obsolescence. According to Jensen Huang, AI computing performance has improved 100,000 times over the past decade. Moore's Law said performance doubles every two years, but AI has far exceeded that.

The 'running' Jensen Huang speaks of does not mean a blind sprint for speed alone. It means the agility to sense change, refuse to fear failure, and pivot boldly toward new opportunities. NVIDIA was able to transform from a graphics card company into an AI computing company because he refused to rest on his success in the PC market and sprinted full speed into a field nobody was watching.

He advises the younger generation: do not avoid pain and hardship. Just as muscles grow stronger after the pain of tearing, intellect grows through failure and trials. When people ask Jensen Huang the secret to his success, he answers plainly. "I have endured tremendous pain and humiliation. That is what made me."

The current AI revolution is approaching faster and on a larger scale than either the Industrial Revolution or the arrival of the internet. Hesitating or stepping backward in front of this enormous wave is dangerous. You must ride it. You must not hesitate to learn new technology and turn it into your own weapon.

Jensen Huang says he uses artificial intelligence every day. He uses ChatGPT. He uses Perplexity.

He uses Gemini. He treats them like personal tutors. When he encounters an unfamiliar subject, he asks the AI: "Explain it to me as if I were a twelve-year-old." The CEO of the world's leading AI company is still learning every day. Can the rest of us afford to stand still?

He says, "Start using artificial intelligence right now." It doesn't matter whether you're a student, a working professional, or retired. There is no need to wait for the future. The future is already here. In 2016, when Jensen Huang personally delivered the first supercomputer to OpenAI, they were dreaming of what was to come. Now that dream is reality. A walker can admire the scenery. But only the runner gets to see the new horizon first.

Even as Jensen Huang receives praise as the world's greatest CEO, he still greets every morning with the mindset that "we are thirty days from going bankrupt." This is not anxiety. It is a healthy sense of urgency. He guards against becoming intoxicated by success. The Roman Empire did not decline because of enemies outside its borders. It fell because of complacency and arrogance within. Jensen Huang relentlessly pushes the organization so that NVIDIA, despite becoming a corporate giant, never loses the hunger of a startup.

The world he envisions is not a dystopia where machines rule over humans. It is an era of creation, one where diseases are conquered, the climate crisis is addressed, and anyone can turn an idea into reality. To bring that about, he steps onto the stage today wearing his leather

jacket. And he tells us: technology is only a tool; it is human will that decides where to run with that tool in hand.

The future is not something to predict. It is something to build. And to build it, we must keep running, never stopping, even at this very moment.

Just as Jensen Huang has done, it is time for each of us to run on our own track until we are gasping for breath. Because that is the path to coexisting with machines while preserving human dignity.

← Previous Chapter

See every chapter

Next Chapter →

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Epilogue: To Those Standing at the Door

← Previous Chapter

See every chapter

As I finish this book, I find myself thinking of a passage I read long ago.

History is not a record of what has already happened, but a mirror for those preparing for what has not yet come.

That sentence stayed with me throughout the writing of Jensen Huang's story.

In October 2025, Jensen Huang visited South Korea.

He came for the Asia-Pacific Economic Cooperation summit held in Gyeongju. Wearing his signature black leather jacket, he took the stage and spoke about Korea's future. He announced that NVIDIA would supply 260,000 Blackwell GPUs to Korea. That evening, he sat in a fried chicken restaurant in Seoul with the chairmen of Samsung Electronics and Hyundai Motor, sharing chicken and beer. The sight of the world's ninth-richest person tipping back a beer glass at a neighborhood chicken joint left a lasting impression on many people.

Watching that scene, I wondered: who is this man, really? How can the head of a company worth five trillion dollars in market capitalization smile so easily? What made him who he is?

This book is my answer to those questions.

Looking into Jensen Huang's life, I discovered several things. His path was never smooth. He was sent to a foreign country at the age of nine and had to survive among tattooed teenagers. His first product was a devastating failure, and his company teetered on the edge of bankruptcy more than once. Wall Street mocked his decisions, and competitors dismissed him.

But he never gave up. Every time he failed, he got back on his feet. Every time he was ridiculed, he kept walking his own road in silence. When he invested ten years in a technology called CUDA, the world said it was useless. He believed otherwise. He believed a time would come when the world would need it.

And that time came.

In 2012, an artificial intelligence model called AlexNet stunned the world. It ran on NVIDIA's GPUs. The technology Jensen Huang had quietly built over a decade was finally showing its power. There is a saying that only the prepared can seize opportunity. Jensen Huang is living proof of that.

I started out in the legal profession and have walked many roads since. I worked as a prosecutor, moved into private practice as an attorney, and served in the National Assembly. Then I encountered a new world called artificial intelligence. At first it was curiosity. But gradually I realized this was not

just a technological shift; it was a turning point for civilization itself.

I asked myself what I could do in this new era, as a legal professional, a policy expert, and a citizen. This book is one result of that reflection. Through Jensen Huang's story, I wanted to convey the essence of the AI era. I believed the human story behind the technology mattered more than the technology itself.

The history of the Roman Empire reveals a pattern worth noting.

Whenever the empire flourished, people from the provinces stood at its center. Trajan came from Hispania. Septimius Severus came from Africa. Because they were not native Romans, they could see Rome more clearly than Romans could.

Because they were not bound by the old frameworks, they were free to try something new.

Jensen Huang is the same. An immigrant born in Taiwan who crossed the ocean to America. He was never part of Silicon Valley's mainstream. Among elites from Stanford and MIT, a graduate of Oregon State University was always an outsider. Yet it was the outsider's eye that saw what others could not: that a graphics chip built for gaming could become the brain of artificial intelligence.

While writing this book, I was moved many times. The story of a nine-year-old boy at a boarding school in Kentucky teaching his illiterate roommate how to read. The story of a young CEO honestly confessing to a Sega executive that their chip was a failure. The story of a businessman approaching sixty personally delivering a supercomputer to OpenAI.

All of these stories share one thing in common: honesty. Jensen Huang never hid his failures. He never pretended to know something he did not. He often uses the phrase intellectual honesty. Admitting you might be wrong, and never stopping to ask questions in search of a better answer. That is what he means by intellectual honesty.

I believe this attitude is something all of us need as we live through the age of artificial intelligence. The world is changing fast. Yesterday's right answer can become today's wrong one. So we must stay humble. We must never stop learning. We must keep moving forward even as we accept that we might be mistaken.

In his commencement speech at Stanford University, Jensen Huang told the graduates: Run, don't walk. Whether you're running to find food or running to avoid becoming food. His words may sound harsh. But they hold a truth. In an era of change, those who hesitate fall behind. Only those who move forward survive.

To the readers turning this final page, I want to say just one thing. We are standing before an enormous door right now. Beyond it lies a new world that artificial intelligence will reshape. No one yet knows exactly what that world will look like.

But one thing is clear. The door is already open. Whether to walk through it is each person's choice. Jensen Huang has already passed through and is running far ahead. You don't need to follow him.

But you can look back at the path he walked and find your own.

A boy born under the sun of Tainan changed the world. You can change it too. Each in your own place, each in your own way. What matters is not giving up. What matters is getting back on your feet every time you fall. That is the way of life Jensen Huang has shown us.

The story ends here, but history goes on. I hope your story becomes part of that history.

New Year's Day 2026, Kim Kyung-jin

← Previous Chapter

See every chapter

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #JensenHuang #NVIDIA #GPU #ArtificialIntelligence #AIChips #AILibrary #Semiconductors #AIIndustry

Support This AI Library Book

If this book was useful, you can support future translations, public editions, and open AI knowledge projects through PayPal.

PayPal

https://paypal.me/kimkjkj

Scan the QR code or open the link above.