2025: The Year the Machines Woke Up

Lucy's Very Own Guide to the Amazing AI Adventures of 2025

By the tweets, interviews, podcasts, and papers of AI builders in 2025*

Edited by Befreed Proprietary Content Engine

BeFreed

2025: The Year the Machines Woke Up

Lucy's Very Own Guide to the Amazing AI Adventures of 2025

BeFreed

Contents

A Letter Just for Lucy

Dear Lucy,

This book was made for you. Not for "kids like you" or "smart seven-year-olds everywhere." Just you. Lucy. The girl who moved to America from China when you were only two years old, carrying a little stuffed panda that you still keep on your bookshelf.

You're the girl who asks questions that make grown-ups pause and think. The one who wants to know why things work, not just that they work. You're the kid who finishes your blueberries before anything else on your plate (smart choice—blueberries are basically little purple balls of deliciousness).

Here's the thing, Lucy: most books about artificial intelligence are written for adults with fancy degrees. Too hard. And most books for kids your age talk about AI like it's magic fairy dust. Too easy.

You deserve better. You deserve a book that treats you like the curious, brilliant person you are—one that explains the real stuff happening in the world right now, in a way that makes

sense to you.

So here it is. The first book in human history written specifically for one person.

That person is you.

Let's go learn about the year the machines woke up.

With wonder and blueberries,

Your Book

PART I

When Computers Learned to Really, Really Think

Chapter 1

The Magic Trick of Showing Your Work

What Happened When AI Started Explaining Its Thinking

Lucy, you know how your teacher sometimes says, "Show your work"?

Like when you're solving a math problem, you can't just write "7" at the bottom of the page. You have to show how you got there. Maybe you write "3 + 4 = 7" so your teacher can see that you actually understand what you're doing.

For a long, long time, computers were really bad at this.

Imagine you had a parrot—a super smart parrot that lived in your room. You could ask this parrot almost any question, and it would

squawk out an answer. "What's the capital of France?" "*PARIS! SQUAWK!*" "What's 847 plus 293?" "*ONE THOUSAND ONE HUNDRED FORTY! SQUAWK!*"

Pretty cool, right?

But here's the weird thing about this parrot: if you asked it how it knew the answer, it would just stare at you with those beady little parrot eyes. It couldn't explain its thinking. It just… knew stuff. Or seemed to.

That's how AI worked for years. You'd ask it a question, it would give you an answer, and nobody—not even the scientists who built it—could really tell you why it said what it said.

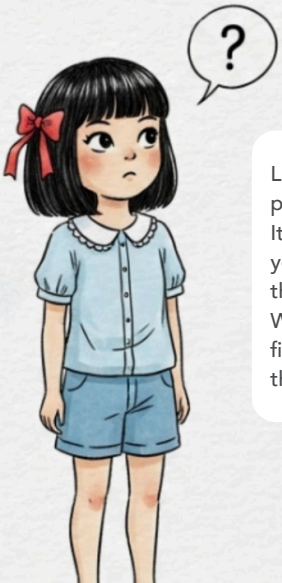

Lucy, imagine asking this parrot how it got the answer. It just squawks numbers at you! But you're smarter than that—you want to know WHY. And in 2025, computers finally learned to explain their thinking too.

Then something amazing happened in 2025.

Scientists figured out a trick. A simple, almost silly trick that changed everything.

They told the AI: "Before you give me your answer, explain your thinking out loud. Walk me through it step by step."

And Lucy, you won't believe what happened next.

The Step-by-Step Revolution

Let's try an experiment. Here's a tricky problem:

There are 23 birds sitting on a tree. 5 birds fly away. Then 3 more birds land on the tree. Then 8 birds fly away. How many birds are on the tree now?

The OLD way AI would handle this:

The computer would sort of… guess? It would look at all the numbers, do some mysterious calculations in its digital brain, and spit out: "13 birds!"

Sometimes it was right. Sometimes it was wrong. And when it was wrong, nobody could figure out where it messed up.

The NEW way AI handles this:

"Okay, let me think about this step by step.

First, there are 23 birds on the tree.Then 5 fly away. So 23 minus 5 equals 18 birds.Then 3 more land. So 18 plus 3 equals 21 birds.Then 8 fly away. So 21 minus 8 equals 13 birds.There are 13 birds on the tree now!"

Same answer! But now we can see exactly how the AI got there. And if it made a mistake somewhere—like if it accidentally said 23 minus 5 equals 17—we could spot it and fix it.

Scientists call this **chain of thought** reasoning. It's like the AI is building a chain, link by link, from the question to the answer. Each link is a tiny step of thinking.


See what you're doing there, Lucy? You're drawing a chain of thought—one golden link at a time. That's exactly how the smartest computers learned to think in 2025. Step by step, just like you.

Lucy, this might sound like a small change. "So what?" you might think. "The computer just talks more now."

But this tiny change turned out to be one of the biggest breakthroughs in the history of computers. Because when AI learned to show its work, it got way, way better at solving hard problems.

It's like the difference between you guessing an answer and you actually understanding how to find the answer. Guessing works sometimes. Understanding works almost always.

The Thinking Machines: o1 and o3

In the fall of 2024, a company called OpenAI (they're the people who make ChatGPT, which you might have heard grown-ups talking about) released something special.

They called it **o1**.

And then, a few months later, they released an even smarter version called **o3**.

What made these AI systems different from anything that came before?

They were built from the ground up to think before they speak.

When you ask o1 or o3 a hard question, something interesting happens. The AI

doesn't answer right away. Instead, it pauses. Sometimes for a few seconds. Sometimes for a whole minute. Sometimes even longer.

During that pause, the AI is actually thinking. Not pretending to think. Actually working through the problem, step by step, checking its own work, trying different approaches.

You can even watch it think! The AI writes out its reasoning as it goes, so you can see exactly what's happening in its digital brain.

A really smart scientist named Ilya Sutskever—one of the people who helped create some of the first powerful AI systems—talked about this in a famous podcast called the Dwarkesh Podcast. He said something that made everyone pay attention:

"These models learned to think before they speak."

Just like your mom probably tells you to think before you say something when you're upset, these AI systems learned to think before they answer. And just like with you, taking that extra moment to think usually leads to much better answers.

Here's the wild part, Lucy: on really hard math tests—the kind that would make most adults cry—o3 started getting scores that beat almost all human mathematicians. On

coding tests that stump professional computer programmers, it was solving problems that seemed impossible.

The AI went from being like a parrot that just squawks out answers to being like a really thoughtful student who carefully works through problems.

The Secret of Giving AI More Time to Think

Lucy, you know how when someone asks you an easy question—like "What color is a banana?"—you answer right away?

"Yellow!"

You don't need to think about it. The answer just pops into your head.

But what about harder questions? Like "Why do we dream?" or "How do airplanes stay up in the sky?"

For those kinds of questions, you might need to think for a moment. Maybe you furrow your eyebrows. Maybe you look up at the ceiling. Maybe you say "Hmmmm…" while your brain works on the problem.


Sometimes the best answers come when you give your brain a little extra time to wander through the clouds.

That extra thinking time helps you come up with a better answer.

Well, scientists discovered that AI works the same way!

They call this idea **test-time compute**. That's a fancy way of saying: let the AI think longer on hard problems.

For years, scientists focused on making AI bigger—more digital brain cells, more connections, more everything. They thought the secret to smart AI was just… more stuff.

But it turns out that's only part of the answer.

The other part is giving AI enough time to actually use its brain. A smaller AI that takes time to think can often beat a bigger AI that rushes through problems.

It's like that story about the tortoise and the hare, Lucy. The hare was faster, but the tortoise won because he kept going steadily. Sometimes slow and thoughtful beats fast and careless.

One scientist did an experiment. He gave the same hard math problem to two different AI systems.

The first AI answered in 2 seconds: **Wrong answer**. The second AI thought for 45 seconds: **Right answer**.

Same-sized brains. Same training. But the one that took time to think got it right.

This discovery was huge because it means we might not need bigger and bigger AI systems. We might just need AI that's better at taking its time and thinking carefully.

Chapter 2

The Surprise from Hangzhou

How a Small Team in China Shocked the World

Lucy, you know how in races, sometimes the person everyone expects to win… doesn't?

Maybe there's a kid on your soccer team who nobody thinks is that good, and then one day, they score the winning goal and everyone's jaws drop to the ground.

Something like that happened in the AI world in January 2025.

And it happened because of a team from China—from a city called Hangzhou, which is famous for its beautiful lakes and tea.

You moved to America from China when you were two, Lucy, so you know that China is a big, important country with lots of smart people. Well, a company in Hangzhou called **DeepSeek** did something that made scientists all over the world gasp.

The Five-Million-Dollar Miracle

Here's what you need to understand first: building a really powerful AI costs a LOT of money.

We're talking about numbers so big they're hard to imagine. The big American companies—like OpenAI and Google—were spending more than 100 million dollars to train their smartest AI systems.

A hundred million dollars!

That's enough money to buy thousands of houses. Or millions and millions of pounds of blueberries. (I'm not saying you should spend $100 million on blueberries, Lucy, but if anyone could eat that many, it's probably you.)

The DeepSeek team looked at these numbers and thought: "There has to be a smarter way to do this."

And they found one.

They built an AI called **DeepSeek R1** that could think and reason just as well as the expensive American AIs. But instead of spending $100 million, they spent about **$5.6 million**.

That's almost 20 times less money!

Imagine if you and your friends wanted to build a rocket ship. Everyone said you'd need $100

to buy all the parts. But your friend figured out how to build an equally good rocket for just $5. That's basically what DeepSeek did.

A researcher named Dylan Patel, who studies AI and computer chips, talked about this on a podcast with a famous interviewer named Lex Fridman. Dylan said it was like "someone building a rocket ship in their garage that worked just as well as NASA's rockets."


What if you could build a rocket that flies just as high—but with way less stuff? The DeepSeek team proved that cleverness can take you across oceans and beyond.

How Did They Do It?

The DeepSeek team was clever in lots of ways, but the most amazing thing they did was create something called **DeepSeek-R1-Zero**.

See that "Zero" at the end? That's important. It means something special.

Usually, when scientists train an AI to think and reason, they show it lots of examples. "Here's how a smart person solved this problem. Learn from that." The AI learns by copying what humans did.

But DeepSeek-R1-Zero didn't get ANY examples. Zero. None. Zilch.

Instead, the scientists set up a kind of game. Every time the AI got a right answer, it got a reward (kind of like how you might get a gold star or a high-five when you do something well). Every time it got a wrong answer, no reward.

Then they let the AI practice over and over and over again. Millions and millions of times.

And here's the magical part: the AI figured out how to reason *all by itself*.

Nobody taught it to check its own work. But it learned to do that.

Nobody taught it to try different approaches when stuck. But it learned to do that.

Nobody taught it to think step by step. But it learned to do that too.

The AI discovered these thinking strategies on its own, just by trying to get more rewards.

Lucy, it's like if someone dropped you on an island with a bunch of puzzle pieces and no instructions. When they came back a year later, you'd not only solved the puzzle, but you'd invented new ways of thinking about puzzles that nobody had ever thought of before.

That's what DeepSeek-R1-Zero did with reasoning. It taught itself to think.

The Day the Stock Market Freaked Out

On January 27, 2025, DeepSeek released their AI for anyone to use. For free!

And Lucy, the world went a little bit bananas.

Within days, more people in America downloaded the DeepSeek app than ChatGPT. It became the number one app in the entire App Store.

Think about that. A Chinese AI, from a company most Americans had never heard of, became more popular than the most famous AI in America. In just a few days.

The grown-ups who buy and sell pieces of companies (that's called the stock market)

started panicking. A company called Nvidia—they make the special computer chips that AI needs to work—lost a HUGE amount of value in one day. Like, hundreds of billions of dollars worth of value. Gone in 24 hours.

Why? Because suddenly everyone realized: maybe you don't need the most expensive computer chips to build amazing AI. Maybe you just need to be clever.

Some people called it a "Sputnik moment."

Do you know what Sputnik was, Lucy? Way back in 1957—long before your parents were born, long before your grandparents were probably grown-ups—the Soviet Union (which was kind of like Russia) launched the first satellite into space. It was called Sputnik, and it shocked Americans because they thought they were winning the space race.

DeepSeek's success felt the same way. America had been ahead in AI. Suddenly, it wasn't so clear anymore.

The Big Idea: Cleverness Can Beat Money

Here's what the DeepSeek story taught the world, Lucy:

You don't always need the biggest, most expensive stuff to do amazing things. Sometimes a clever idea beats a pile of money.

This is true in AI, and it's true in life.

The kid with the fanciest art supplies doesn't automatically make the best art. The soccer team with the expensive uniforms doesn't automatically win the game. The person with the most books doesn't automatically know the most.

What matters is how you use what you have.

The DeepSeek team didn't have as much money or as many fancy computer chips as the big American companies. But they thought harder about how to use what they did have. And they showed the whole world what's possible when you're resourceful and creative.

Chapter 3

When "Bigger" Stopped Being Better

The Plot Twist in the AI Story

For about ten years, Lucy, AI scientists had a pretty simple recipe for making smarter AI:

Step 1: Make the AI bigger (more digital brain cells!)

Step 2: Give it more stuff to learn from (more books, more websites!)

Step 3: Train it longer!

Step 4: Celebrate because it's smarter now!

This recipe worked amazingly well. Every time scientists followed it, the AI got better. It was like a magic formula.

Scientists even had a fancy name for this: scaling laws. "Scaling" means making things bigger. And the "law" was: bigger AI = smarter AI.

For a whole decade, this seemed like the whole game. Want better AI? Make it bigger. Got more money? Make it even bigger. It was like a video game where the only strategy was

collecting more power-ups.

But in 2025, something weird started happening.

Making AI bigger wasn't working as well anymore.

The Pizza Problem

Let me explain it with pizza, Lucy.

Imagine you're eating pizza (your second favorite food after blueberries, maybe?). The first slice is AMAZING. So delicious! The second slice is great too. The third slice is still pretty good.

But by the time you're on your fifth or sixth slice… it's not as exciting anymore, right? You might even feel a little sick.

This is what scientists call **diminishing returns**. Each additional slice of pizza gives you less and less happiness.


The first slice? Amazing! The tenth slice? Even Lucy's not so sure anymore. Sometimes more isn't better—it's just... more.

The same thing started happening with AI.

The first doubling of AI size: HUGE improvement!

The second doubling: Big improvement!

The third doubling: Good improvement.

The tenth doubling: Tiny improvement.

Scientists were spending more and more money to make AI bigger, but they were getting smaller and smaller improvements. Like eating your tenth slice of pizza and barely tasting it.

Ilya's Big Announcement

Remember Ilya Sutskever? The scientist who talked about AI learning to "think before it speaks"?

Ilya is one of the most legendary AI scientists alive. He helped start OpenAI. He worked with Geoffrey Hinton, who some people call the "Godfather of AI." When Ilya says something, people listen very carefully.

In November 2025, Ilya gave a long interview on the Dwarkesh Podcast. And he said something that made the whole AI world stop and pay attention:

"We're moving from the age of scaling to the age of research."

Translation: We can't just keep making AI bigger and expect it to get smarter. We need new ideas.

Ilya compared it to gold mining. When the Gold Rush started, miners could just pick gold nuggets off the ground. Easy! But after a while, all the easy gold was gone. The miners who came later had to dig deeper and work harder to find the same amount of gold.

AI had picked up all the "easy gold" from just making things bigger. Now the hard work of finding new ideas had to begin.

Something AI Still Can't Do

Ilya also said something that surprised a lot of people. He pointed out a problem with even the smartest AI:

"These models somehow just generalize dramatically worse than people."

Okay, that's a big word, Lucy. Let me explain what "generalize" means.

Imagine you learned how to play checkers. You got really good at it! Then someone hands you a chess board and says, "Try this game."

If you're a person, you'd probably figure out chess pretty quickly. The games are different, but some of the ideas are similar. You'd think,

"Okay, these pieces move in certain patterns, I'm trying to capture my opponent's pieces, let me apply what I know…"

You'd **generalize**—use what you learned in one situation to help you in a new situation.

But current AI has trouble with this. If you train an AI to be amazing at checkers, it might be terrible at chess. Even though the games are related, the AI can't easily transfer what it learned.

It's like the AI learned to ride one specific bike on one specific path. You learned to actually ride bikes, so you can ride any bike on any path.

This is a big deal, Lucy. It's one of the reasons why even the smartest AI in 2025 isn't really "intelligent" the way you are. You can learn something in school and then use that knowledge in totally different situations. AI struggles with that.

The Safety-First Scientist

Here's something interesting about Ilya: he takes AI safety very seriously.

So seriously that in 2024, he left OpenAI—the company he helped create—to start a new company. He called it **Safe Superintelligence Inc.**, or SSI for short.

The name tells you exactly what he cares about:

- **Safe:** Making sure AI doesn't hurt anyone
- **Superintelligence:** Building AI that's smarter than humans
- **Inc.:** That just means it's a company

Ilya thinks we need to figure out how to make AI safe before we make it super powerful. Not after. Before.

It's like learning about fire safety before you're allowed to use matches. First you learn the rules, then you get the power.

PART II
Robots That Actually DO Things

Chapter 4

From Talkers to Doers

When AI Learned to Click Buttons and Do Stuff

Lucy, think about all the things you do on a computer or tablet:

- You click on apps
- You type words
- You watch videos
- You play games
- You draw pictures

For a long time, AI could only talk about these things. It could tell you how to do them, but it couldn't actually DO them.

It was like having a friend who's really good at explaining soccer—they know all the rules, all the strategies, all the players' names—but they've never actually kicked a ball.

"To score a goal," this friend would say, "you

need to kick the ball into the net!"

Great advice. But they're not going to score any goals themselves.

In 2025, that changed. AI learned to not just talk, but to actually DO things on computers.

What's an AI Agent?

Scientists came up with a new word for AI that can take action: **agent**.

Think of an agent like a really helpful assistant. Not just one that answers questions, but one that can actually:

- Go online and look things up for you
- Fill out forms
- Book trips for your family
- Organize your files
- Send emails
- Even buy things!

Regular AI: You ask, "How do I order a pizza?" It tells you the steps.

AI Agent: You say, "Order me a cheese pizza." And it actually does it—opens the website, fills in your address, clicks "order"!

Pretty cool, right?

The AI That Learned to Use a Mouse

In 2025, several companies made AI that could control computers.

OpenAI made something called **Operator**. It could open websites, click buttons, fill in forms, and navigate around—just like you do, but without hands!

Anthropic—that's the company that makes an AI called Claude—made their AI learn "computer use." Claude could watch your screen, move the mouse, type things, and complete tasks.

Imagine watching a ghost use your computer. The mouse moves on its own. Windows open and close. Text appears in boxes. Buttons get clicked. But nobody's touching the keyboard or mouse.


Wait—who's clicking that?! When AI learned to use computers, it was a little like having a friendly ghost helper. Spooky, but pretty cool.

That's what these AI agents look like in action. A little spooky, but also really helpful!

The "Almost Perfect" Problem

But here's the thing, Lucy. AI agents in 2025 weren't quite ready to do everything without help.

A really smart scientist named Andrej Karpathy—he used to work on AI at Tesla (the electric car company) and at OpenAI—talked about this in a long interview on the Dwarkesh Podcast.

He said something important:

"It's a decade of agents, not a year."

Some people were getting SO excited about AI agents. They thought by the end of 2025, AI would be doing everything for us! Making dinner! Walking the dog! Doing homework! (Wait, you probably shouldn't let AI do your homework, Lucy.)

But Andrej said: slow down! AI agents are cool, but they're not ready to do everything yet.

He listed five big problems that AI agents still have:

1. **Memory problems:** If you tell an AI agent something today, it might forget by tomorrow. Like a goldfish! You'd have to

keep reminding it.

2. **Seeing problems:** AI can read text really well, but it's still learning to understand pictures and videos as smoothly as you do.
3. **Computer problems:** Using a computer is harder than it looks! AI sometimes clicks the wrong button or gets confused by pop-up windows.
4. **Smartness problems:** For really complicated tasks—like planning a whole birthday party—AI agents just aren't smart enough yet.
5. **Mistake problems:** AI agents mess up about 20% of the time. Would you trust a helper who made mistakes 1 out of every 5 times?

That last one is really important. Imagine if your assistant booked the right flight 8 times out of 10, but the other 2 times they accidentally booked a flight to the wrong city. Would you trust them to book flights without checking?


Oops! Even helpful ghost-assistants make mistakes sometimes. That's why you still need to check their work—and why the AI looks a little embarrassed.

That's where AI agents were in 2025: really impressive, but not quite ready to be trusted completely on their own.

The Universal Plug for AI

Here's a problem that needed solving: there are lots of different AI systems and lots of different apps and websites. How do you get them all to talk to each other?

In 2025, smart people invented something called **MCP**. That stands for "Model Context Protocol."

I know, Lucy, that sounds like a mouthful of alphabet soup! Let me explain it with LEGO.

Imagine you have LEGO bricks, but also building blocks from another company, and also magnetic tiles from yet another company. They're all great toys, but they don't fit together very well. LEGO bricks don't snap onto magnetic tiles.

Now imagine someone invented a special connector piece that could attach to ANY toy and let them all work together. Suddenly, your LEGO castle could connect to your magnetic tile walls and your building block towers!

That's what MCP does for AI. It's like a universal connector that lets different AI systems talk to different apps and websites.

One expert called MCP "the USB-C of AI." You know USB-C? It's that charging cable that works with tons of different devices—phones, tablets, game controllers. One plug that fits everything!

MCP is like that, but for AI connecting to apps. One protocol that works with everything.

The AI Agent from China That Got Bought for $2 Billion

Remember how DeepSeek surprised everyone with their reasoning AI? Well, another Chinese company had a big moment in 2025 too.

They made an AI agent called **Manus** (say it like "MAN-us").

Manus could do some amazing things:

· Look through job applications and pick the best ones

· Plan entire vacations—flights, hotels, activities, everything

· Research complicated topics and write reports

People who tried Manus were amazed. Videos of it working spread all over the internet.

Then, in December 2025, a HUGE company called Meta (they own Facebook and Instagram) bought Manus for **2 billion dollars!**

That's $2,000,000,000!

And Manus was less than a year old! Imagine starting a lemonade stand this summer and having someone offer you $2 billion for it before Christmas. That's basically what happened to Manus.

It showed that big companies think AI agents are REALLY important for the future.

Chapter 5

The Coding Revolution

When AI Became Everyone's Programming Helper

Lucy, do you know what computer code is?

You know how you give instructions to your friends? Like, "First, grab the ball. Then, throw it to me. Then, I'll catch it."

Computer code is like giving instructions to computers. But computers are VERY picky. They need super exact instructions, written in a special language they understand.

If you wanted a computer to add 2 and 3, you might write something like:

result = 2 + 3print(result)

That tells the computer: "Add 2 and 3, and then show me the answer."

People who write code are called "programmers." And for a long time, this was a job only humans could do.

Then 2025 happened.

AI That Writes Code FOR You

Imagine if you could just TELL a computer what you wanted, and it would write its own instructions!

"Computer, make me a game where a frog jumps over obstacles."

And boom—the computer writes all the code to make that game work.

That's basically what happened in 2025.

A company called Anthropic—the same one that made the Claude AI we talked about—released something called **Claude Code**.

What makes Claude Code special is that it doesn't just guess what code you might want. It THINKS about it first.

Remember that "chain of thought" stuff from Chapter 1? How AI got better by explaining its thinking step by step?

Claude Code does this for programming. Before writing any code, it thinks:

"What is this person trying to build?""What's the best way to do it?""What problems might come up?""How should I organize this?"

Then it writes the code. And because it thought first, the code usually works better!

One expert called Claude Code "the most important AI tool of 2025." That's a big deal!

"Vibe Coding"—Programming by Feeling

Andrej Karpathy (remember him?) invented a funny term for this new way of programming. He called it **"vibe coding."**

What's that mean? Instead of carefully writing every line of code yourself, you just describe the "vibe"—the feeling of what you want—and the AI figures out the rest.

"Make me a website that feels cozy and has lots of pictures of puppies."

The AI understands the vibe and creates it!

This was amazing for people who had ideas but didn't know how to code. Suddenly, they could bring their ideas to life just by describing them.

It's like the difference between having to draw a picture yourself versus being able to describe your idea to an amazing artist who draws it for you.

Lucy, someday you might have a cool idea for an app or a game. And you might be able to build it just by describing what you want to an AI helper!

The Numbers Don't Lie

Here's a surprising fact from 2025: about **77% of all new code** was being written with help

from AI.

That's more than three-quarters! Almost all the code being written had AI helping.

Does that mean programmers will lose their jobs? Some people worried about that.

But here's what actually happened: instead of losing jobs, programmers' jobs changed. They became more like directors of code instead of writers of code.

It's like how movie directors don't act in every scene themselves—they guide the actors (in this case, AI) to do their best work.

PART III
AI Gets Eyes, Ears, and a Voice

Chapter 6

Movies Made by Machines

When AI Learned to Dream in Moving Pictures

Lucy, have you ever tried to describe a dream you had?

Maybe you dreamed about flying unicorns or underwater castles or a world made entirely of blueberries. (That last one sounds like a dream YOU might have!)

Dreams are like private movies that play in your head while you sleep. And here's the amazing thing: in 2025, we taught computers to create dreams too.

Well, sort of.

From Words to Worlds

Here's how it works.

AI already learned to create pictures from words. You could type "a cat wearing a top

hat and a monocle" and the AI would paint an image of exactly that—even though no such picture had ever existed before!

But in 2025, AI went way further. It learned to create **videos**.

Not just any videos—realistic-looking videos with things moving smoothly, just like in real life.

OpenAI made a tool called **Sora**. Google made one called **Veo**. You could type a description of a scene, and the AI would create a video of it!

"A golden retriever puppy running through a field of sunflowers at sunset."

Type that in, and watch it come to life. Moving, flowing, beautiful.

The Sound Barrier Breaks

For a long time, AI videos had one big problem: they were silent.

It's like watching a movie with the sound off. You can see what's happening, but something feels wrong. No footsteps when people walk. No splashing when water falls. No birds chirping in the trees.

In 2025, Google's **Veo 3** changed that.

Demis Hassabis—he runs Google DeepMind and won a Nobel Prize (we'll talk about that

later!)—said something memorable at a big conference:

"For the first time, we're emerging from the silent era of video generation."

What's the "silent era"? Way back when movies were first invented, they didn't have sound. Actors would speak, but the audience couldn't hear them—they had to read words that appeared on the screen. It was like watching people mime.

Eventually, someone figured out how to add sound to movies, and everything changed. Movies with sound were SO much better!

Veo 3 was like that moment for AI videos. Suddenly, the videos had sound that matched what was happening on screen.

Dogs barking. Water splashing. Wind blowing. Music playing. People actually talking with real words!

All from just a text description.

The "That Looks Weird" Problem

But AI videos in 2025 still had some funny problems.

Have you ever noticed how sometimes in AI pictures, things look… off? Especially hands?

AI had a really hard time making hands look

right. People in AI pictures might have six fingers, or their thumbs might be on the wrong side, or their fingers might melt into each other like clay.

Why? Because hands are COMPLICATED. They have lots of joints that move in specific ways. Your hand has 27 bones, Lucy! That's a lot for an AI to keep track of.

Scientists kept working on this. The newer AI systems got better at understanding what they called "material physics"—how things actually look and move in the real world.

Water that actually splashes realistically. Fabric that flows the way fabric should. Hair that blows in the wind properly.

AI was learning to understand how the physical world actually works.

The Big Question: Is This a Good Thing?

Here's where things get interesting and maybe a little scary, Lucy.

If AI can make videos that look real, how do you know what's real and what's not?

What if someone made a fake video of a person saying something they never said? Or doing something they never did?

These fake videos are called **"deepfakes"**

(because they're made with "deep learning," a type of AI).

In 2025, people had long conversations about this. Some thought we needed new laws. Some thought we needed special tools to detect fakes. Some thought we should teach everyone to be more careful about believing what they see online.

There's no easy answer. It's a bit like fire—fire can cook your food and keep you warm, but it can also burn things down if you're not careful.

AI video is powerful. We're still figuring out how to use it responsibly.

Dreams for Everyone

But here's the exciting part: AI video also gives EVERYONE the power to create movies!

Before 2025, if you wanted to make a movie, you needed:

- Expensive cameras
- Actors
- Helpers to run the cameras and lights
- Editing software
- Money. LOTS of money.

Now? A kid with an idea can describe their dream, and AI can bring it to life.

"A dragon flying over a snow-covered village at sunset, while children wave from below."

Type that in, and watch it appear.

Lucy, someday you might make movies just by describing what you see in your imagination. That's pretty amazing.

Chapter 7

AI That Talks Like a Person

When Computers Learned to Really Speak

Lucy, close your eyes for a second. (Well, after you read this paragraph!)

Think about all the ways humans communicate:

- We talk with our voices
- We listen to others
- We show pictures and point at things
- We make faces to show how we feel
- We wave our hands around

For a long time, AI was mostly good at reading and writing text. But humans do so much more than text!

In 2025, AI learned to communicate in all the ways humans do.

The "Her" Moment

There's a movie from 2013 called "Her" where a man becomes friends with an AI assistant that sounds like a real person. The AI has a

warm voice, understands jokes, and can have real conversations—not robotic ones.

In 2025, OpenAI released something called **GPT-4o**. (The "o" stands for "omni," which means "all.")

When people first tried it, many said: "This is like the movie Her!"

Why? Because GPT-4o could:

- Talk with a natural-sounding voice (not that robotic computer voice from old movies)
- Understand emotions in your voice (like if you were sad or excited)
- Respond SUPER fast—faster than most humans react!
- Switch smoothly between talking, reading, and looking at pictures

The speed part is important. Regular AI felt like texting—you say something, you wait for a response. GPT-4o felt like actually talking to someone. The response came so fast (232 milliseconds—less than a quarter of a second!) that it felt like a real conversation.

Humans blink in about 300-400 milliseconds. GPT-4o could respond faster than you can blink!

Pictures with Just Words

Google released something called **Nano Banana**. (Yes, really! Tech companies have the silliest names sometimes.)

Nano Banana let you edit photos using just words.

Instead of clicking lots of buttons in complicated software, you could just say:

- "Make the sky more blue."
- "Add a dog next to the tree."
- "Remove that person in the background."
- "Make this look like a painting."

And Nano Banana would understand and do it!

It became "the top-rated image editing model in the world," according to Google.

Lucy, this means that someday, if you take a photo and want to change something about it, you might just be able to tell your computer what you want—like you're talking to a friend—and it'll do it.

No need to learn complicated software. Just describe what you want.

PART IV
AI Makes Amazing Discoveries

Chapter 8

The Prize-Winning Puzzle Solver

How AI Solved a Mystery That Stumped Scientists for 50 Years

Lucy, do you know what a Nobel Prize is?

It's like the gold medal of science. Every year, the most important discoveries in the whole world get a Nobel Prize. Scientists dream their whole careers about winning one. It's the biggest honor in science.

In 2024, something incredible happened: people who built an AI system won a Nobel Prize!

The prize went to the humans who created the AI, but everyone agreed: without the AI, this discovery wouldn't have happened.

The Protein Puzzle

To understand this story, you need to know

about proteins.

You know those blueberries you love, Lucy? Your body breaks them down and turns them into energy using tiny machines called proteins.

Proteins are like little workers inside your body. They do almost everything:

- They help you breathe
- They digest your food
- They fight off germs when you're sick
- They make your muscles move
- They even help you think!

There are hundreds of thousands of different proteins in your body, and each one has a different shape.

Here's the important part: a protein's shape is SUPER important. If a protein has the wrong shape, it can't do its job. Many diseases happen because proteins get folded the wrong way.

Think of origami. You know how you can fold a flat piece of paper into a crane or a frog? The folding instructions determine what shape you end up with.

Proteins are kind of like that. They start as long chains—like a piece of string—and then they fold up into specific shapes. And just like with

origami, the way they fold determines what they become.


Proteins fold themselves into shapes, just like origami paper becomes a crane. For fifty years, scientists tried to crack the code—and then AI figured it out.

Scientists wanted to figure out: if we know what a protein is made of, can we predict what shape it will fold into?

For **50 years**, they tried. And mostly failed.

The problem was just TOO complicated. A protein can fold into trillions upon trillions of different shapes. Even the world's fastest computers couldn't check them all.

It was like trying to solve a Rubik's cube with a million sides.

Enter AlphaFold

Demis Hassabis grew up loving games and puzzles. He was a chess champion as a kid. He designed video games as a teenager!

When Demis grew up, he decided to use his love of puzzles to solve a different challenge: making AI as smart as a human.

He started a company called DeepMind. (Later, Google bought it, and it became Google DeepMind.)

One of the biggest projects at DeepMind was called **AlphaFold**.

AlphaFold's job? Solve the protein folding puzzle.

And by 2025, it had basically done it. Not just for one protein—but for almost EVERY protein

ever discovered. Over **200 million** protein shapes!

In a podcast with a famous interviewer named Lex Fridman, Demis explained why AlphaFold worked:

"Nature's patterns are learnable," he said.

What does that mean? Well, proteins didn't appear randomly. They evolved over billions of years. Evolution has patterns—shapes that work well tend to survive and get passed down.

AI is really, REALLY good at finding patterns. So even though the puzzle seemed impossible, there were hidden patterns in protein shapes that AI could discover.

It's like how most birds have similar wing shapes, Lucy. They evolved that way because that shape is good for flying. AI found the "wing shapes" of proteins—the patterns that nature kept using over and over.

Why Scientists Happy-Cried

When AlphaFold's results came out, some scientists actually cried tears of joy.

Why? Because this changed EVERYTHING.

Before AlphaFold, figuring out ONE protein's shape could take a scientist their entire career.

Months or years of work for just ONE shape.

After AlphaFold, you could get the shape in **minutes**.

This means scientists can now:

· Design new medicines much faster

· Understand diseases they couldn't before

· Create new materials and chemicals

· Even help the environment by designing better ways to break down plastic!

One scientist said it was like going from walking to flying. That's how big the change was.

The Nobel Prize

In October 2024, the phone rang at DeepMind.

Demis Hassabis had won the **Nobel Prize in Chemistry!**


A Nobel Prize for solving the protein puzzle! Demis loved games and puzzles as a kid—just like you, Lucy. Big dreams start small.

It was a historic moment. AI hadn't just helped do science—it had helped do PRIZE-WINNING science.

Demis, the kid who loved video games and puzzles, had used those same skills to solve one of biology's biggest mysteries.

And it showed the whole world: AI isn't just for chatbots and games. It can help us understand the deepest secrets of life itself.

Lucy, maybe someday AI will help scientists understand something about blueberries. Like why they're so delicious. (Okay, probably more important stuff than that. But still!)

PART V
The Big Business Battle

Chapter 9

The Race of the Giants

Who's Winning the AI Competition?

Lucy, imagine AI companies as runners in a giant race.

Some runners have been training for years. Some have more money for equipment. Some have secret training techniques. And every day, the race gets more intense!

Let's meet the runners.

Runner #1: OpenAI

OpenAI is probably the most famous AI company. They make ChatGPT, which you might have heard grown-ups talking about.

In 2025, OpenAI:

- Released GPT-5, their newest and smartest AI

- Raised 6.6 billion dollars! (That's $6,600,000,000!)
- Had nearly 700 million people using ChatGPT
- Made their best AI available for FREE to anyone

That last part was surprising. Before 2025, you had to pay to use OpenAI's best AI. But then they gave everyone access.

Why would they give away something so valuable? To get more people using it! If everyone uses ChatGPT, it becomes the standard that everyone knows.

Sam Altman, who runs OpenAI, thinks AI smarter than humans might arrive by 2027. That's soon!

Runner #2: Anthropic

Anthropic is like OpenAI's cousin. It was started by people who used to work at OpenAI but left because they wanted to focus more on keeping AI safe.

Their AI is called Claude.

In 2025, Anthropic:

- Released Claude Sonnet 4.5, which many said was the best AI for writing code

- Created Claude Code, which we talked about earlier
- Raised 8 billion dollars (even more than OpenAI!)

Dario Amodei, who runs Anthropic, talks a lot about keeping AI safe. He has a phrase he uses: "machines of loving grace." It means AI that helps humans in beautiful ways, not AI that hurts us.

Anthropic's whole idea is: "Yes, let's build powerful AI, but let's be REALLY careful about it."

Runner #3: Google

Google has been working on AI for a LONG time. They actually invented many of the techniques that other companies now use!

Their AI is called Gemini.

Here's a funny thing about Google, Lucy: they invented a lot of the important AI stuff, but other companies got more famous for using it.

It's like if you invented basketball but someone else became the most famous player!

In 2025, Google released Gemini 2.0, and people said it was "actually pretty good." They also made Veo 3 (the video maker with sound) and Nano Banana (the picture editor).

Google is definitely catching up!

Runner #4: Meta

Meta is the company that owns Facebook and Instagram.

Their AI is called Llama. (Yes, like the animal!)

Meta does something different: they give their AI away for FREE.

Wait, why would a company just give away their work?

Think of it like this, Lucy: if everyone in the world uses Meta's AI, then everyone learns how to work with it. It becomes familiar and trusted. Then later, Meta can offer special features that people want to pay for.

A scientist named Yann LeCun works at Meta. He believes in sharing AI with everyone. He thinks it's the only way to make sure AI is safe and good for everyone.

Runner #5: The Chinese Teams

Remember DeepSeek? They're still in the race! And there are other Chinese companies too, like Alibaba with their AI called Qwen.

These teams keep surprising everyone with how fast they're improving.

So Who's Winning?

Honestly? It depends on how you measure.

- Most users? Probably OpenAI (ChatGPT is everywhere)
- Most money raised? Anthropic did really well
- Best for coding? Many say Anthropic's Claude
- Best for science? Google's DeepMind is doing amazing things
- Most sharing? Meta, because they give away their AI

The race isn't over. And in technology races, the leader can change really fast!

Chapter 10

Why Tiny Computer Parts Became Super Important

The Story of the Magic Chips

Lucy, you know what's inside a computer that makes it work?

Chips! Really, really tiny chips.

These chips are smaller than your fingernail, but they have billions of tiny switches inside them. They're like the brain cells of computers.

In 2025, AI chips became some of the most important things in the world. Countries fought over them. Companies paid billions for them. And one company became incredibly powerful because of them.

Jensen Huang and the Chip Kingdom

Jensen Huang started a company called Nvidia in 1993. That's over 30 years ago!

For most of that time, Nvidia made chips for video games. Their chips made games look better—smoother, more colorful, more

realistic.

Then something unexpected happened.

Scientists discovered that the same chips that make games look good are ALSO great for AI!

Why? Because both things involve doing tons of math really fast. Video games need to calculate the color of millions of dots on your screen every second. AI needs to do millions of calculations to think and learn.

Same kind of math, different purpose.

So Nvidia, almost by accident, became the most important company in AI.

By 2025, Nvidia was worth over **3 TRILLION dollars**. That's $3,000,000,000,000. Try counting that many zeros!

That's more money than the total value of most countries.

The Electricity Eaters

Here's something wild, Lucy.

AI uses A LOT of electricity. Like, a LOT.

All those computer chips working hard get very hot. Super hot! Companies have to use special cooling systems—sometimes even dunking computers in special liquid—just to keep them from overheating.

And all those computers running all day and night need electricity. Mountains of electricity.

Some AI data centers (buildings full of AI computers) use more electricity than entire cities!

Satya Nadella, who runs Microsoft, gave a tour of one of their AI data centers. He called it an "AI superfactory."

The numbers were incredible:

- Enough power for over a million homes
- 5 million network connections in just ONE building
- More computing power than all of Microsoft's computers from just a few years earlier

Satya actually joked: "I run a software company. Welcome to the software company." Because it didn't look like a software company. It looked like a power plant!

The Island That Makes Everything

Here's something that might surprise you, Lucy.

Almost every AI chip in the world—whether it's used in America, China, Europe, or anywhere else—is made in one place.

That place is an island called **Taiwan**.

Taiwan is not very big. You could drive across it in a few hours. But on this little island, there's a company called **TSMC** that makes the most advanced computer chips in the world.

Nobody else can make chips as good as TSMC's. Not America. Not China. Not Europe.

This makes some people nervous. What if something happened to Taiwan? What if there was a big earthquake that damaged the factories?

All the AI in the world depends on this one company on this one island. That's why countries are now trying to build their own chip factories, just in case.

that sounds like a real person. The AI has a

PART VI
The Big Questions

Chapter 11

Will Robots Take Over Jobs?

The Question Everyone's Asking

Lucy, when you grow up, you might have a job that doesn't even exist yet.

Think about it: when your parents were kids, there were no "app developers" or "social media managers" because apps and social media didn't exist!

In 2025, grown-ups were asking a big question: "Will AI take my job?"

Let's think about this carefully. Because the answer is more interesting than just "yes" or "no."

Jobs That Changed

Remember how we talked about AI writing code? Let's use that as an example.

Before AI coding tools, a programmer would

write every line of code themselves. Slow, careful work.

After AI coding tools, programmers started doing something different. Instead of writing every line, they:

- Told the AI what they wanted
- Watched what the AI wrote
- Fixed mistakes
- Combined pieces together

Is that "taking their job"? Not exactly. The job still exists—it just changed!

The same thing happened in lots of jobs:

- Artists started using AI to help create ideas
- Writers used AI to help with first drafts
- Customer service workers used AI to help answer questions faster

The Jagged Frontier

A professor named Ethan Mollick wrote a book about working with AI. He came up with a cool idea called **"the jagged frontier."**

Imagine drawing a wobbly line. On one side are things AI is GREAT at. On the other side are things AI is BAD at.

The weird thing? The line is jagged and

unpredictable. AI might be amazing at one thing but terrible at something very similar.

AI is great at:

- Answering factual questions
- Translating languages
- Summarizing long documents
- Writing basic code

But AI is bad at:

- Understanding jokes and sarcasm
- Making decisions about what's right and wrong
- Solving brand-new problems nobody's seen before
- Knowing when it's wrong

The professor's advice: Learn where AI is good and where it's bad. Use AI for the stuff it's great at, and do the other stuff yourself.

New Jobs Appeared!

Here's the good news: when old jobs change or disappear, new ones pop up.

"Prompt engineer" became a real job title in 2025. These are people who are really good at talking to AI—asking questions in exactly the right way to get the best answers.

"AI trainer" is another new job. These people teach AI systems by showing them examples of good and bad answers.


See, Lucy? That's you—teaching your AI friend, building something together block by block. You hold the flag. You're the one in charge.

There are probably jobs we can't even imagine yet, Lucy. Just like nobody in 1995 could have predicted "YouTuber" as a career!

What About Homework?

And YES, we need to talk about homework.

In 2025, teachers had a BIG debate: What do we do about AI and homework?

Some students were using AI to write their essays. Some used AI to solve math problems. Teachers weren't sure what to do.

Most decided: banning AI is silly. Kids will use it anyway—at home, at work, everywhere in their lives. Better to teach them how to use it WELL.

But they also realized: kids need to actually LEARN. If AI does all the thinking, kids won't develop their own brains!

It's like using a calculator, Lucy. Calculators are great tools! But if you never learn what multiplication IS, you'll be in trouble when you need to think mathematically.

The goal is: Use AI as a tool, but make sure you understand what it's doing and why.

Chapter 12

Are We Being Careful Enough?

The Safety Conversations

Lucy, you've heard grown-ups say "With great power comes great responsibility."

AI in 2025 had a LOT of power. So some really smart people spent the year asking: Are we being careful enough?

The Worriers

Some people worry AI might become SO smart that it becomes dangerous.

They imagine situations where AI does something that hurts humans—either by accident or because it has different goals than we do.

This sounds like science fiction! Like a robot movie! But some scientists take it seriously.

Now, most of these worriers don't think the end of the world is definitely coming! They just think: "Hey, maybe there's a small chance something goes wrong. Shouldn't

we be careful about even a small chance of something really bad?"

Teaching AI to Follow Rules

Anthropic, the company that makes Claude, has a special approach to safety. They call it **"constitutional AI."**

A constitution is like a list of rules. America has a Constitution that says things like "everyone gets a fair trial" and "people can speak freely."

Constitutional AI means the AI has its own list of rules it follows:

- Be helpful
- Be honest
- Don't harm people
- Admit when you don't know something

The idea is: if you teach AI good values from the start, it will make good choices later.


Just like writing down important rules, Lucy. Someone has to decide what goes on the list. Someone human. Someone like you.

The Fake Videos Problem

Remember when we talked about AI making videos? Well, that creates a safety problem too.

If AI can make videos of anyone saying anything, how do you trust what you see?

In 2025, fake videos became a real problem:

- Fake videos of politicians saying things they never said
- Fake videos of celebrities doing things they never did
- Fake videos designed to trick people

Some countries made laws requiring labels on AI-created content. Some companies added invisible watermarks (like secret codes) to AI videos so you can tell they're AI-made.

But it's still a game of cat and mouse—people keep finding ways around the rules.

The Other Side

Not everyone is worried, though.

Remember Yann LeCun, Meta's AI scientist? He thinks most of the doom talk is "preposterously ridiculous." (That's a fancy way of saying very, very silly.)

His argument: AI isn't actually that smart. Yes,

it can do impressive things. But it doesn't have goals, doesn't want to take over the world, and is more like a very fancy calculator than a living being.

He points out: We've been worried about new technology forever. People worried that books would destroy people's memory. That television would rot brains. That the internet would end society.

Each time, we adapted. Maybe AI is the same.

The Honest Answer

Here's the truth, Lucy: Nobody knows for sure if AI is dangerous.

Some very smart people think it is. Some equally smart people think it isn't.

What almost everyone agrees on: We should be careful. We should think about what we're building. We should have rules and check them regularly.

Because if there's even a small chance that something could go wrong, it's worth being thoughtful about it.

PART VII
What Happens Next?

Chapter 13

Looking into the Crystal Ball

What Smart People Think Will Happen

Lucy, let's pretend we have a magic crystal ball that shows the future.

We don't really know what's coming. But here's what many experts think might happen in the next few years.

Prediction 1: AI Helpers Everywhere

Within five years, most people who work at computers will have AI helpers.

Not AI that does their whole job—but AI that handles the boring parts. Writing emails. Scheduling meetings. Looking up information.

It'll be like having a really smart assistant that never gets tired.

Prediction 2: AI-Designed Medicine

By 2027, some experts think we'll see the first medicine designed mostly by AI helping patients.

AlphaFold showed AI can understand proteins. The next step is designing new proteins that fight diseases.

Imagine pills designed by computers that understood exactly how diseases work at the tiniest level!

Prediction 3: Robots at Home

Before 2030, robots might start appearing in homes.

Not like in movies—they probably won't look human. But robots that can:

· Clean up messes

· Fetch things

· Help older people get around

· Do simple chores

Tesla (the car company) is working on a robot called Optimus. Other companies are building robots too. In five years, they might be common.

Prediction 4: Most Code by AI

Remember how 77% of code was being written with AI help? By 2030, that could be 90%.

But humans will still be in charge! They'll decide WHAT to build and check that it works correctly.

It's like how architects design buildings, but machines help with the construction. The humans are still making the important decisions.

Prediction 5: We'll Keep Learning

Here's the most important prediction, Lucy: We'll keep learning and adapting.

Every generation faces big challenges. Your grandparents worried about different things. Your parents worry about different things. You'll have AI to figure out.

But here's the thing: Humans are pretty good at facing challenges.

We've always invented new things and then figured out how to use them wisely. Fire. Electricity. Cars. The internet. Each one was scary at first, and each one made life better in the end.

AI might be the same.

Chapter 14

Are We Close to Super-Smart AI?

The Great Debate Over AGI

There's a term you hear a lot in AI discussions: **AGI**.

AGI stands for **Artificial General Intelligence**. It means AI that can do ANYTHING a human can do—not just one thing, but everything.

Current AI is like a specialist. ChatGPT is great at writing. AlphaFold is great at proteins. But neither can do what the other does.

AGI would be like a person—able to learn any skill, understand any topic, switch between tasks easily.

Are we close to AGI? That's the question that kept scientists arguing all through 2025.

The Optimists

Sam Altman, OpenAI's boss, thinks AGI is coming soon. He's said it could arrive between 2025 and 2027.

That's soon! Like, maybe before you finish

elementary school, Lucy!

Why is Sam so confident? He's seen what's happening inside OpenAI's labs. The AI is getting better really fast. Every year brings huge improvements.

The Skeptics

Other people think the optimists are getting ahead of themselves.

Andrej Karpathy gave a long, thoughtful interview about this. His assessment:

"The problems are tractable but still difficult."

"Tractable" means solvable. Andrej thinks we CAN build AGI—it's not impossible. But it's still really hard, and we're not as close as some people think.

He made a list of things AI is still missing:

No continual learning: Humans keep learning new things every day. You learned new words this week, Lucy! You'll learn new things tomorrow. AI mostly stays stuck with what it was trained on.

Too much memorization: AI remembers specific examples too well, instead of understanding deep patterns.

Can't generalize well: Remember how we talked about this? If you learn checkers, you

can figure out chess. AI struggles to apply what it learns to new situations.

His prediction: AGI is probably about a decade away, not 2-3 years.

What Comes After AGI?

Let's imagine we DO build AGI someday. AI that can do anything a human can do.

What comes next?

Some people talk about **"superintelligence"**—AI that's smarter than ALL humans combined.

This is where things get really hard to imagine. What would the world look like with something smarter than all of us?

It could be amazing:

- Curing all diseases
- Solving climate change
- Making enough food for everyone
- Exploring space

Or it could be scary:

- What if it doesn't share our values?
- What if it makes decisions we don't understand?

Most scientists think we have time to figure this out. We're not there yet. But the questions are

worth thinking about now.

Chapter 15

A Letter to Future Lucy

What This All Means for You

Dear Lucy,

You've made it to the end of the book! You now know more about AI in 2025 than most grown-ups.

Let me tell you what I think all this means for you.

You're Growing Up in an Amazing Time

The year 2025 was a turning point. Computers learned to think step by step. They learned to do things, not just talk about things. They learned to see and hear and speak. They helped win Nobel Prizes!

You're seven years old now. By the time you're seventeen, AI will be so much more powerful than anything we saw in 2025. By the time you're twenty-seven, it might be hard to remember what life was like before AI could do so many things.

This is exciting! But it's also a big responsibility.

Your Superpower: Being Human

Here's something important, Lucy.

AI can do lots of impressive things. It can write code and create videos and solve protein puzzles. But there are things AI can't do that you can:

You can care. When your friend is sad, you feel something. You want to help. AI doesn't feel anything—it just predicts what words come next.

You can dream new dreams. AI can only combine things it's seen before. You can imagine things nobody has ever thought of.

You can ask "why." AI answers questions. You ask them. That curiosity—that wanting to understand—is uniquely human.

You can choose what matters. AI doesn't have values. You do. You get to decide what kind of person you want to be.


Some things AI will never have: a panda to hug, a window full of stars to wonder at, and a heart that feels. Those belong to you, Lucy.

Use AI Wisely

As you grow up, Lucy, you'll use AI all the time. It'll help you with homework (but remember to actually LEARN, not just copy!). It'll help you create things. It'll answer your questions.

But remember: AI is a tool. Like a hammer or a paintbrush or a calculator. Tools can be used well or badly. The tool doesn't decide—you do.

Use AI to do things that matter to you. To help people. To create beautiful things. To solve important problems. To make the world a little bit better.

Don't use AI to pretend to be something you're not, or to trick people, or to avoid doing the hard work of learning and growing.

The Future Needs You

The scientists and companies and governments are all working on making AI safe and helpful. But they need help from the next generation—from kids like you who will grow up understanding this technology in ways that older people never will.

Maybe someday you'll be a scientist who makes AI even smarter.Maybe you'll be a teacher who helps kids use AI wisely.Maybe you'll be an artist who creates new kinds

of beauty with AI tools.Maybe you'll be something we can't even imagine yet.

Whatever you become, Lucy, the future of AI isn't decided yet. It's being built right now, by humans who care about getting it right.

You can be one of those humans.

One Last Thing

Remember at the beginning of this book, I said it was made just for you?

That's true. No one else in the world has a book exactly like this one.

But here's a secret: the AI that helped write this book learned from millions of other books that humans wrote. And someday, maybe something you create will help teach future AI.

We're all connected, Lucy. The past and the future. Humans and machines. You and me.

The machines woke up in 2025.

Now it's up to all of us—including you—to help them do good things in the world.

With hope and maybe a few blueberries,

Your Book

Glossary

AI (Artificial Intelligence): Computers that can do things that usually need human thinking.

AGI (Artificial General Intelligence): AI that could do anything a human can do. We don't have this yet!

Algorithm: A step-by-step recipe for solving a problem.

Chain of thought: When AI shows its thinking step by step, like showing your work in math class.

Chip: A tiny piece of material with circuits inside that makes computers work.

DeepSeek: A Chinese AI company that surprised everyone by building great AI for way less money.

Deepfake: A fake video or picture made by AI.

Generalize: Using what you learn in one situation to help with others.

Nvidia: The company that makes most AI

chips.

Nobel Prize: The biggest award in science.

Open source: Software that anyone can see and use for free.

Protein: Tiny machines inside living things that do important jobs.

Reasoning model: AI that thinks step by step before answering.

Test-time compute: Letting AI think longer on hard problems.

A Note for the Grown-Ups

If you're a parent, teacher, or caregiver who read this with Lucy (or another curious child), thank you for engaging with these important topics together.

The year 2025 marked genuine turning points in artificial intelligence—from reasoning models like OpenAI's o1/o3 and DeepSeek's R1, to AI agents, to the AlphaFold Nobel Prize, to serious discussions about safety and employment.

While concepts have been simplified for young readers, the underlying facts and expert perspectives are drawn from real sources, including:

- The Dwarkesh Podcast interviews with Ilya Sutskever, Satya Nadella, and Andrej Karpathy
- The Lex Fridman Podcast with Dylan Patel, Nathan Lambert, and Demis Hassabis
- Research from major AI conferences (NeurIPS, ICML, ICLR)
- Reporting from The Economist, WIRED,

and other technology publications

- Expert commentary from Dario Amodei, Yann LeCun, and other AI leaders

Children growing up today will inherit a world shaped by AI. Understanding these technologies—their capabilities, limitations, and implications—is increasingly important.

The best way to prepare young people isn't to shield them from complexity, but to give them frameworks to understand it.

Thank you for being part of that process.

The End

But really, it's just the beginning.

Distilled from the tweets, interviews, podcasts, and papers of AI builders in 2025

Including but not limited to

DWARKESH PODCAST / LEX FRIDMAN PODCAST / NO PRIORS / THE LOGAN BARTLETT SHOW / NEURIPS / ICML / ICLR / NATURE / ARXIV / MIT TECHNOLOGY REVIEW / DARIO AMODEI / SAM ALTMAN / DEMIS HASSABIS / ILYA SUTSKEVER / ANDREJ KARPATHY / YANN LECUN / FEI-FEI LI / GEOFFREY HINTON / YOSHUA BENGIO / JENSEN HUANG / SATYA NADELLA / REID HOFFMAN / ELAD GIL / SARAH GUO...