

A NEATVE BOOK

Brilliant, Not Perfect

The Future of AI and Agents: Where They Will Shine and Where They Will Still Struggle



Honest

no hype

Plain

simple words

Balanced

the good and the
gaps

Practical

what it means for
you

Brilliant, Not Perfect

The Future of AI and Agents:
Where They Will Shine and Where They Will Still
Struggle

*A calm, plain-language look at what is coming, what is not
and how to get ready.*

Brilliant, Not Perfect: The Future of AI and Agents

First edition, 2026.

© 2026 Neatve. All rights reserved.

This ebook is licensed for the personal use of the buyer only. No part of it may be copied, shared, resold, uploaded or given away, in whole or in part, without written permission from the publisher, except for short quotes in reviews.

This book shares opinions and informed guesses about the future. Nobody can predict the future with certainty and things in AI change quickly. It is not legal, financial, career or medical advice.

Leyla is a made-up example. Claude is a product of Anthropic. ChatGPT is a product of OpenAI. Other names mentioned belong to their owners. This book is independent and is not connected with or endorsed by any of these companies. All illustrations are original.

Contents

	Before We Start	5
PART ONE: WHERE WE ARE		
Chapter 1	A Quick Look at Where We Are Now	7
Chapter 2	From Chatbots to Agents: What Changed	47
Chapter 3	How to Think About the Future Without Getting Fooled	
PART TWO: WHERE AI WILL BE AMAZING		
Chapter 4	Building Software	
Chapter 5	Research, Science and Medicine	
Chapter 6	Office Work and the Paperwork of Life	
Chapter 7	Learning and Teaching	
Chapter 8	Talking, Translating and Customer Service	
Chapter 9	Creativity as a Partner	
Chapter 10	Money, Shopping and Small Business	
Chapter 11	A Day in the Near Future	
Chapter 12	Robots in the Real World	
PART THREE: WHERE IT WILL STILL STRUGGLE		
Chapter 13	Truth and Trust	
Chapter 14	Judgment, Taste and Responsibility	
Chapter 15	The Human Touch	
Chapter 16	Common Sense and Local Knowledge	
Chapter 17	When Agents Talk to Agents	
Chapter 18	The Hidden Costs of AI	
Chapter 19	The Real Risks	
PART FOUR: YOU IN THE AGENT AGE		
Chapter 20	Jobs That Grow and Jobs That Change	

Chapter 21	Skills That Will Matter Most	
Chapter 22	Kids, Schools and Growing Up With AI	50
Chapter 23	Using AI Wisely Starting Today	
Chapter 24	The Simple Scorecard	
Chapter 25	Final Words	

Before We Start

Hi again.

If you've picked up this book you're probably curious about the same thing everyone seems to be talking about. Where is all this AI stuff going? Will it take my job? Will it cure diseases? Will robots be folding my laundry in five years? Or is it all hype that will fade like so many tech trends before it?

I want to have an honest chat with you about that.


Not the scary version you see in headlines. Not the shiny version you hear from people selling something. Just a calm, plain-language look at where AI and agents are likely to be brilliant and where they're likely to keep struggling for a long time.

I'll be upfront about one thing. Nobody knows the future. Not me, not the big tech bosses, not the experts on TV. People have been wrong about technology in both directions for as long as there's been technology. Some things arrived much faster than anyone guessed. Others were "just five years away" for fifty years.

So this book isn't a list of predictions carved in stone. It's more like a map drawn in pencil. I'll show you the patterns I see, explain why I think they matter and give you simple ways to think about the future yourself so nobody can fool you.

By the end I hope you'll feel less anxious and more prepared. The future with AI isn't something that just happens to us. We get to shape how we use it.

Let's start by looking at where we are right now.



Part One: Where We Are

A quick, honest look at today before we look at tomorrow.

Chapter 1: A Quick Look at Where We Are Now

Let me take you back just a few years.

Not long ago, talking to a computer felt clumsy. You'd ask a voice assistant a simple question and it would misunderstand you or read out a random website. Chatbots on company websites were basically menus pretending to be people. "I'm sorry, I didn't understand that. Please choose from the following options."

Then things changed very quickly.

Tools like ChatGPT and Claude showed that a computer could hold a real conversation. It could write an email, explain a hard idea, help with homework, fix a bug in code and summarize a long document. It wasn't perfect. It made mistakes. But it felt different. It felt like talking to something that actually understood what you meant, at least most of the time.

What AI can do today

Here's a rough picture of what everyday AI tools can already do today:

- Write and edit almost any kind of text
- Answer questions and explain things at any level
- Read and summarize long documents
- Look at pictures and describe them
- Have spoken conversations
- Translate between many languages
- Write computer code, sometimes a lot of it
- Search the web and pull together what it finds

- Make images from a description

That's already a lot. A few years ago most people would have called some of this science fiction.

The newer thing: agents

The biggest shift happening right now is the move from AI that *talks* to AI that *does*. These doing-AIs are called agents.

A regular chatbot answers you and stops. An agent takes a goal and works on it step by step. It might search the web, open files, write code, test it, fix mistakes, fill in forms and come back when the job is done or when it needs your help.

Agents are still early. They make mistakes. They get stuck. They sometimes do silly things a human never would. But they're getting better fast and they're already quietly doing real work in offices, software teams and small businesses around the world.

Why it feels like it's moving so fast

People often ask why AI suddenly seems to improve every few months. There are a few simple reasons.

Companies are spending enormous amounts of money on it. Computers built for AI keep getting more powerful. Researchers keep finding clever new tricks. And millions of people using these tools every day helps the companies learn what works and what doesn't.

It's a bit like a snowball rolling downhill. Each improvement makes the next one easier.

But here's something important. Fast progress in some areas doesn't mean fast progress in all areas. Some problems are much harder than others. That's really what this book is about: figuring out which is which.

Chapter 2: From Chatbots to Agents: What Changed

Let me explain the difference between a chatbot and an agent with a simple story.

Imagine you want to plan a weekend trip.

With a chatbot you ask "What are some nice places to visit near me?" It gives you a list. You pick one. You ask "What hotels are there?" It gives you some names, but it might not know current prices. You go to a booking site yourself. You compare. You book. You check the weather. You plan the drive. The chatbot helped, but you did most of the work.

With a good agent you say "Plan me a relaxing weekend trip within three hours' drive, under 400 dollars, somewhere with nature and good food. Show me two options before booking anything." The agent searches, compares places, checks prices, looks at reviews, maybe checks the weather forecast and then comes back with two clear plans. You pick one. It might even help you book it after you say yes.

Same brain underneath. Very different experience.

The three things that make an agent

An agent is basically an AI with three extra things.

Tools. It can use things beyond chat: a web browser, your files, a calendar, a code editor, a spreadsheet, other apps.

Steps. It can plan a job, do one part, look at the result and decide what to do next. It keeps going instead of stopping after one reply.

Memory. It can keep track of what it's done, what it learned and what's left to do, at least for the length of the task.

When you put those together you get something that feels less like a clever search engine and more like a junior assistant.

Why agents are a big deal

Most of the value in work isn't in answering one question. It's in getting a whole job done. Booking the trip. Fixing the software. Filling in the forms. Researching the market. Sorting the emails.

That's why so many companies are excited about agents. If AI can do whole jobs, even simple ones, it changes how work gets done.

Why agents are also tricky

Here's the catch. When an AI only talks, its mistakes are mostly harmless. You read a wrong answer, you notice, you move on.

When an AI takes actions, mistakes matter more. Booking the wrong date. Sending an email to the wrong person. Deleting the wrong file. Spending money it shouldn't.

So the big challenge with agents isn't just making them smarter. It's making them **reliable** and **safe** enough to trust with real actions. Keep that word in mind: reliable. It's going to come up again and again in this book.

Chapter 3: How to Think About the Future Without Getting Fooled

Before we look at where AI will shine and where it will struggle, I want to give you a few simple tools for thinking about the future. These will help you long after you've finished this book, because the headlines will keep coming.

Tool one: separate the demo from daily life

Tech companies love demos. A robot doing a backflip. An AI building a whole app in one minute. A video that looks like it was filmed by a Hollywood crew.

Demos are real, but they're often the best case, not the normal case. The robot might have fallen over forty times before the video you saw. The app might have needed lots of cleaning up afterwards.

So when you see something amazing, ask yourself: "Does this work every time, for normal people, in messy real life?" That question alone will save you from a lot of hype.

Tool two: easy for us isn't easy for AI (and the other way round)

This one surprises people. Things that are hard for humans are sometimes easy for AI. Things that are easy for humans are sometimes really hard for AI.

A computer can beat world champions at chess and write a decent poem in seconds. But folding a crumpled towel, making a cup of tea in a stranger's kitchen or knowing when a friend is quietly upset? Those are still very hard for machines.

Researchers have noticed this for decades. The things our bodies and brains do without thinking, like moving around, handling objects and reading a room, turn out to be incredibly complicated. We just don't notice because we're so good at them.

Tool three: "good enough" beats "perfect"

A lot of technology wins not because it's perfect but because it's good enough and cheap. Early phone cameras were terrible compared with real cameras. But they were always in your pocket, so they won.

AI will spread fastest in areas where "pretty good and very fast" is better than "perfect but slow and expensive." It will spread slowest where one mistake is a disaster.

Tool four: follow the checking

Here's my favorite way to guess where AI will take off. Ask: **how easy is it to check if the AI did a good job?**

If it's easy to check, AI will improve quickly and people will trust it quickly. Code either runs or it doesn't. A math answer is right or wrong. A translation can be checked by someone who speaks both languages.

If it's hard to check, AI will improve more slowly and people will trust it less. Is this the right medical decision? Is this a good long-term business strategy? Is this advice actually wise? These are hard to check even for experts.

Keep this idea in your back pocket. It explains a lot of what's coming.

Tool five: people and rules move slower than tech

Even when a technology works, it takes time for people to trust it, for companies to change how they do things and for laws to catch up. Think about how long it took for online banking or video calls to feel normal for everyone.

So when someone says "AI will replace X in two years," ask whether they're only thinking about the technology or also about people, habits, money and laws. Real change usually needs all of those.




Part Two: Where AI Will Be Amazing

The areas where I think AI and agents are likely to get really, really good.

Chapter 4: Building Software

If I had to bet on one area where AI will be superb, it's writing computer code.



Remember tool four from the last chapter? Follow the checking. Code is one of the easiest things in the world to check. You run it. It works or it doesn't. You can write tests that automatically say "yes, this is correct" or "no, try again." That gives AI a clear way to learn and improve.

It's already happening. Many software developers now use AI coding assistants every day. Some write large parts of their programs with AI help. People who never learned to code are building simple apps just by describing what they want.

What this will look like

In the coming years I expect AI agents to handle more and more of the routine parts of software work. Fixing common bugs. Updating old code. Writing tests. Building standard features like login pages, forms and dashboards. Turning a written description into a working first version.

For small businesses and individuals, this is huge. Things that once needed a team of developers and months of time might take a person with a good idea and a few weekends.

What will still need humans

But software isn't only typing code. It's deciding **what** to build and **why**. Understanding what customers really need. Making trade-offs between speed, cost and safety. Designing big systems that will still work years from now. Taking responsibility when something breaks.


So I don't think software people disappear. I think the job shifts. Less typing, more thinking. Less "how do I write this" and more "is this the right thing to build and is it safe?"

A small warning

When building gets very easy, lots of people build things quickly without checking them properly. That means more apps with security holes, more bugs and more half-finished products. The people who stand out will be the ones who build carefully, not just quickly.

Chapter 5: Research, Science and Medicine

This is the area that gets me most excited and also the one where I want to be most careful not to overpromise.



Science involves a mountain of reading, sorting and testing. There are millions of research papers. No human can read them all. There are huge piles of data from experiments, telescopes, hospitals and labs. There are countless possible combinations to test when looking for a new medicine or material.

These are exactly the kinds of jobs AI is good at. Reading fast. Spotting patterns. Suggesting ideas. Testing lots of possibilities quickly on a computer before scientists try the best ones in a real lab.

Where it's already helping

AI has already helped scientists predict the shapes of proteins, which are tiny building blocks inside living things. That was a problem researchers had struggled with for decades and the work was considered so important that it was recognized with a Nobel Prize. AI is also helping search for new materials, study diseases, read medical images and sift through huge amounts of research.

What I expect

I expect AI to become a normal research assistant in almost every lab. It will help scientists read everything, suggest experiments, analyze results and write up findings. Discovery could speed up in

some fields quite a lot.

In medicine I expect AI to get very good at things like reading scans, flagging warning signs early, helping doctors with paperwork so they can spend more time with patients and helping people understand their own health information in plain words.

Why I'm still careful

Here's the honest part. A new idea from a computer is not a cure. Every new medicine still has to be tested carefully in real people over years to make sure it's safe and actually works. AI can make the early steps faster, but it can't skip the testing. And it shouldn't.

Also, the human body is incredibly complicated. Many ideas that look great on a computer fail in real life. So when you see a headline like "AI discovers cure for X," read it as "AI helped find a promising idea that now needs years of testing." That's still exciting. It's just more realistic.

Chapter 6: Office Work and the Paperwork of Life

Let's talk about the boring stuff. The emails. The forms. The reports. The spreadsheets. The meeting notes. The invoices. The scheduling. The "can you just quickly update this document" requests.

For many people this kind of work eats up a huge chunk of every day. And it's exactly where AI agents are likely to become incredibly useful and quite quickly.

Why this area is ripe

Most office tasks follow patterns. Read this, write that. Take information from here, put it there. Summarize this meeting. Reply to this common question. Check this form for missing details. These are clear, repeatable jobs, often with a right answer that's easy to check.

That's the sweet spot for agents.

What I expect

I expect a lot of office admin to become quietly automatic over the next several years. Things like:

- Meetings that summarize themselves, with action items sent to the right people
- Inboxes that sort themselves and draft replies for you to approve
- Forms and applications that fill themselves in from documents you already have
- Reports that assemble themselves from your data every week
- Invoices and receipts that get read, filed and matched without anyone typing

For you personally, it might mean an assistant that handles the paperwork of life. Comparing insurance options, finding the right form, remembering renewal dates, disputing a wrong charge, organizing your documents. The dull stuff that nobody enjoys.

The human part that stays

Even here, some things stay human. Deciding what matters. Handling the tricky exceptions. Talking to the upset customer. Noticing that something feels off even though the numbers look fine. Making the call when there's no clear rule.


I think the people who do best will be the ones who learn to manage these assistants well. Setting them up, checking their work and stepping in when something unusual comes up.

A thought about jobs

I won't pretend this has no downside. Some jobs that are mostly routine admin will change a lot and some will shrink. That's a real worry and I'll talk about it more honestly in Part Four. But history also shows that when boring work gets automated, new kinds of work tend to appear. The key is to prepare rather than panic.

Chapter 7: Learning and Teaching

Here's an area where I'm genuinely hopeful.



Imagine having a patient, friendly tutor available any time of day. One who never gets tired of your questions, never makes you feel stupid, explains things in five different ways until one clicks and adjusts to exactly what you already know.

For most of history only rich families could afford personal tutors. AI could make something like that available to almost anyone with a phone.

What it could look like

A child stuck on fractions at 9 p.m. gets a step-by-step explanation with little practice questions. A nurse studying for an exam gets quizzed on the parts she keeps getting wrong. A 60 year old learning to use a computer gets calm, simple instructions without feeling embarrassed. Someone learning a new language gets to practice speaking without fear of being judged.

Teachers could get help too: making lesson plans, creating practice exercises at different levels, marking simple work and spotting which students are falling behind.

Why it's not that simple


But learning isn't just information. It's motivation, confidence, friendship, discipline and having someone who believes in you. A great teacher notices when a kid is having a hard time at home. A great coach knows when to push and when to be kind.

There's also a real risk. If an AI always gives the answer, some people will stop thinking for themselves. Learning needs a bit of struggle. The best AI tutors will be designed to guide you to the answer, not just hand it over.

So I see AI becoming a brilliant helper in learning, not a replacement for teachers. The best future is probably one where AI handles the explaining and practicing and teachers have more time for the human parts only they can do.

Chapter 8: Talking, Translating and Customer Service

AI is already very good with language. It's what these systems were built on. So it makes sense that anything involving conversation, writing and translation is an area where AI will keep getting better.



Translation

Real-time translation is getting close to something out of science fiction. You speak in your language, the other person hears theirs. Menus, signs and documents can be translated just by pointing your phone at them.

I expect this to keep improving until language barriers are much smaller in everyday life: travel, business, healthcare and making friends across borders.

Will it be perfect? No. Jokes, poetry, slang, cultural hints and deep emotion are hard to translate even for skilled humans. For important things like legal contracts or medical consent, you'll still want a human expert to check. But for everyday life, the change will be enormous.

Customer service

You've probably already talked to an AI when contacting a company. Some are frustrating. Some are surprisingly good. I expect them to keep getting better at handling common questions quickly: tracking orders, changing bookings, explaining bills and fixing simple problems.

Done right, this is good for everyone. No more waiting on hold for a simple question.


But here's what I hope companies remember. When something goes really wrong, when you're upset, confused or dealing with something serious, you want a human. A company that hides behind AI and makes it impossible to reach a person will lose trust. The best ones will use AI for the easy stuff and make real people easier to reach for the hard stuff.

Voice

Talking to computers out loud is becoming natural. I think many people will use voice more than typing, especially older people, busy parents, people who find typing hard and anyone driving or cooking. That makes AI more accessible to more people, which is a quietly important change.

Chapter 9: Creativity as a Partner

This is one of the most debated areas, so let me be careful and fair.



AI can now make images, music, videos, stories and designs from a simple description. Some of it is impressive. Some of it is bland. A lot of it looks a bit the same.

Where it's genuinely useful

I think AI will be a fantastic creative **partner**. It's great for brainstorming, trying lots of ideas quickly, getting past a blank page, making rough drafts, fixing small problems and doing the boring technical parts of creative work.

A songwriter might use it to try different chord ideas. A small business owner might make a decent logo draft. A writer might use it to check their story for plot holes. A filmmaker might use it to plan scenes before spending money on filming.

Where it struggles

But great art usually comes from somewhere. From a person's life, their pain, their joy, their strange point of view, their taste. AI doesn't have a life. It learns from what already exists, so it tends toward what's common and safe.

That's why so much AI-made content feels a little empty. It's often technically fine but it doesn't say anything new.

I expect that as AI content floods the internet, people will actually value real human creativity **more**, not less. Handmade things got more special when factories made everything cheap. Something similar might happen here.

The fairness question

There are also real, unresolved questions about fairness. Many artists and writers are upset that AI learned from their work without permission or payment. Laws about this are still being argued in courts and governments around the world. I don't think this will be settled quickly and I think it's right that people keep asking hard questions about it.

Chapter 10: Money, Shopping and Small Business

Money is one of the areas where agents will quietly change daily life, probably faster than people expect.

Shopping

Right now when you shop online you search, scroll, compare and read reviews yourself. In the future I expect more people to simply tell an assistant what they need. "Find me a good pair of waterproof walking boots under 150 dollars, in my size, with good reviews from people with wide feet." The agent searches, compares and shows you the best two or three.

Some people will even let agents buy routine things automatically, like household basics, within limits they set.

This sounds convenient and it will be. But it raises interesting questions. If an agent picks what you buy, companies will try hard to influence the agent instead of you. Advertising might start aiming at AI assistants. You'll want to know whose side your assistant is really on.

Personal money

I expect AI to help a lot of people understand their money better. Spotting forgotten subscriptions. Warning you before a bill is due. Explaining a confusing letter from the bank. Helping you build a simple budget. Suggesting when you might be overpaying.

This could be especially helpful for people who never had anyone to teach them about money. Plain-language explanations, available any time, without judgment.

But I'd be very careful about letting AI make big money decisions for you. Investing, borrowing and big purchases involve risk and personal goals. AI can help you understand options. It shouldn't be the one deciding. And it certainly isn't a licensed financial advisor.

Small businesses

This is where I'm most optimistic. Small businesses have always had a disadvantage against big companies. Big companies can afford marketing teams, legal teams, accountants, designers and software developers. A corner bakery can't.

AI agents could level the field a bit. A small business owner could have help with their website, social posts, bookkeeping, customer emails, scheduling and paperwork for a small monthly cost. That means more time for what they actually love: baking the bread, fixing the bikes, cutting the hair.


I think we might see more tiny businesses, run by one or two people with a team of AI helpers, doing work that used to need ten staff.

A word of caution

When everyone has powerful AI, some advantages disappear. If every business has a nice website and good marketing copy, those stop being special. What stays special? Real quality, real service, real relationships and a real story. The human stuff again.

Chapter 11: A Day in the Near Future

Let's play a little game of imagination. I'm going to describe an ordinary day a few years from now. Not a sci-fi fantasy. Just a normal Tuesday, based on where the technology seems to be heading. Some of this might arrive sooner, some later and some might never happen the way I picture it. That's fine. It's a sketch, not a promise.



Meet Leyla. She's 38, works as a project coordinator for a construction company and has two kids.

Morning

Leyla wakes up at 6:45. While she makes coffee she asks her assistant out loud what her day looks like. It tells her she has three meetings, one of which was moved last night and that her son's school sent a message about a trip next week that needs a signature and a small payment.

"Want me to fill in the trip form and put the payment on your card? I'll show you before sending," it says. She says yes. It shows her the form on her phone. She checks it and taps approve. Thirty seconds, done.

It also mentions that her car insurance renews in two weeks and that it found two cheaper quotes with similar cover. It doesn't switch anything. It just puts a short comparison in her inbox for later.

Work

At work, Leyla's company uses a team of AI agents. One reads every delivery note and invoice and matches them to orders. Another watches the project schedule and warns her when something is likely to run late. Another drafts the weekly update for the client.

Around 10 a.m. the schedule agent flags a problem. A steel delivery is delayed, which will push back two other jobs. It suggests three options for reshuffling the week and shows the cost and risk of each.

This is where Leyla comes in. She knows something the agent doesn't: one of the subcontractors is going through a hard time and has been unreliable lately. She picks the option that doesn't depend on him, calls the site manager herself to explain and asks the agent to update everyone.

The agent did the heavy lifting. Leyla made the call.

Afternoon

After lunch she has a meeting with a new client. Her assistant listens (everyone agreed to this at the start) and afterwards sends a summary with action items. She fixes one thing it misunderstood, a joke the client made that the AI took seriously and sends it out.

Her daughter texts that she's stuck on her science homework. Leyla suggests she ask her school's AI tutor, which is set up to guide her with questions rather than just give her the answer. Twenty minutes later her daughter texts "got it!!"

Evening

At home, a small robot vacuum does the floors like they already do today. There's no humanoid robot cooking dinner. Leyla cooks with her kids, which she actually enjoys. Her assistant suggests a recipe based on what's in the fridge and reminds her that her mum's birthday is Saturday.

Before bed she asks it to help plan a surprise for her mum. It asks her a few questions about what her mum loves and comes back with three ideas. She picks one and adds her own twist.

What to notice


Look at what the AI did in Leyla's day. Paperwork. Scheduling. Comparing. Summarizing. Matching. Tutoring. Reminding. Suggesting options.

Now look at what Leyla did. Approved. Judged. Knew the people. Made the phone call. Caught the misunderstanding. Cooked with her kids. Added the personal touch.

That, I think, is the most likely shape of the near future for most people. Not robots replacing us. Assistants taking the boring parts so we have more time and energy for the parts that matter.

Chapter 12: Robots in the Real World

When most people think "AI of the future" they picture robots. Robot helpers at home. Robot workers in factories. Robot delivery. So let's talk about them.



I'm putting this chapter at the end of the "amazing" section for a reason. Robots are coming and in some places they'll be amazing. But they'll arrive more slowly and more unevenly than chatbots did.

Where robots already work well

Robots have worked well for years in controlled places. Car factories. Warehouses. Places where the floor is flat, the lighting is the same, the objects are the same shape and there are no toddlers or dogs running around.

In these places I expect robots to keep getting smarter and more flexible, partly thanks to AI. Warehouses, farms, factories and delivery hubs will likely use more and more of them.

Why homes are so hard

Now think about your home. Stuff everywhere. Different shaped objects. Soft things, slippery things, fragile things. Pets. Kids. Stairs. A cup that could be empty or full of hot coffee. A sock under the sofa.

For a robot, this is chaos. Every home is different and every day is different. Remember tool two from Chapter 3: things that are easy for us, like picking up laundry, are some of the hardest things for machines.

There's also cost and safety. A robot that's strong enough to be useful is strong enough to hurt someone if it makes a mistake. It has to be very, very reliable before people will trust it around their families.

What I expect

I expect humanoid and home robots to keep improving and to show up first in simple, specific jobs, like carrying things, cleaning floors or helping in care homes with simple tasks under supervision. I expect lots of impressive videos. But I'd be cautious about anyone promising a robot that does all your housework cheaply and safely any time soon.

It will happen gradually. Then one day you'll realize it's become normal. That's usually how it goes.




Part Three: Where It Will Still Struggle

The honest part. The places where AI is likely to keep falling short.

Chapter 13: Truth and Trust

If you've used AI for a while you've probably caught it making something up. A fake fact. A book that doesn't exist. A quote nobody said. A link to nowhere. Said with total confidence.



People call this "hallucination." I think it's one of the biggest problems AI will keep wrestling with.

Why it happens

Here's a simple way to understand it. These systems learned by reading enormous amounts of text and getting very good at predicting what words usually come next. Most of the time that produces correct, sensible answers. But sometimes it produces something that **sounds** right without **being** right.

It's a bit like a very well-read person at a dinner party who never likes to say "I don't know." Usually they're great company. Occasionally they confidently tell you something completely wrong.

Will it get better?

Yes, it's already getting better. AI can now search the web and show sources. It can double check its own work. It can say "I'm not sure" more often. Companies are working hard on this because it matters so much.

But I don't expect the problem to disappear completely any time soon. And here's the tricky part: as AI gets right more often, people trust it more and check it less. So the rare mistakes can actually become more dangerous because nobody catches them.

Fakes and trust

There's a bigger trust problem too. AI can now make fake photos, fake voices and fake videos that look real. This will make it harder to know what's true online. Scams will get more convincing. Fake news will be cheaper to make.

I expect we'll all need new habits. Checking sources. Being suspicious of shocking content. Using code words with family for money requests. Trusting known, reliable sources more than random viral posts.

The technology to spot fakes will improve too. But it will be a constant race.

What this means for you

The simple rule for the future: **the more something matters, the more you should check it.** AI for a birthday poem? Don't worry. AI for your health, money, legal rights or anything you'll share publicly? Check it with real sources or real people.

Chapter 14: Judgment, Taste and Responsibility

Here's something AI will likely always find hard: knowing what **should** be done, not just what **can** be done.

Judgment

Life is full of decisions with no clear right answer. Should this company take a big risk? Should this patient try a treatment that might help but might also hurt? Should I take this job and move my family?

These decisions depend on values, feelings, relationships, timing and tiny details that are hard to put into words. AI can lay out options and point out things you missed. That's genuinely helpful. But it doesn't live with the results. You do.

Taste

Taste is knowing what's good. Which design feels right. Which sentence has the right rhythm. Which product idea will make people smile. Which joke will land in this room with these people.

AI has learned what's common. That's not the same as what's good. It can make many options, but choosing the best one and pushing it from "fine" to "wonderful" is still a very human skill. I think taste becomes more valuable in a world full of AI-made "fine."

Responsibility


This might be the most important point in this whole section. **An AI can't be responsible.** It can't be fired, sued, put in jail or feel guilty. It doesn't lose sleep over a mistake.

So in any serious area, like medicine, law, money, building bridges or running a country, a human will need to stay responsible. That means a human needs to understand, check and approve what the AI does. Even if AI does most of the work, someone has to sign their name.

That's not a weakness to fix. It's a feature of how a fair society works. We need someone to hold accountable.

Chapter 15: The Human Touch

Some things we want from other people, not machines. Even if a machine could do them.



Care

When you're sick, scared or grieving, you want someone who truly cares. A nurse who holds your hand. A friend who sits with you. A doctor who looks you in the eye.

AI can help carers with paperwork, reminders and information. That could give them more time for the human part. But the caring itself? I don't think most people will want that from a machine and I don't think they should have to.

Friendship and connection

Some people already talk to AI companions and find comfort in them. For lonely people, a kind voice at 3 a.m. can help and I don't want to dismiss that.

But I worry about AI replacing human connection instead of supporting it. Real relationships are messy. People disagree with you, need things from you, surprise you and grow with you. An AI designed to always agree and always be available isn't the same. It might even make real relationships feel harder by comparison.

I hope the future uses AI to help people connect with each other more, not less.

Trust built over time

A lot of business and life runs on trust between people. You go back to the mechanic who didn't overcharge you. You recommend the accountant who explained things kindly. You buy from the small shop because you know the owner.

That kind of trust is built slowly through real experiences. I think it becomes even more valuable when so much else is automated. People will pay for humans they trust.

Presence

Finally, some things are about being there. A teacher in a classroom. A leader in a crisis. A parent at a school play. A performer on stage. The value is partly in the physical, shared, real-time presence of a person. That's very hard to replace and I'm glad it is.

Chapter 16: Common Sense and Local Knowledge

Here's a funny thing about AI. It can explain quantum physics and then make a mistake a ten year old never would.

Common sense

Common sense is all the little things we know without being told. Water makes things wet. You don't put a laptop in the oven. If someone says "I'm fine" in a certain tone, they're not fine. If the shop is closed on Sunday, you can't pick up your order on Sunday.

AI has absorbed a lot of this from reading and it's much better than it used to be. But it still slips up in odd ways, especially in unusual situations it hasn't seen described much. Agents doing real tasks will sometimes do things that are technically correct and practically silly.

Local and personal knowledge

AI knows a lot about the world in general. It knows very little about **your** world. Your town. Your street. Your workplace politics. Your customer who always pays late but is worth keeping. The unwritten rules of your community. The fact that the "shortcut" road floods every spring.

This kind of knowledge often lives only in people's heads. It's rarely written down anywhere AI could learn it. And it's often exactly what makes the difference in real decisions.


The messiness of real life

Real life is full of exceptions. The form doesn't have a box for your situation. The customer's problem isn't on the list. The rules say one thing but everyone knows it really works another way.

Humans are great at handling these weird edge cases. We improvise. We ask around. We bend a rule kindly when it makes sense. AI agents are getting better, but they tend to be most useful in the normal cases and least reliable in the strange ones. And real life has a lot of strange ones.

Chapter 17: When Agents Talk to Agents

Here's something that sounds like science fiction but is already starting: AI agents talking to other AI agents.



Imagine your assistant wants to book you a table at a restaurant. Instead of filling in a web form, it talks directly to the restaurant's booking agent. They agree on a time in a second. Or your company's buying agent negotiates with a supplier's selling agent about price and delivery.

Why this could be great

A huge amount of time in life and business is spent on back-and-forth. Emails about when to meet. Calls about prices. Forms that ask for the same information you've given ten times before.

If agents can handle a lot of that between themselves, things could move much faster. Less waiting. Less admin. Fewer "just following up" emails.

Why it could get messy

But think about what can go wrong.

Mistakes can spread. If one agent misunderstands something and passes it to another, which passes it to another, a small error can grow. Like a rumor passing through a group of friends.

Tricks and manipulation. An agent negotiating for you might be tricked by a cleverer agent on the other side. Or a website might hide instructions meant to fool visiting agents into doing something they shouldn't. Security people are already worried about this.

Nobody's really watching. When humans do business, there's usually someone who notices when something feels off. When agents talk to agents at high speed, strange things could happen before any person notices.

Who's responsible? If your agent agrees to something bad with another company's agent, who is to blame? You? The company that made your agent? The other side? The law hasn't figured this out yet.

What I expect

I expect agent-to-agent systems to grow, first in simple and low-risk areas like booking appointments, sharing standard information and routine ordering. For bigger things, I expect humans to stay in the loop with clear limits: "Negotiate the price, but don't agree to anything over this amount without asking me."

That phrase, "without asking me," is going to be one of the most important features of good agents.

Chapter 18: The Hidden Costs of AI

When you type a question into an AI assistant, it feels almost free and weightless. Just words on a screen. But behind every answer there's a real, physical world. It's worth understanding because it will shape how AI grows.

Energy and data centers

AI runs in huge buildings called data centers, filled with powerful computers. These computers use a lot of electricity, both to run and to stay cool. Training the biggest AI systems takes enormous amounts of computing power and serving millions of people every day adds up too.

As AI use grows, so does the demand for electricity. This is why you might see news about tech companies building new power plants, signing deals with energy providers or placing data centers near cheap power. In some places, local communities have raised concerns about the strain on power grids and water used for cooling.

Will this get better?

Partly, yes. The cost of running AI for each answer has been falling quickly as chips get better and the software gets more efficient. Smaller AI systems that run on your phone or laptop are getting more capable too, which means some tasks won't need a giant data center at all.

But there's a catch economists have noticed for a long time. When something gets cheaper and more efficient, people often use much more of it. So total energy use can still go up even as each answer

gets cheaper.

The human costs

There are also human costs that are easy to forget. Many AI systems were improved with the help of people who label data, check answers and review harmful content, often for low pay. Artists and writers whose work was used to train systems are asking for fairness. These are real people behind the magic.

Why this matters for the future

These hidden costs will shape where AI grows fastest. Things that are useful enough to be worth the cost will spread. Silly uses might get more expensive or be limited. Countries and companies with access to cheap energy and chips will have an advantage.

As a normal user, you don't need to feel guilty every time you ask an AI something. But it's good to know that AI isn't magic floating in the cloud. It's a real industry with real costs and the choices made about it affect real people and places.

Chapter 19: The Real Risks

A book about the future of AI wouldn't be honest without talking about the risks. I'll keep this calm and practical, not scary.

Jobs

The worry everyone has. Will AI take my job?

My honest view: AI will change almost every job and some jobs will shrink or disappear. Jobs made mostly of routine screen work are the most exposed. At the same time, new jobs will appear, many we can't imagine yet. That's what happened with past big technologies.

But "new jobs will appear" can be cold comfort if you're the one whose job changes. The transition can be painful and it won't be fair to everyone. That's why how governments, companies and communities handle it matters a lot. Training, support and giving people time to adjust are important.

Security and scams

AI makes it easier for bad actors to write convincing scam messages, fake voices and trick people. Agents that can take actions could also be tricked into doing harmful things if they read hidden instructions on websites or in emails. Security will be a big and ongoing challenge.

Privacy

AI works best when it knows a lot about you. That's useful and also risky. Who holds all that information? How is it protected? Who can see it? These questions will matter more and more.

Power


A handful of very large companies build the most powerful AI systems. That concentrates a lot of power. Governments around the world are trying to figure out sensible rules and different countries are taking different approaches. How this plays out will shape everyone's future.

The bigger worries

Some serious researchers worry about much bigger long-term risks from very powerful AI systems that might be hard to control. Others think these worries are overblown. I think it's healthy that people take safety seriously and that some companies and governments are investing in it. You don't have to lose sleep over it, but it's reasonable to want careful people working on it.

What you can do

You can't fix all of this alone. But you can stay informed, protect your own information, be careful with scams, keep learning and support sensible, fair rules. That's more than most people do.



Part Four: You in the Agent Age

What all this means for your work, your skills and your life.

Chapter 20: Jobs That Grow and Jobs That Change

Let's make this personal. What does all this mean for work?

Instead of listing specific jobs that will "survive," which is impossible to predict, let me give you a simple way to look at any job, including yours.

The three questions

Ask these three questions about your work.

How much of my day is routine screen work? Writing standard emails, filling forms, moving information around, making standard reports. The more of your day looks like this, the more AI will change it.

How much of my work is with my hands in unpredictable places? Plumbers, electricians, nurses, chefs, mechanics, carers, builders. Remember how hard the physical world is for robots. These jobs are likely to be among the most stable for a long time.

How much of my work is about people, trust and judgment? Leading a team, calming a patient, negotiating a deal, teaching a class, making the final decision. These parts are likely to become more valuable, not less.

Most jobs are a mix. The routine parts will shrink. The human parts will grow.

Jobs likely to grow

In general I'd expect growing demand for:

- People who can set up, manage and check AI systems
- Skilled trades and hands-on work
- Care work: health, elderly care, childcare
- Roles built on trust and relationships
- People who combine deep knowledge of one field with AI skills
- Anyone who helps others adapt: trainers, coaches, advisors

Jobs likely to change the most

The ones that change most are mostly routine and mostly on a screen. Some parts of data entry, basic customer support, simple content writing, basic translation and routine admin. That doesn't mean these people have no future. It means the smart move is to start adding skills now rather than waiting.

The new kind of worker

I think many jobs will look a bit like being a manager of AI helpers. You'll decide what needs doing, hand parts to AI, check the results, handle the tricky bits and take responsibility for the outcome. The people who get comfortable with that early will have a real advantage.

Chapter 21: Skills That Will Matter Most

If you're wondering what to learn or teach your kids, here's my honest list. Notice that most of these are very human skills.

Asking good questions. AI gives answers. Knowing what to ask and how to ask it becomes a superpower.

Checking and thinking critically. Being able to spot when something's wrong, ask "how do we know?" and look for sources.

Judgment and taste. Knowing what's good, what matters and what to do when there's no clear answer.

Communicating with people. Listening, explaining, persuading, calming, leading. Machines can draft words. Humans still need to connect.

Working with AI tools. Not deep technical skill, just being comfortable using them well, the way we all learned to use email and phones.

Deep knowledge of something. AI knows a bit about everything. Someone who knows a lot about one area, plus how to use AI in it, is very valuable.

Hands-on skills. Fixing, building, cooking, caring, making. Real-world skills are harder to automate.

Learning how to learn. The tools will keep changing. People who enjoy learning new things will do fine. People who want to learn once and stop will struggle.


Staying calm with change. Honestly this might be the most important one. A lot of people will feel anxious. Being able to stay steady, curious and practical is worth a lot.

A simple habit

Here's one small habit I'd suggest. Once a week, spend 20 minutes trying AI for something new in your life or work. Just play. See what it's good at and where it fails. Over a year you'll build a real feel for these tools and that feel will help you spot opportunities and dangers that others miss.

Chapter 22: Kids, Schools and Growing Up With AI

Kids today are the first generation growing up with AI that talks back. That's a big deal and I want to spend a chapter on it, because parents and teachers are asking a lot of questions right now.



The worry about cheating

The first reaction many schools had was panic about cheating. If AI can write an essay in ten seconds, how do you know what a student really learned?

It's a fair worry. But I think schools will adapt, as they did with calculators and the internet. More writing in class. More talking and explaining out loud. More projects where students show their thinking. More tasks where using AI is allowed but students must explain and improve what it produced.

The opportunity

The opportunity is huge. A child who is shy about asking questions in class can ask an AI tutor anything. A child who learns differently can get explanations shaped for them. A child in a place without enough teachers can still get help.

AI could be the patient helper many kids never had.

What kids really need to learn

In a world with AI, I think kids need to learn some old-fashioned things more than ever:

- **How to think for themselves.** Forming an opinion before asking the AI.
- **How to check information.** "How do we know this is true?"
- **How to write and speak clearly.** Because clear thinking shows up in clear words.
- **How to work with people.** Teamwork, kindness and handling disagreement.
- **How to stay focused.** In a world full of distractions, attention is precious.

And some new things:

- **How to use AI well.** Asking good questions, knowing its limits and using it as a helper, not a crutch.
- **How to spot fakes.** Knowing that photos, voices and videos can be made up.

A note for parents

If you're a parent, you don't need to be an expert. Try AI tools together with your kids. Talk about what it gets right and wrong. Laugh at its mistakes. Make it a shared thing rather than a secret thing.

And keep an eye on how they use AI companions or chat apps. Kids deserve real friendships, real play and real conversations. AI should add to that, not replace it.

Chapter 23: Using AI Wisely Starting Today

We've spent most of this book talking about the future. Let's finish Part Four with some simple habits you can start today, so you're ready for whatever comes.

Use it, don't fear it

The people who struggle most with new technology are usually the ones who avoid it until they have no choice. The people who do best are the ones who try it early, in small, safe ways. You don't need to become a tech expert. Just start playing. Ask AI to explain something you've always wondered about. Ask it to help plan a meal. Ask it to rewrite an email. Get comfortable.

Keep your brain in charge

Before asking AI for an answer, take ten seconds to think about what you believe. Then compare. This small habit keeps your own thinking strong. It's a bit like walking instead of always taking the car. You

stay fit.

Check what matters

Make it a rule: the higher the stakes, the more you check. Use AI freely for ideas and drafts. Double check it for facts, money, health, law and anything public.

Protect your information

Be careful about what you share. Don't paste passwords, bank details or other people's private information into AI tools. Check the privacy settings. Use strong passwords and two-step login on important accounts, because scams are getting smarter.

Agree a family code word

This one sounds odd but I really recommend it. Agree on a secret word with close family. If someone ever calls with an urgent money request in a voice that sounds like your child or parent, ask for the word. Voice cloning scams are real and this simple trick beats them.

Keep learning

Spend a little time each week learning something new, with or without AI. A new skill. A new tool. A new idea. The world will keep changing and learning is how you keep up without feeling left behind.


Invest in people

Finally and most importantly, invest in your relationships. Friends, family, neighbors, coworkers. In a world where more and more is automated, human connection becomes more precious. It's also

what helps us get through big changes together.

Chapter 24: The Simple Scorecard

Let's put it all together in one place. Here's my honest, pencil-drawn scorecard for the coming years. Remember, these are informed guesses, not promises.



Where AI is likely to be super and where it will likely still lack.

Area	How good AI is likely to get	Why
Writing code	Excellent	Easy to test and check
Office admin and paperwork	Excellent	Repetitive, clear patterns
Translation	Very good	Huge amounts of language data
Research help	Very good	Great at reading and spotting patterns
Tutoring and explaining	Very good	Patient and available anytime
Customer service (simple)	Very good	Common questions repeat
Creative drafts and ideas	Good	Fast and varied but often bland
Medical support	Good but slow	Must be carefully tested and regulated
Robots in factories and warehouses	Good and growing	Controlled environments
Robots in homes	Slow	Messy, unpredictable, safety critical
Always telling the truth	Still struggling	Mistakes remain hard to remove fully

Area	How good AI is likely to get	Why
Judgment and big decisions	Struggling	Depends on values and responsibility
Taste and originality	Struggling	Learns from what's common
Care, empathy and presence	Struggling	People want people
Local and personal knowledge	Struggling	Often not written down anywhere
Handling weird exceptions	Struggling	Real life is full of surprises

The pattern

Look at the table again and you'll notice a pattern.

AI shines where things are **checkable, repeatable, digital and low-risk when wrong**.

AI struggles where things are **hard to check, unusual, physical, personal and high-risk when wrong**.

If you remember nothing else from this book, remember that pattern. It will help you make sense of every headline for years to come.

Chapter 25: Final Words

We've covered a lot. The amazing parts. The weak parts. The risks. The jobs. The skills.

I want to leave you with a feeling more than a fact.

The future with AI isn't something that simply happens to us like weather. It's something people build, choose and shape every day. Companies decide how they use it. Governments decide the rules. Schools decide what to teach. And you decide how you use it in your own life.

You don't need to be an expert to have a say. You just need to stay curious, stay a bit skeptical and keep your humanity at the center.

AI will probably be one of the most powerful tools people have ever made. It could help us learn more, work less on boring stuff, discover new medicines and connect across languages. It could also be used carelessly or harmfully. Both are possible. Which one wins depends on people.

So here's my hope for you. Use these tools. Learn them. Let them take the boring work off your plate. But protect the things that make you human: your judgment, your care for others, your taste, your relationships and your ability to say "wait, that doesn't seem right."

Those are the things the future will need most.

Thank you for reading. I hope you feel a little more ready for what's coming.

The future is not written yet.

We get to help write it.