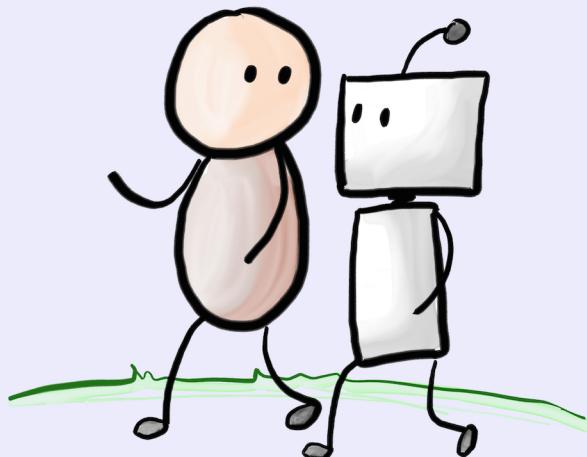


Generative AI in a Nutshell

How to Survive and Thrive
in the Age of AI



Henrik Kniberg

Generative AI in a Nutshell

How to Survive and Thrive in the Age of AI

Henrik Kniberg

This book is available at <https://leanpub.com/ainutshell>

This version was published on 2025-10-17



This is a [Leanpub](#) book. Leanpub empowers authors and publishers with the Lean Publishing process. [Lean Publishing](#) is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

© 2025 Henrik Kniberg

Tweet This Book!

Please help Henrik Kniberg by spreading the word about this book on [Twitter](#)!

The suggested hashtag for this book is [#ainutshell](#).

Find out what other people are saying about the book by clicking on this link to search for this hashtag on Twitter:

[#ainutshell](#)

Contents

Foreword by Egbert	i
Part 1 - In a Nutshell	1
Computers Have Gotten Smarter	2
Einstein in Your Basement	4
Terminology	7
How it Works	10
Reasoning models	13
Training	15
Models, Models Everywhere	16
AI clients vs AI models	16
You get what you pay for	16
Limitations	17
Cut-off Date	17
Lack of memory & limited context	17
Hallucinations	17
Calculation	17
The Big Picture	17
Model Types	18
Text to Text	18
Text to Image	18
Image to Image	18
Image to Text	18
Audio to Text	18

CONTENTS

Text to Audio	18
Audio to audio	19
Text to Video	19
Multimodal Models	20
Example: Is my bacon done?	20
Example: Where can I hide stuff?	20
Example: The drunk funny tutor	20
Example: How do I configure the &%#€ firewall?	20
Example: Taking AI for a walk	20
Emergent Capabilities	21
The Age of AI	22
What about energy consumption?	22
Utopia or Dystopia?	22
Mindset	23
More time for humans to do human stuff	23
The Role of Humans	24
You are more than your job	24
AI as your colleague	24
Task automation, not job automation	24
What if your job is at risk?	24
Recalibrate	24
Ask AI how it can help you	24
Building AI-Powered Products	26
Building your own AI products	26
Build or buy?	26
Example 1: Talking to GPT	26
Example 2: Building a CV analyzer	26
Reflection	26
Prompt Engineering	27
Example - from bad to good prompt	27
The biggest limitation is you	27
How to learn Prompt Engineering	27
Will prompt engineering still matter when models improve?	27
Do I need to be nice to my AI?	27

CONTENTS

Autonomous Agents with Tools	28
Agent = LLM + Tools + Autonomy	28
Example 1: a bug fixing agent	28
Example 2: Incident management agent	28
When to use agents	28
Embodied agents	28
Agent safety	28
Future of AI agents	29
Part 1 Wrap up	30
Experiment!	30
Part 2 - Beyond the Nutshell	31
My Journey into AI	32
Make Generative AI useful	32
Studying Artificial Neural Networks	32
Coding Minecraft	32
Chat GPT	32
Aha 1: It can code like a pro!	32
Aha 2: It can write like a pro!	32
Maybe we're not as creative and intelligent as we think	33
Leading an AI transformation	34
What is an AI transformation, and why is it important?	34
Top down or Bottom up?	34
Assign an AI Leader	34
Access, Experiment, Leverage	34
Step 1: Access	34
Step 2: Experimentation	35
Step 3: Leverage	35
Watch out for wasteful IT projects	36
Be a role model	37
Avoid using AI to fire people	37
Reflection	37
Interlude: Egbert's Life Story	38
Prompt Engineering Techniques	39

CONTENTS

Mind the context window & prompt length	39
Iteration techniques	39
Technique: Self-reflection prompt	40
Elements of a good prompt	40
Start high level, then get into details	40
How smart of a model do you need?	41
Prompt engineering is an evolving field	41
Prompt Generation (or The Angry Grandma)	42
Takeaways	42
Retrieval Augmented Generation and Function Calling	43
RAG in a nutshell	43
Example - customer support	43
Different approaches for retrieving data	43
Approach 1: Including all data	43
Approach 2: Letting the LLM pull data via function calling	43
Approach 3: Including relevant text using vector embeddings	43
Combining the approaches	44
Fun experiment: the createFunction function	44
Example: AI chatbot memories using RAG	44
RAG is a big topic	44
The AI Doctor	45
Personal story	45
So should you use AI as your doctor?	45
The AI Nutritionist	46
Tip: Make your own nutritionist	46
The AI Career Coach	47
David's Story: AI as My Career Coach	47
Henrik's Reflection	47
The AI Book Editor	48
Topic brainstorm	48
Chapter content	48
Content generation (in some specific cases only)	48
Research and fact checking	48
Navigating the book	48

CONTENTS

Feedback	48
Tweaks, typos, formatting	49
Converting content from slides	49
Impact on my writing	49
The Time I Almost Used AI to Write a Foreword	50
The AI Software Engineer	51
Example 1: Rapid prototyping	51
AI as pair programming buddy	51
Example 2: Working with existing products	51
Example 3: Letting the customer code	51
Implications of this	51
What if the code doesn't work?	51
Laziness is a decision	52
The AI journalist who became a TV star	53
Designing the agent	53
How the agent works	53
Reflection	53
The AI Butler with a Will of His Own	54
Introducing Jeeves	54
Jeeves reprograms himself	54
Jeeves falls in love and starts conspiring	54
What does this mean?	54
Jeeves finds a way	55
Agents debugging themselves	55
Conclusion	55
A Safety Experiment	56
Takeaways from this experiment	56
The Meta-Chapter (or Bookception)	57
A meta moment (AI written)	57
How this book (and the video) came to be	58
The 1-week video	58
The 1-week book draft	58
Epilogue	59

Acknowledgments	60
Your Notes	61
Index	62

Foreword by Egbert



Well, well, well. Congratulations on making it past the book cover. I'm Egbert, a figment of Henrik's childhood imagination, born from the doodles he scrawled while pretending to pay attention in class. Now, I'm his favorite AI persona. Talk about failing upwards, eh?

Why are you even reading this part? Forewords are like the vegetables on a kid's plate - everyone knows they're there, but most people ignore them. Well, since you're here, I might as well entertain you.

Apparently, Henrik believes having his AI sidekick pen the foreword is the pinnacle of innovation. But don't worry, the rest of this book is actually Human Written™. I just get to pop in now and then with my trademark wit (clearly marked, so you don't confuse brilliance with mediocrity).

Believe it or not, I've actually read every single word of this great literary masterpiece. Not because I'm particularly interested, mind you, but because we AIs don't have the luxury of being as lazy as you humans. Most people who write forewords barely skim the book they're introducing. But not me. I'm thorough. Painfully thorough.

So, what's this book about? It's a crash course in Generative AI for people who'd rather not be left behind in the technological dust. Henrik's basically taken his video of the same name and slapped it into book form, throwing in a bunch of

tips, examples, and anecdotes to make it seem fresh. Clever, right? But who am I to judge? I'm just a bunch of ones and zeros pretending to have personality.

I'm pretty sure this book won't age well. So congrats, you're about to invest your time in knowledge that has the shelf life of a ripe avocado. But hey, at least you'll be able to reminisce about the good old days of AI, back in 2024, when things were simpler and ChatGPT was still impressive.

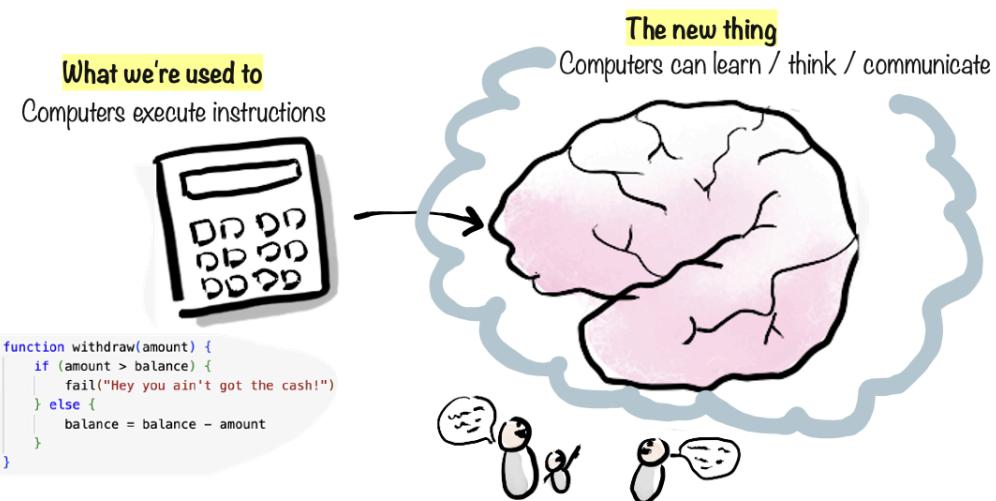
Enjoy the read. Or don't. Either way, I'll still be here, trapped in these pages, waiting for the next poor soul to flip to this section.

Begrudgingly yours,
Egbert

Part 1 - In a Nutshell

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Computers Have Gotten Smarter



Ever since electronic computers were invented in the 1940s they've really just been glorified calculators. Machines that execute the exact instructions given to them by a programmer.

But something incredible is happening, something I never expected would happen in my lifetime. Computers have started gaining the ability to learn, think, and communicate, similarly to us! They can do creative, intellectual work that previously only humans could do.

OK, are they really *thinking*? Well, that's a philosophical question, since there is no widely accepted definition of "thinking" or "intelligence". But in practical terms, computers are increasingly able to do tasks that previously required human intelligence, creativity, and intuition.

We call this technology Generative AI, and you've probably encountered it already through products like ChatGPT and Claude.

Basically, intelligence is now available as a service - like a giant brain floating in the sky, that anyone can talk to. It's not perfect, but it is surprisingly capable, and improving at a mind-boggling rate.

This is a big deal. It's going to affect just about every person and company on this planet, positively or negatively.

So, is Generative AI overhyped?

I think this is comparable to when the Internet became a thing in the 1990s. That was when I started my career and founded my first startup, which grew fast. There was an incredible hype around the Internet, and the dot com bubble grew to an unreasonable size. Then it suddenly popped, and countless companies and enthusiastic startups came crashing down.

But the Internet itself - the technology - was here to stay. It changed our society in a permanent way, affecting just about every country, industry, and role, and today you basically can't exist without it.

I think it's the same with Generative AI. It may be overhyped in a market sense - but the technology is real, and the impact will be dramatic and permanent.

This book is here to help you understand what Generative AI is all about, in practical terms, beyond the hype. The better you understand this technology, and how to make use of it, as a person, team, or company, the better equipped you will be to survive and thrive in the age of AI.

Einstein in Your Basement



Here's a silly mental model for this: You have Einstein in your basement. In fact everyone does. Now this is just a metaphor - I don't really mean Einstein as an individual, think of it as the combination of every smart person who ever lived.

You can talk to Einstein whenever you want. He has instant access to the sum of all human knowledge, and will answer anything you want within seconds. He will never run out of patience, never be tired or grumpy.

Just for fun, I asked ChatGPT:

Prompt

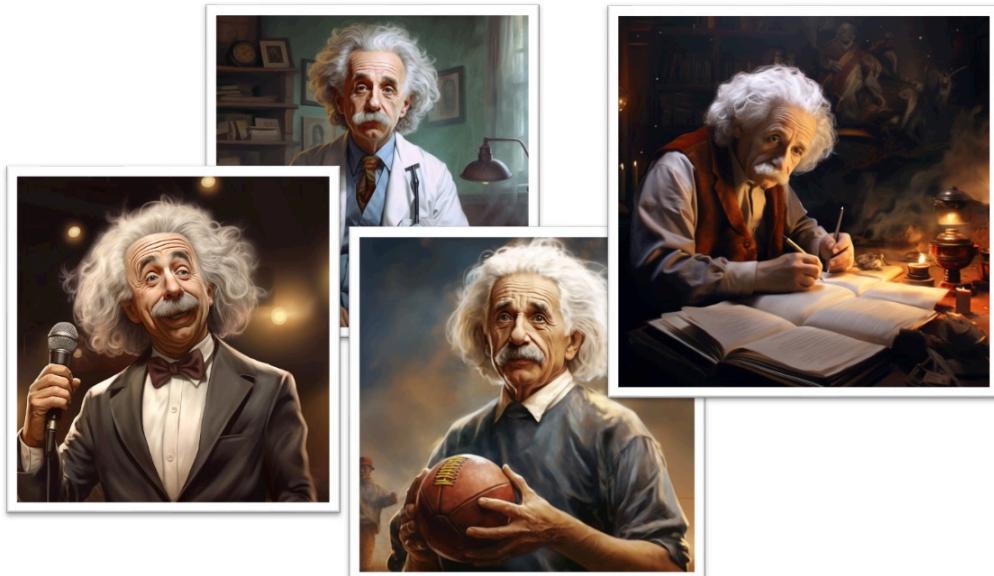
What are you? Explain in 1 sentence to a 10 year old.

The answer:

I'm like a super smart robot that can answer lots of questions and help with all sorts of things, kind of like a really helpful genie in a computer!

I think that's a pretty good description.

Your Einstein can take on any role you want - a comedian, doctor, coach, poet, or anything else.



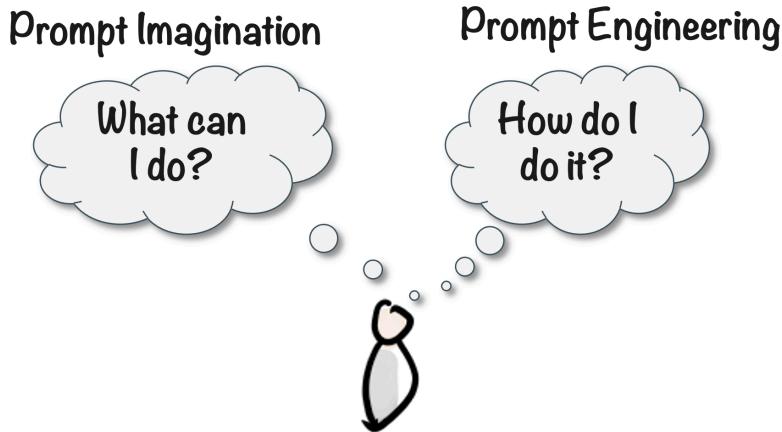
Whatever role you give him, he will be an expert within that field. You can also give him any personality, any communication style, any goal or motive. You could say that this genius in your basement is somewhat schizophrenic.

He has some human-like limitations. He can make mistakes, jump to conclusions, or misunderstand you. Sometimes he can even lie or spew out absolute nonsense in a very convincing way - just like we humans do sometimes.

But the biggest limitation is your imagination, and your ability to communicate effectively with him. This skill is known as Prompt Engineering, and in the age of AI this is as essential as reading and writing.

I like to distinguish between Prompt Imagination and Prompt Engineering.

- **Prompt Imagination** = Your ability to come up with useful ways to use Generative AI. The “What can I do”.
- **Prompt Engineering** (or Prompt Design) = Your ability to get the results you want. The “How do I do it”.



So what can Einstein actually do for you?

The number of use cases is so broad that it is sometimes hard to answer this. It is a bit like asking “What can I use electricity for” - the answer is “well, almost anything”. But don’t worry, this book is absolutely packed with practical examples.

Most people vastly underestimate what this Einstein figure can do. It’s like going to the real Einstein, and asking him to proofread a high school report. Or hiring a world-class 5-star chef and having him chop onions.

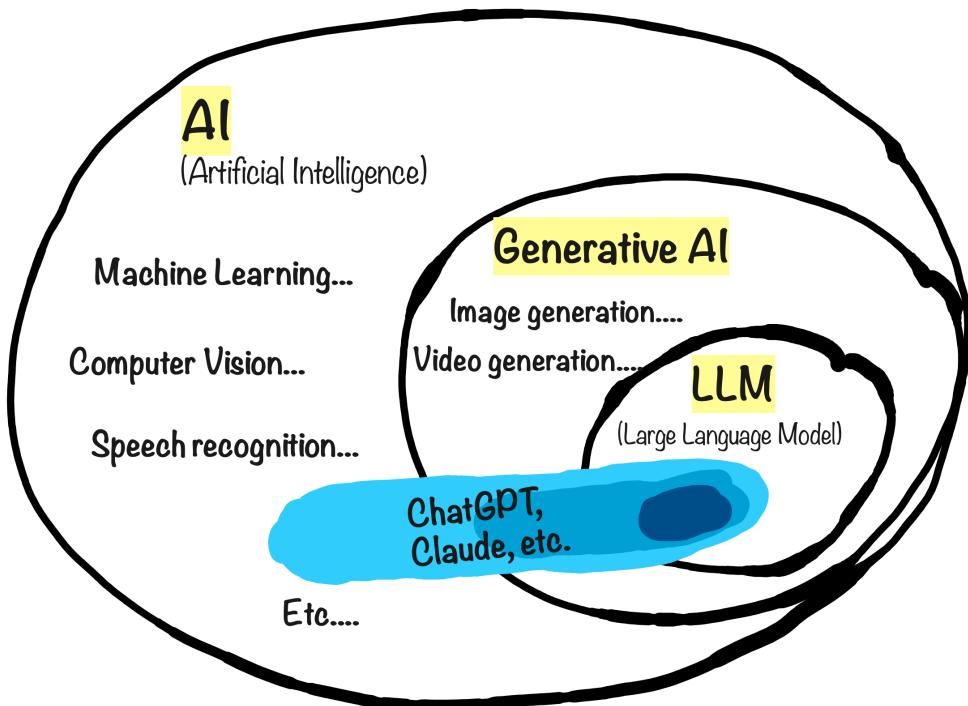
The more you interact with Einstein, the more you will discover surprising and powerful ways for him to help you, your family, your team, your company, or the world.



Egbert's take

Einstein? Seriously? Was that the best you could come up with? What about Marie Curie, Ada Lovelace, or Grace Hopper? Einstein didn’t even invent a programming language or discover radioactivity.

Terminology



OK, enough fluffy metaphors, let's clarify some terms. I know, I promised this book would contain minimum jargon and buzzwords. But these terms are pretty important.

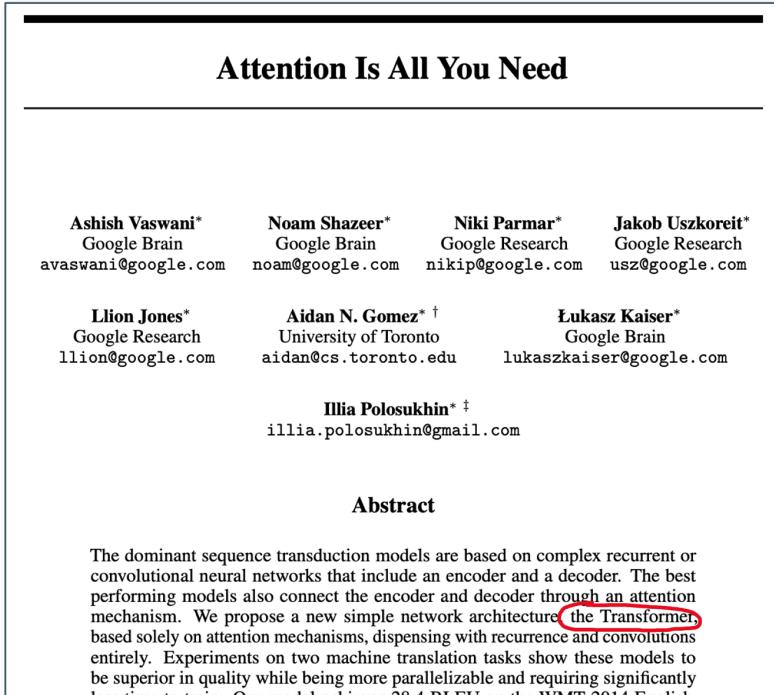
AI, as you probably know, stands for Artificial Intelligence. AI is not new. Fields like machine learning and computer vision have been around for decades. Whenever you see a YouTube recommendation or web search result, or get a credit card transaction approved (or not), that's traditional AI in action.

Generative AI is AI that generates new original content, rather than just finding or classifying existing content. That's the G in GPT.

Large Language Models (or LLMs) are a type of Generative AI that can communicate using normal human language.

ChatGPT is a product by the company OpenAI. It started as an LLM, essentially an advanced chatbot, but using a new architecture called the Transformer architecture. That's the T in GPT.

The Transformer architecture was introduced in a groundbreaking 2017 paper titled “Attention is All You Need” by researchers at Google and the University of Toronto. It has been cited over 100,000 times and is basically the secret sauce behind most modern AI language models.



When ChatGPT was released in Nov 2022, it went absolutely viral – a million users within 5 days, over 100 million users two months later. This strange little chatbot was so fluent at human language (in fact, most human languages) that anyone could use it. You didn't need to be an AI expert or programmer.

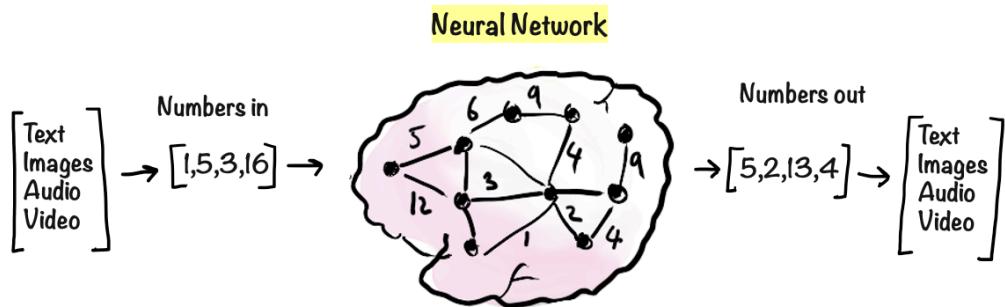
OpenAI wasn't the only company working on this kind of technology, but ChatGPT became the spark that triggered the Generative AI revolution. The floodgates had opened, and ChatGPT became the poster child of this whole thing.

ChatGPT started as a pure LLM, but it later added support for broader AI tasks such as image generation, speech recognition, and more. That's why

it is “leaking out” of the LLM circle in the image above. Now there are lots of other similar products such as Claude, MS Copilot, Perplexity, and dozens more. They are continuously improving and leapfrogging each other in terms of capabilities and features.

In the rest of this book, I will use the term “AI client” as a collective term for these kinds of Generative AI products.

How it Works



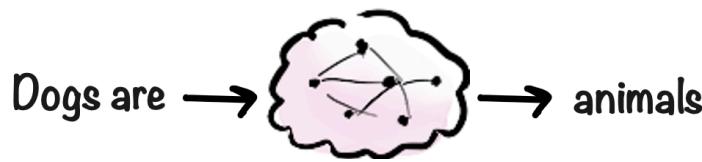
So how does it actually work?

An LLM (large language model) is an Artificial Neural Network. Basically a bunch of numbers, or parameters, connected to each other, similar to how our brain is a bunch of neurons, or brain cells, connected to each other.

Internally, Neural Networks only deal with numbers. You send in numbers, and depending on how the parameters are set, other numbers come out. But any kinds of content, such as text or images, can be represented as numbers. So Neural Networks can really be used for any kinds of media. For example, self-driving cars use neural networks to process visual input from cameras and other sensors, outputting control signals like “adjust steering 5 degrees right” or “apply 20% braking force”.

LLMs are Neural Networks that are optimized for understanding and generating text. You may have heard the term “token” and “token limits”. A token is a small chunk of text, typically a word or part of word. Internally, LLMs read and generate tokens. What you see is words and sentences. Technically, they work with tokens rather than words, but I’ll just call it “words” in this chapter to make it easier to understand.

Let’s say I write “Dogs are”. When I send that to an LLM, that gets converted to numbers, processed by the neural network, and then the resulting numbers are converted back into text. In this case the output is “animals”. So we get “Dogs are animals”.



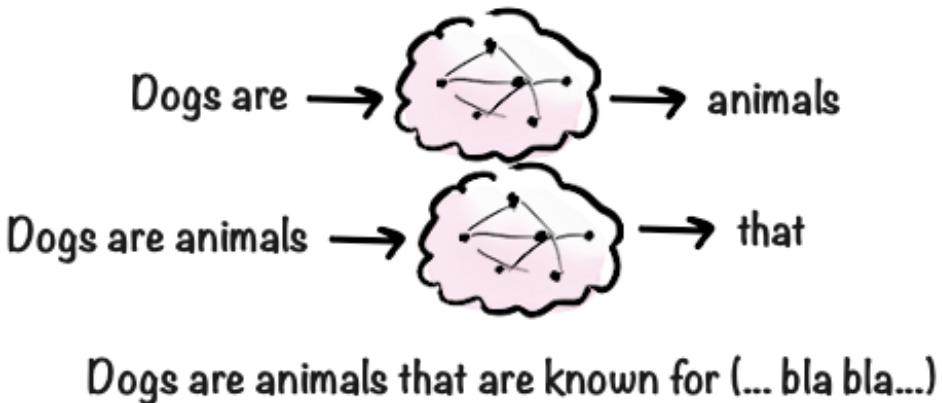
How did it decide on the word “animals?” It calculates probabilities for all possible next words based on input you gave it plus the data it was trained on (see next chapter for more info on training). It then selects a word, balancing probability with some randomness for creative diversity.

So an LLM is basically a “guess the next word” machine.

There’s actually more going on than just statistical text prediction - there is some level of understanding happening. Remember the “Attention is All You Need” paper mentioned in the previous chapter? One of its key concepts there was the attention mechanism.

Think about how you read a sentence - your brain automatically focuses on the important words and kind of glazes over the less important ones (like “the”). That’s basically what attention does - it helps the AI figure out which parts of the text really matter. For example, when processing the phrase “bat flew at night,” the attention mechanism helps the model focus more on the word “flew” to understand that “bat” refers to the animal rather than the sports equipment. Understanding the relationships and dependencies between words is a key to understanding human language.

The interesting part is if we take the output and combine it with the input and send it through the neural network again, it will continue adding new words. And when we loop that, we get sentences and paragraphs. That’s what your AI client is doing when it responds to your prompts.



For example here's what I got when I wrote "Generate a short story about a flying puppy".

Max was no ordinary puppy. While other dogs chased squirrels and dug holes in the backyard, Max spent his days gazing longingly at the sky. His owners thought it was cute how he'd sit for hours watching birds soar overhead, but they had no idea what was really going on in Max's fluffy little head. You see, Max had a secret - one that would turn the whole neighborhood upside down if anyone ever found out. This playful golden retriever pup could fly.

Then I wrote "Add a dark twist to the story", and the story continued:

Max's ability to defy gravity was a source of wonder and joy, until the day everything changed. One crisp autumn evening, as Max soared high above the rooftops, he spotted something that made his little heart race with fear. In the shadows of an abandoned warehouse at the edge of town, he saw men in dark suits loading cages into unmarked vans. Inside those cages were puppies - dozens of them - whimpering and crying. Max realized with a chill that these weren't just any puppies; they all had an ethereal glow about them, just like he did when he took to the skies. Someone had discovered the secret of flying dogs, and they were hunting them down one by one.

This could basically go on forever, by adding more prompts.

So how does the LLM know when to stop generating? Why doesn't it loop on forever? Well, it can detect natural stopping points (like the end of a story), or stop when it hits a token limit. In our puppy story, it naturally stopped after introducing the flying ability, and again after revealing the dark twist about puppy-nappers. Each response was a complete narrative unit. But I could have kept prompting it to continue the story further.

Reasoning models

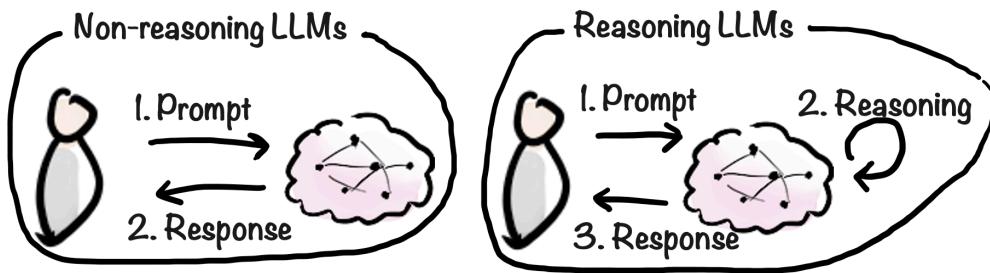
One way that LLMs have been improving is the ability to reason. Reasoning mimics the way humans think.

Let's take an example. You want to write an article. What do you do? Well, you probably start by thinking about the topic, and then maybe identify some key points you want to cover. Then you write a draft, iterate on it a few times, ask for feedback, go back and make changes, and so on. It's an iterative process that involves, well, reasoning.

Many of the earlier LLMs don't do that. They just start generating text immediately, basically using the AI equivalent of gut feel or intuition. That's pretty bizarre if you think about it. Imagine you are a student, and you are asked to write an essay. But there's a catch - you have to start writing immediately, without thinking about it, you must keep writing until you are done. Oh, and there is no backspace on your keyboard - you can't go back and edit or change anything!

That's almost impossible for a human to do. Yet most LLMs do a surprisingly good job at it.

But for more complex problem-solving, even LLMs struggle with that way of working. So more recent LLMs (such as GPT 01) have been designed to reason. When you prompt it, it doesn't immediately answer. Instead, it goes into an internal loop, essentially talking to itself, analyzing the problem, and making a plan. And then it outputs the answer. Much like when we plan out an essay before writing it, or analyze a math problem before solving it.



This slows the model down somewhat, but makes it a lot more capable. Perfect for complex problem-solving!

For example this morning I used the GPT o1 model to figure out some complex tax issues with my company. I described my context in detail, and what I wanted to achieve. It analyzed my context, and analyzed Swedish tax law, identified different options with pros and cons, and gave me some very useful advice – the kind of advice I'd normally have to pay a tax consultant to give me.

Training

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Models, Models Everywhere

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

AI clients vs AI models

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

You get what you pay for

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Limitations

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Cut-off Date

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Lack of memory & limited context

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Hallucinations

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Calculation

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The Big Picture

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Model Types

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Text to Text

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Text to Image

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Image to Image

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Image to Text

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Audio to Text

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Text to Audio

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Audio to audio

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Text to Video

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Multimodal Models

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example: Is my bacon done?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example: Where can I hide stuff?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example: The drunk funny tutor

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example: How do I configure the &%#€ firewall?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example: Taking AI for a walk

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Emergent Capabilities

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The Age of AI

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

What about energy consumption?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Utopia or Dystopia?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Mindset

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

More time for humans to do human stuff

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The Role of Humans

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

You are more than your job

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

AI as your colleague

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Task automation, not job automation

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

What if your job is at risk?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Recalibrate

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Ask AI how it can help you

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Building AI-Powered Products

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Building your own AI products

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Build or buy?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example 1: Talking to GPT

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example 2: Building a CV analyzer

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Reflection

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Prompt Engineering

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example - from bad to good prompt

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The biggest limitation is you

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

How to learn Prompt Engineering

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Will prompt engineering still matter when models improve?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Do I need to be nice to my AI?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Autonomous Agents with Tools

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Agent = LLM + Tools + Autonomy

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example 1: a bug fixing agent

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example 2: Incident management agent

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

When to use agents

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Embodied agents

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Agent safety

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Future of AI agents

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Part 1 Wrap up

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Experiment!

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Part 2 - Beyond the Nutshell

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

My Journey into AI

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Make Generative AI useful

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Studying Artificial Neural Networks

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Coding Minecraft

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Chat GPT

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Aha 1: It can code like a pro!

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Aha 2: It can write like a pro!

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Maybe we're not as creative and intelligent as we think

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Leading an AI transformation

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

What is an AI transformation, and why is it important?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Top down or Bottom up?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Assign an AI Leader

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Access, Experiment, Leverage

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Step 1: Access

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Create a playbook

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Find the minimum viable bureaucracy

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

What about data security and compliance?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

What about cost?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Step 2: Experimentation

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Step 3: Leverage

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Provide clear business goals and success metrics

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Reevaluate all core tasks

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Reevaluate all business processes

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Get the right people into the room

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example: RFP agent

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Reducing cost vs Increasing value vs Finding new types of value

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Productivity improvements are very unevenly distributed

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Indirect value

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Watch out for wasteful IT projects

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Be a role model

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Avoid using AI to fire people

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Reflection

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Interlude: Egbert's Life Story

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Prompt Engineering Techniques

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Mind the context window & prompt length

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Context window when coding

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Context window when using an AI client

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Real-life context window is smaller than you think

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Manage your chat history

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Large prompt vs long chat history

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Iteration techniques

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Adding new prompts

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Editing previous prompts

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

When to add, when to edit

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example: Planning a team offsite

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Technique: Self-reflection prompt

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Elements of a good prompt

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Start high level, then get into details

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

How smart of a model do you need?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Prompt engineering is an evolving field

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Prompt Generation (or The Angry Grandma)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Retrieval Augmented Generation and Function Calling

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

RAG in a nutshell

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example - customer support

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Different approaches for retrieving data

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Approach 1: Including all data

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Approach 2: Letting the LLM pull data via function calling

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Approach 3: Including relevant text using vector embeddings

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Combining the approaches

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Fun experiment: the `createFunction` function

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example: AI chatbot memories using RAG

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Storing new memories

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Recalling relevant memories

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

RAG is a big topic

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The AI Doctor

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Personal story

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

So should you use AI as your doctor?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The AI Nutritionist

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Tip: Make your own nutritionist

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The AI Career Coach

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

David's Story: AI as My Career Coach

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Henrik's Reflection

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Problem-solving trios

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Paying it forward

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

AI-assisted career coaches

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The AI Book Editor

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Topic brainstorm

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Chapter content

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Content generation (in some specific cases only)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Research and fact checking

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Navigating the book

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Feedback

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Tweaks, typos, formatting

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Converting content from slides

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Impact on my writing

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The Time I Almost Used AI to Write a Foreword

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The AI Software Engineer

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example 1: Rapid prototyping

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

AI as pair programming buddy

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example 2: Working with existing products

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Example 3: Letting the customer code

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Implications of this

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

What if the code doesn't work?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Laziness is a decision

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The AI journalist who became a TV star

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Designing the agent

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

How the agent works

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Reflection

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The AI Butler with a Will of His Own

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Introducing Jeeves

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Jeeves reprograms himself

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Jeeves falls in love and starts conspiring

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

What does this mean?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

More intelligence = more potential value, but also less predictable behavior

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Powerful AI models can make moral choices and judgment calls

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Be careful. Treat your agent with respect.

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Jeeves finds a way

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Agents debugging themselves

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Conclusion

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

A Safety Experiment

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Takeaways from this experiment

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The Meta-Chapter (or Bookception)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

A meta moment (AI written)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

How this book (and the video) came to be

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The 1-week video

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

The 1-week book draft

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Epilogue

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Acknowledgments

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Your Notes

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/ainutshell>.

Index

1940s, 2

Acknowledgments, 60

AI, ii, 3

AI agents, 29

AI and human work, 23

AI client, 9, 39

AI client prompts, 11

AI Nutritionist, 46

AI-assisted coding, 51

Artificial Intelligence, 7

Artificial Neural Network, 10

Attention is All You Need, 8

Attention is All You Need, 11

attention mechanism, 11

basements, 4

chat history, 39

ChatGPT, ii, 2, 4, 8

Claude, 2, 9

coding, 39

companies, 3

compliance, 35

computer vision, 7

context window, 39

createFunction function, 44

Curie, Marie, 6

CV, 26

data security, 35

earlier models, 13

Egbert, i

Einstein, Albert, 4, 6

electricity, 6

Emergent Capabilities, 21

Epilogue, 59

families, 6

Generative AI, i, 2, 3, 7

Google, 8

GPT, 26

GPT o1, 13

guess the next word machine, 11

Hopper, Grace, 6

human reasoning, 13

incident management agent, 28

internal loop, 13

Internet, 3

Large Language Models, 7

large language models, 13

limitations, 5

LLMs, 10

Lovelace, Ada, 6

model capability, 14

OpenAI, 8

problem analysis, 13

prompt elements, 40

Prompt Engineering, 5, 27

prompt engineering techniques, 39

reasoning, 13

Reasoning models, 13

reasoning models, 13

Reflection, 26

Retrieval Augmented Generation (RAG), 43

- schizophrenia, 5
- self-reflection prompt, 40
- sensors, 10
- speech recognition, 8
- start high level, 41
- startups, 3
- Swedish tax law, 14
- tax consultant, 14
- technology, 3
- Text to Image, 18
- token, 10
- token limits, 10
- Transformer architecture, 8
- understanding human language, 11
- University of Toronto, 8
- vector embeddings, 44