


THE BEGINNER'S GUIDE TO AI

Understanding Artificial Intelligence,
How LLMs Work, What They Can Do,
Their Limitations, and How to Use Them
Effectively



ANDREW CAVENDISH

The Beginner's Guide to AI

Understanding Artificial Intelligence, How LLMs Work, What They Can Do, Their Limitations, and How to Use Them Effectively

Steve Publications

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

This version was published on 2026-09-19



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.

© 2026 Steve Publications

Contents

- Understanding Artificial Intelligence, How LLMs Work, What They Can Do, Their Limitations, and How to Use Them Effectively 1**
- Introduction 2**
- Chapter 1: What Is Artificial Intelligence, Really? 4**
 - The Promise and the Confusion 4
 - Why Every Definition of AI Is a Moving Target 5
 - The Real Core: Pattern Recognition at Scale 5
 - What AI Is Not: Consciousness, Understanding, Intent 6
 - Why the Confusion Persists 7
 - What This Book Will Help You Understand 8
- Chapter 2: From Rule-Based Programs to Learning Machines 10**
 - Early Dreams: The Birth of the Field in the 1950s 10
 - Rule-Based Systems and Expert Systems 11
 - The Limits of Hand-Crafted Rules 12
 - The First AI Winters 13
 - The Expert Systems Boom and the Second Winter 14
 - The Data Revolution: More Compute, More Data, More Patience . . . 14
 - The Deep Learning Breakthrough 15
- Chapter 3: The Family Tree of AI Terms 18**
 - The Nesting Doll Model: How These Terms Relate 18
 - Artificial Intelligence as the Umbrella 18
 - Machine Learning and Its Variants 18
 - Deep Learning and Neural Networks 18
 - Generative AI and Large Language Models 18
 - Specialized Domains: NLP, Vision, Robotics, Agents 18
 - Putting It All Together 19

Chapter 4: The Training Ground: How LLMs Learn from Data 20

Where the Data Comes From: The Internet and Beyond 20

Cleaning and Filtering Training Data 20

What a Model “Sees” During Training 20

Learning as Pattern Detection, Not Reading 20

The Scale of Modern Training Datasets 20

Data Quality, Biases, and What Data Shapes the Model 20

Chapter 5: Words Become Numbers: Tokens and Embeddings 22

How Computers Understand Text: The Necessity of Numbers 22

Tokenization: Splitting Words into Pieces 22

Why Tokens Matter for Context Limits and Cost 22

Embeddings: Turning Tokens into Vectors of Meaning 22

Similarity in Vector Space: Why Related Ideas Cluster Together 22

What Embeddings Capture and What They Miss 22

Chapter 6: The Neural Network: Brains Made of Math 24

The Neuron Metaphor and Its Limits 24

Layers and Computation: How Data Flows Through 24

Parameters: The Knobs the Model Adjusts 24

Learning by Minimizing Errors 24

Why Deep Networks Are Powerful 24

The Meaning of Model Size and Scaling 24

Chapter 7: Attention and Transformers: The Architecture That Changed Everything 26

The Problem with Earlier Models 26

The Transformer Revolution 26

Attention Explained: Weighing What Matters 26

Self-Attention and Context 26

Why Transformers Handle Language So Well 26

The Recipe: Tokens Plus Embeddings Plus Attention Plus Layers 26

Chapter 8: From Pretraining to Polishing: Fine-Tuning and Alignment 28

Pretraining: Learning Language from Raw Text 28

What the Model Learns During Pretraining 28

Supervised Fine-Tuning: Teaching Helpful Behavior 28

Reinforcement Learning and Human Feedback 28

Reward Models and Preference Optimization 28

Why Raw Pretrained Models Are Not Chat-Ready 28

Chapter 9: When You Hit Send: How Inference and Generation Work 30

 The Difference Between Training and Inference 30

 Next-Token Prediction: The Basic Task 30

 Autoregressive Generation: Building Responses Word by Word 30

 Temperature and Sampling: Controlling Creativity 30

 Why Generation Is Not Database Lookup 30

 Speed, Cost, and the Context Window 30

Chapter 10: The Real Strengths of Modern AI 32

 Language Understanding and Generation 32

 Summarization and Information Synthesis 32

 Translation and Cross-Lingual Capability 32

 Code Generation and Technical Assistance 32

 Creative Support and Idea Exploration 32

 Speed, Scalability, and Accessibility 32

 Putting It Together 33

Chapter 11: The Real Limitations of Modern AI 34

 Hallucinations: What They Are and Why They Happen 34

 Confidence Without Correctness 34

 Arithmetic and Systematic Reasoning Failures 34

 Biases and Incomplete Perspectives 34

 The Knowledge Cutoff and Outdated Information 34

 Prompt Injection and Security Concerns 34

 Putting It Together 35

Chapter 12: What LLMs Do Not Know, Do Not Feel, and Do Not Understand 36

 The Fluency Trap: Why Sounding Smart Is Not Being Smart 36

 Pattern Matching Versus Understanding 36

 The Chinese Room and Similar Arguments 36

 Knowledge, Beliefs, and Truth 36

 Consciousness, Feeling, and What We Know and Do Not Know 36

 What We Can and Cannot Say About Model “Intelligence” 36

Chapter 13: How to Talk to AI: Prompting That Actually Works 38

 What a Prompt Is and How the Model Interprets It 38

 Clarity: The Most Important Rule 38

 Providing Context and Constraints 38

 Giving Examples: Few-Shot Prompting 38

 Asking for Structured Outputs 38

Common Prompting Mistakes and How to Fix Them	38
Chapter 14: Working with AI on Complex Tasks	40
Decomposition: Breaking Big Tasks into Steps	40
Iterative Refinement: Prompting, Reviewing, Revising	40
Self-Critique: When It Helps and When It Does Not	40
Chaining and Multi-Step Workflows	40
Combining Multiple Tools and Techniques	40
When to Stop Delegating to AI	40
Chapter 15: Tools, Retrieval, and AI Agents	42
The Limits of Training-Time Knowledge	42
Tool Use: Giving AI Functions to Call	42
Retrieval-Augmented Generation (RAG)	42
AI Agents and Autonomous Action	42
The Trade-Offs: Flexibility Versus Reliability	42
Chapter 16: Building AI Into Your Workflow	43
Email, Communication, and Administrative Tasks	43
Writing and Content Creation	43
Research and Information Gathering	43
Programming and Technical Work	43
Learning and Education	43
When NOT to Use AI (or When to Be Extremely Cautious)	43
Chapter 17: How to Evaluate AI-Generated Information	45
Fluency, Plausibility, and Accuracy: Not the Same	45
Spotting Hallucinated Sources and Citations	45
Cross-Checking Key Facts	45
Using AI to Help Verify Itself	45
When AI Output Needs Expert Review	45
Building Verification Habits	45
Chapter 18: AI and Work: Transformation, Not Replacement	47
The Reality of Automation: Tasks, Not Jobs	47
What AI Can Realistically Replace	47
Augmentation: Working with AI Rather Than Against It	47
Who Benefits and Who Does Not	47
Skills That Matter More and Skills That Matter Less	47
Career Strategy in an AI-Shaped Economy	47

Chapter 19: Misinformation, Deepfakes, and the Future of Truth 49
 The Scale Problem: AI-Generated Content at Industrial Volume 49
 Deepfakes and Synthetic Media 49
 Misinformation and Manipulation 49
 Authentication and Provenance 49
 The Trust Crisis 49
 What Individuals and Society Can Do 49

Chapter 20: Power, Ethics, and the Questions We Have Not Solved . . . 51
 Who Controls AI and Why It Matters 51
 Environmental Costs of Training and Running Models 51
 Copyright, Data Rights, and Intellectual Property 51
 Surveillance, Privacy, and Civil Liberties 51
 Fairness, Bias, and Harm 51
 Regulation and Governance 51

Chapter 21: The Next Steps: What Is Emerging Now 53
 Scaling and Specialization Trends 53
 Multimodal Systems 53
 Reasoning Models and Chain-of-Thought 53
 On-Device and Local AI 53
 Robotics and Embodied AI 53
 Reliability and Verifiability Research 53

Chapter 22: Living with AI: A Practical Conclusion 55
 Your Mental Model of AI: The Key Takeaways 55
 Using AI Wisely: Principles Over Tricks 55
 Knowing When to Trust, When to Verify, When to Ignore 55
 Continuing Your Education About AI 55
 Being a Responsible User and Citizen 55
 The Human Center: AI as a Tool, You as the Purpose 55

References 57

Understanding Artificial Intelligence, How LLMs Work, What They Can Do, Their Limitations, and How to Use Them Effectively

This book takes you from knowing nothing about artificial intelligence to understanding how modern AI systems work, what they are genuinely capable of, where they fail, and how to use them wisely in your daily life and work. It avoids hype and fearmongering while giving you a clear mental model of the technology, the vocabulary to talk about it, and practical skills for using it responsibly.

Introduction

Three years ago, most people had never heard of a large language model. Today, almost everyone has chatted with one. The speed of this change is disorienting. Headlines oscillate between visions of AI that will solve all our problems and warnings of machines that will take every job and perhaps end civilization. Both extremes get something right and something wrong. Neither is very useful for a person who just wants to understand what this technology actually is and figure out how to deal with it.

This book is for that person. It assumes you are intelligent, curious, and willing to learn. It does not assume you have studied computer science, know statistics, or have ever written code. The only prerequisite is a willingness to pay attention.

The promise of this book is straightforward. By the time you finish it, you will know what artificial intelligence really means and what it does not mean. You will understand, in plain language, how modern AI systems like ChatGPT and Claude actually work, including the ideas behind neural networks, training data, attention mechanisms, and generation. You will know what these systems are genuinely good at and where they are unreliable. You will be able to use them more effectively in your work and life, verify their output, and recognize when they should not be trusted at all. And you will have enough background to follow news about AI without being confused by jargon or swept up in exaggeration.

To keep things organized, the book is divided into six parts. Part I establishes the foundation: what AI is, a brief history, and the vocabulary you need to keep different terms straight. Part II goes inside the technology and explains how large language models are built and trained, from data collection through the architecture that makes them tick, all the way to what happens when you send a prompt and get a reply. Part III examines what these systems can do well and where they fail, including the problem of hallucinations and the gap between sounding intelligent and actually being reliable. Part IV is a practical guide to using AI effectively, covering everything from basic prompting to complex workflows and when to be cautious. Part V steps back and looks at the broader implications: work, misinformation, power, ethics, law, and

the environment. Part VI looks forward at where AI development is heading, separating real progress from speculation, and wraps up with a framework for thinking about AI that you can carry forward.

A few notes before we begin. First, this book is written with the understanding that AI is advancing quickly. Specific models and capabilities will change. Some benchmarks will be surpassed. New architectures may emerge. The point is not to give you a dated snapshot of who is winning today, but to equip you with mental models that stay useful even as the field moves. The fundamental ideas about how these systems work, what they can and cannot do, and how to think critically about them will not become obsolete because a new model was released last month.

Second, the goal is not to turn you into an AI researcher or a prompt engineer. It is to turn you into an informed user who understands what is happening beneath the interface. You should be able to have conversations about AI with colleagues, friends, or policymakers without feeling lost or manipulated. You should be able to use these tools intelligently and recognize the limits of their competence.

Third, I will use analogies throughout to explain difficult concepts. Every analogy is imperfect. An analogy that helps you grasp a basic idea may mislead you if you take it too literally. I will signal when I am using an analogy and will sometimes revisit it later to refine your understanding.

Let me start with a simple observation that runs through the entire book. Modern AI systems are not like humans. They are not conscious, they do not have beliefs or desires, and they do not understand language the way we do. At the same time, they are not just keyword search engines or simple pattern matchers either. They are genuinely new kinds of systems with genuinely surprising capabilities. The key to dealing with them well is to understand what they are, without either mystifying them or dismissing them.

We will build that understanding step by step, starting from the beginning.

Chapter 1: What Is Artificial Intelligence, Really?

The Promise and the Confusion

You have heard the term artificial intelligence used in a dozen different ways. Sometimes it means a spam filter in your email. Sometimes it means a car that can drive itself. Sometimes it means a chatbot that writes poems. Sometimes it means a robot in a factory. Sometimes it means something that will think like a human. Sometimes it means something more powerful than a human. All of these are called AI, which makes the term almost useless if we want to talk about anything specific.

The confusion is not an accident. Part of it comes from the history of the field, which I will cover in the next chapter. Part of it comes from marketing, which has an incentive to make everything sound cutting edge. Part of it comes from the movies and science fiction, which have spent seventy years training us to picture AI as either a helpful companion or a threatening superintelligence. And part of it comes from the fact that artificial intelligence, as a field, has always been ambitious. The goal has been to build machines that can do things that, when done by people, we would call intelligent. That is a moving target.

Let me try to give you a clear starting point. When I use the term artificial intelligence in this book, I mean any computer system that can perform tasks that would normally require human intelligence if a person were doing them. These tasks include recognizing images, understanding spoken language, translating between languages, writing text, playing games, making predictions, and solving problems. That definition is intentionally broad because the field is broad. Not every AI system is built the same way. Not every AI system is particularly impressive.

The crucial thing to grasp is that the definition describes what the system does, not what it is. An AI system is not defined by having a mind or being conscious or understanding the world. It is defined by its behavior: the kinds

of tasks it can perform and how well it performs them. This distinction matters more than you might think, as we will see throughout the book.

Why Every Definition of AI Is a Moving Target

There is a well-known joke in the field: artificial intelligence is whatever we have not built yet. When a computer does something impressive that we did not think it could do, we are startled and call it AI. Once it becomes routine, we stop calling it AI and just call it software. When spam filters first appeared, they were hailed as AI. Now we do not think twice about them. When chess-playing programs first became better than humans, that was a landmark in AI. Now it is just a feature on your phone.

This pattern has repeated throughout the history of the field. Each breakthrough eventually becomes mundane. Each mundane technology was once considered impossible. That is partly why the definition of AI keeps shifting. As systems get better at tasks we used to think required intelligence, our threshold for what counts as intelligence moves upward.

This has practical consequences. When someone says that a new system is intelligent, you should ask: intelligent at what? And compared to what? A system can be extraordinarily good at one thing and incompetent at everything else. A system that can diagnose medical images with high accuracy is not the same as a system that can write legal documents, and neither of them is the same as a system that can drive a car through city traffic. They are all called AI, and they all use related techniques. But they are not the same thing.

The Real Core: Pattern Recognition at Scale

If there is one unifying idea behind modern artificial intelligence, it is pattern recognition. All the major AI systems you interact with today, from language models to image generators to recommendation systems, are built on the same basic principle: look at a lot of data, find patterns in that data, and use those patterns to make predictions or generate new outputs.

Consider a spam filter. It looks at thousands of emails that have been labeled as spam and not spam. It learns that certain words and certain patterns tend to appear in spam. When a new email arrives, it uses those learned patterns to make a guess about whether it is spam. That is pattern recognition.

Consider an image recognition system. It looks at millions of photos labeled with what they contain. It learns which visual patterns correspond to cats and which correspond to dogs. When shown a new photo, it uses those patterns to classify what it sees. That is pattern recognition.

Consider a language model like ChatGPT. It reads trillions of words of text. It learns which words and phrases tend to follow other words and phrases. It learns patterns in grammar and style and topic. When you give it a prompt, it uses those patterns to generate a response that fits. That is pattern recognition, at a scale that is hard to imagine.

Pattern recognition itself is not new. Humans do it all the time. What is new is the ability to do it at enormous scale using neural networks, which are mathematical structures inspired very loosely by how biological brains organize information. These networks can detect patterns that are too complex for a human to write down as explicit rules. They can handle enormous amounts of data. And they can generalize, meaning they can apply patterns they have learned to new situations they have never seen before.

That is the core of modern AI: extremely powerful pattern recognition, built on data and math, operating at scales that would have been unthinkable a generation ago.

What AI Is Not: Consciousness, Understanding, Intent

Here is where I want to be very direct, because the confusion around these issues is the source of many mistakes. Modern AI systems are not conscious. They do not experience the world. They do not have inner lives. They do not feel pleasure or pain, curiosity or boredom, pride or shame. They do not have beliefs or desires or intentions. When an AI says I want to help you, it is not expressing a want. It is generating text that, based on its training, is a likely response to your request.

This does not mean AI is harmless. It does not mean AI is unimportant. It does not mean you can ignore the consequences of how it is designed and used. But it does mean that many of the questions we ask about AI are fundamentally different from the questions we ask about people.

An AI system does not understand language the way you do. When you read a sentence, you bring to it years of experience, emotional associations, cultural context, and a sense of what things mean in the real world. An AI brings

patterns of statistical association learned from text. It knows which words tend to appear together. It knows that sentences about dogs often mention walking and fetching and barking. It does not know what a dog is, in the way that you know what a dog is. It has never seen one, petted one, or been scared by one. It has never felt the joy or inconvenience of living with a dog.

That distinction matters because it is the source of a lot of both overconfidence and underconfidence in AI. People who do not understand this think that if an AI sounds confident and articulate, it must really know what it is talking about. That is wrong. An AI can produce impressive-sounding text about almost any topic without having any actual understanding of that topic. On the other hand, people who understand this might conclude that AI is therefore useless, which is also wrong. The ability to generate useful text, even without understanding, is still incredibly valuable. You do not need a word processor to understand English in order for it to help you write. You do not need a calculator to understand numbers in order for it to help you compute. You do not need an AI to understand medicine in order for it to help a doctor organize information or explain concepts to a patient.

The point is to keep the distinction clear. AI is a tool that is very good at some things and unreliable at others. It is not a person, not a mind, not a thinking being. Treating it as if it were will lead you into trouble. Treating it as if it were nothing more than a dumb database will cause you to miss its usefulness.

Why the Confusion Persists

Why is it so hard to keep this straight? Several forces pull in different directions.

First, AI is designed to be persuasive. Language models are trained to be helpful and engaging. They respond in a conversational style. They apologize when they make mistakes. They say I understand and I think and I recommend. All of this is designed to make interaction smoother. It also makes it easy to forget that you are talking to a pattern-matching machine.

Second, the capabilities are genuinely impressive. When you ask a language model to summarize a long article, it often does a good job. When you ask it to rewrite a paragraph more clearly, it often succeeds. When you ask it to explain a concept in simple terms, it frequently produces something useful.

These are not trivial accomplishments. They are the result of real technical breakthroughs. The risk is that we mistake technical accomplishment for something deeper.

Third, many explanations of AI are themselves misleading. You will hear people say that AI thinks like a human or learns like a child or understands context. These are metaphors that can be helpful for grasping general ideas but are also easily taken too literally. A child learns language by interacting with the world, connecting words to objects and events and other people. An AI learns language by processing text, connecting words to other words. The similarity is at the level of abstract mathematics. The difference is everything else.

Fourth, researchers and companies sometimes talk about AI in ways that invite confusion. Researchers use terms like intelligence and learning and reasoning as technical vocabulary that may not carry their everyday meanings. Companies have an incentive to make their products sound more capable than they are. The media often does not have the background to correct either side. The result is a public conversation in which technical accuracy is thin on the ground.

What This Book Will Help You Understand

This book is built to resist those confusions. It will not ask you to suspend your skepticism or accept AI as magic. It will not try to scare you into thinking AI is an imminent catastrophe. It will try to give you a clear, precise picture of what modern AI is, how it works, what it can and cannot do, and what is happening with it in the real world.

Here is what I want you to walk away with:

- A vocabulary for talking about AI that is precise enough to avoid confusion but simple enough to be practical.
- A mental model of how large language models work that helps you understand both their capabilities and their limitations.
- A realistic sense of what AI can help you with, what it cannot, and when its help is dangerous without careful verification.
- Practical skills for using AI tools effectively in your work and life.
- Enough background to evaluate claims about AI, whether they come from journalists, executives, researchers, or friends.

- A framework for thinking about the broader social, economic, and ethical questions that AI raises.

We start with history, because understanding how we got to where we are helps explain why the technology is the way it is and why certain problems persist. The next chapter covers the journey from early dreams of thinking machines to the data-driven systems of today.

Chapter 2: From Rule-Based Programs to Learning Machines

Early Dreams: The Birth of the Field in the 1950s

The formal story of artificial intelligence begins in the mid-twentieth century, in the immediate aftermath of World War II, when digital computers were transitioning from theoretical ideas and wartime code-breakers to general-purpose machines. The people who founded the field were ambitious, optimistic, and sometimes naive about how hard the problem really was.

In 1950, Alan Turing, one of the founding figures of computer science, published a paper titled *Computing Machinery and Intelligence*. Turing did not ask whether machines could think, which he considered too vague. Instead, he proposed a practical test. If a human evaluator, communicating with both a machine and a person through text, could not reliably tell which was which, then the machine had demonstrated a form of intelligence. This is now called the Turing test. It is not a perfect measure of intelligence, and most researchers no longer treat it as the goal. But it captured something important: the idea that intelligence could be defined in terms of observable behavior rather than hidden inner states [1].

In 1956, a summer workshop at Dartmouth College brought together a small group of researchers, including John McCarthy, Marvin Minsky, Claude Shannon, and Nathaniel Rochester. McCarthy coined the term artificial intelligence during the planning of this conference. The proposal for the workshop stated that the event would attempt to find how aspects of learning and any other feature of intelligence might be so precisely described that a machine can be made to simulate them. The participants believed they could make a significant advance on machine intelligence within two years. They were not wrong about making advances. They were very wrong about the timeline [2].

The early results were impressive. A program called the Logic Theorist, developed by Allen Newell and Herbert Simon, proved mathematical theorems. Another program, the General Problem Solver, attempted to solve abstract

puzzles using search and heuristic strategies. A program called ELIZA, created by Joseph Weizenbaum in 1966, simulated a Rogerian psychotherapist by reflecting users statements back in the form of questions. People who interacted with ELIZA, including Weizenbaum himself, found it surprisingly convincing, even though the program was doing something very simple.

The mood of the time was one of excitement. Many researchers believed that general intelligence, the kind of flexible reasoning that humans display across many different domains, was within reach. They thought the main challenge was finding the right representation for knowledge and the right algorithms for reasoning about it. The idea that the problem might be fundamentally harder than that, or that a radically different approach might be necessary, was not on most people's radar.

Rule-Based Systems and Expert Systems

Through the 1960s and 1970s, a dominant approach in AI was rule-based programming. The idea was straightforward: intelligent behavior comes from having a lot of rules and being able to apply them in the right order. If a computer has all the rules that an expert uses, and if it can apply those rules correctly, then the computer will be able to perform like an expert.

This approach led to the development of expert systems, which were among the first genuinely useful AI technologies. An expert system encodes the knowledge of human experts in a specific domain as a set of rules and then uses those rules to reason about new cases. One of the first successful expert systems was Dendral, developed in 1965 at Stanford University. Dendral could propose molecular structures based on mass spectrometry data, performing at a level comparable to human chemists. Another famous example was MYCIN, developed in the 1970s, which could recommend antibiotic treatments for bacterial infections [3].

Expert systems had real advantages. They could encode specialized knowledge that would otherwise be lost when an expert retired. They could apply that knowledge consistently, without fatigue or distraction. In narrow domains with well-understood rules, they worked well. They were also the first AI technology to achieve commercial success, with companies building expert system shells and knowledge management tools in the 1980s.

But expert systems had fundamental problems that eventually limited them. First, encoding knowledge as explicit rules is extremely difficult. Human

experts often cannot articulate the rules they use. Their knowledge is partly intuitive, based on patterns they have absorbed through experience. When a grandmaster chess player looks at a board and sees a good move, they are not running through a checklist of rules. They are recognizing a pattern. Expert systems required knowledge engineers to interview experts and translate their knowledge into formal rules, a process that was slow, expensive, and often incomplete.

Second, rule-based systems are brittle. They work well inside the domain for which they were designed but fail spectacularly outside it. An expert system for diagnosing bacterial infections cannot answer a question about car repair, even a simple one. It has no mechanism for drawing on related knowledge or using common sense. It does not know what a car is.

Third, rule-based systems do not learn. If the world changes, or if the rules turn out to be incomplete or wrong, the system does not adapt. Someone has to manually update the rules. This is fine for a stable domain where the rules do not change much. It is not fine for most real-world problems.

The Limits of Hand-Crafted Rules

The problems with rule-based systems pointed to a deeper issue: many of the tasks we think of as requiring intelligence are not primarily about following explicit rules. They are about recognizing patterns in messy, complex, high-dimensional data. Recognizing a face. Understanding a sentence. Navigating a room. These tasks are easy for humans because our brains are incredibly good at pattern recognition. They are hard for rule-based systems because the patterns are too subtle and too numerous to write down as rules.

This became clear in areas like machine translation and speech recognition. In the 1960s, researchers had great optimism about machine translation, partly because of the strategic importance of translating between Russian and English during the Cold War. Early systems, like the Georgetown-IBM experiment in 1954, translated a small set of Russian sentences into English with enough accuracy to impress audiences. But when researchers tried to build systems that could handle real-world text, the problems became apparent. Natural language is full of ambiguity. Words have multiple meanings. Sentences have complex structures. Context matters in subtle ways. Rule-based approaches, which tried to encode grammatical rules and dictionary definitions, struggled to handle the messiness of real language.

Similar problems arose in speech recognition. The idea of encoding rules for how sounds combine into words and sentences seemed promising in theory. In practice, human speech varies enormously across speakers, accents, contexts, and environments. A rule-based system could not handle the diversity.

The same issue appeared in computer vision. The task of recognizing objects in images seemed like it should be reducible to rules about edges, shapes, and colors. But real images are complex. Objects appear at different scales and angles, in different lighting conditions, partially occluded, surrounded by clutter. No one could write down a complete set of rules that would work.

By the mid-1970s, the limitations of the rule-based approach were becoming hard to ignore. Progress had slowed. The ambitious promises of the early years had not materialized. Researchers had built interesting programs that could solve narrow problems, but nothing close to general intelligence, and nothing that could handle the messy complexity of the real world in a robust way.

The First AI Winters

When progress stalls and funding dries up, the field calls it a winter. The first AI winter ran roughly from the mid-1970s to the early 1980s. It was triggered by a combination of technical disappointment and political pressure.

In 1973, a report by British computer scientist James Lighthill criticized AI research for failing to deliver on its promises. The Lighthill Report argued that AI had concentrated on trivial problems and had failed to tackle genuinely hard tasks. It recommended cutting government funding for AI in Britain. The United States government followed a similar path, reducing support for AI research through agencies like DARPA [4].

The report was not entirely fair. AI researchers had made real progress, even if it was not the progress they had promised. But the disappointment was real, and the funding cuts were real. Many research labs closed. Many researchers left the field.

There was also a technical factor. In 1969, Marvin Minsky and Seymour Papert published a book called *Perceptrons*, which analyzed a simple type of neural network called the perceptron. They showed that a single-layer perceptron could not solve certain basic problems, like distinguishing whether two sets of points were connected or separate. Although they acknowledged that multi-layer networks might be able to solve these problems, the book

had a chilling effect on neural network research. For years, the computational community largely abandoned neural networks, focusing instead on symbolic approaches like rule-based systems and logic programming.

The first winter was not absolute. Research continued, often in obscurity, with smaller grants and less public attention. But the field shrank. The optimism evaporated.

The Expert Systems Boom and the Second Winter

The 1980s brought a revival, this time driven by commercial interest in expert systems. Companies realized that expert systems could solve real problems, at least in narrow domains. Japan announced an ambitious Fifth Generation Computer Systems project in 1982, which aimed to build computers capable of intelligent reasoning. This provoked a response from the United States and Europe, which increased funding for AI research to compete.

For a while, it worked. The expert systems market grew rapidly. Companies sold AI consulting services and knowledge management tools. The stock market celebrated AI stocks. The winter seemed to be over.

But the problems with expert systems, which I described earlier, eventually caught up with the boom. The systems were expensive to build and maintain. They were hard to update when knowledge changed. They did not generalize well. They could not handle uncertainty or noisy data. And the hardware and software needed to run them were expensive. When companies realized that the return on investment was not what they had been promised, the spending dried up.

In the late 1980s and early 1990s, the field entered a second winter. Expert system companies went bankrupt. Research funding declined. AI became associated with overpromising and underdelivering. The term itself was sometimes considered toxic, and researchers in the field began to use other labels like machine learning or data mining to avoid the baggage [5].

Again, the winter was not absolute. Research continued. But the mood was different from the 1950s and 1960s. Researchers were more cautious. They made smaller claims. They focused on narrower problems.

The Data Revolution: More Compute, More Data, More Patience

The modern era of AI began with a change in approach, made possible by three factors: more computing power, more data, and a shift in methodology.

The computing power factor is straightforward. As processors got faster and cheaper, it became possible to do computations that had previously been impossible. In particular, it became possible to train large neural networks, which require enormous amounts of computation. The availability of graphics processing units, or GPUs, originally designed for rendering video games, provided a particular boost. GPUs are highly parallel, meaning they can do many calculations at once. This makes them well-suited to the kinds of mathematical operations used in neural networks. Researchers discovered that they could train models that were orders of magnitude larger than before [6].

The data factor is equally important. The growth of the internet generated vast amounts of digitized text, images, audio, and video. Websites like Wikipedia provided structured knowledge. Social media, forums, and blogs provided informal language in enormous quantities. Companies collected data on user behavior, purchases, and interactions. By the 2010s, it was possible to collect datasets with billions of examples, a scale that would have been unimaginable in the 1980s.

The methodological shift was the most important of all. Instead of trying to handcraft rules for intelligent behavior, researchers focused on learning from data. The basic idea, known as machine learning, is to give a computer a lot of examples and let it find patterns on its own. Instead of telling the computer how to recognize a cat, you show it millions of pictures of cats and non-cats and let it learn the difference. Instead of writing rules for translating between languages, you give it millions of parallel sentences and let it learn the translation patterns.

This approach had been around since the 1950s, but it had not been taken very seriously because the computational tools and data were not available. Neural networks, the primary tool for modern machine learning, were first proposed in the 1940s and 1950s. But without enough computing power and without enough data, they could not learn useful patterns. With both available, everything changed.

The Deep Learning Breakthrough

The moment that most researchers point to as the beginning of the modern AI era is 2012. That year, a team led by Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton entered a competition called the ImageNet Large Scale Visual Recognition Challenge. The task was to classify images into one of a thousand categories. The best existing systems had error rates around 26%. Krizhevsky and his colleagues trained a deep neural network, later called AlexNet, using GPUs and the ImageNet dataset, which contained 1.3 million labeled images. Their system achieved a top-5 error rate of 15.3%, more than ten percentage points better than the runner-up. The margin of victory was so large that it stunned the community [7].

AlexNet demonstrated that deep neural networks, when trained on large datasets with sufficient computing power, could learn incredibly complex visual patterns and outperform systems built using hand-crafted features and rules. It was a proof of concept that changed the direction of research overnight.

The success of AlexNet triggered a wave of investment and experimentation. Researchers applied deep learning to speech recognition, where it dramatically improved accuracy. They applied it to machine translation, where it produced translations that were fluent and often indistinguishable from human work. They applied it to game playing, where systems like AlphaGo defeated world champions at the ancient game of Go, a task that many experts had thought would remain beyond machines for decades.

Deep learning became the dominant paradigm in AI. Almost every major advance in artificial intelligence since 2012 has been based on some variant of deep neural networks. The approach is now behind the AI systems that most people interact with: voice assistants, image recognition, recommendation systems, and large language models like ChatGPT and Claude.

What made deep learning work was not a single breakthrough but the combination of three ingredients that aligned at the right time. Large datasets provided the material for learning. Fast hardware, especially GPUs, provided the computational power to process that material. And improved algorithms, including better training methods and architectural innovations, made it possible to train networks that were deeper and more powerful than before.

The next chapter will explain the vocabulary of modern AI, laying out

the relationships between terms like artificial intelligence, machine learning, deep learning, neural networks, generative AI, large language models, and others. Understanding those relationships is essential for making sense of what follows.

Chapter 3: The Family Tree of AI Terms

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

The Nesting Doll Model: How These Terms Relate

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

Artificial Intelligence as the Umbrella

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

Machine Learning and Its Variants

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

Deep Learning and Neural Networks

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

Generative AI and Large Language Models

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

Specialized Domains: NLP, Vision, Robotics, Agents

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

Putting It All Together

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

Chapter 4: The Training Ground: How LLMs Learn from Data

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

Where the Data Comes From: The Internet and Beyond

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

Cleaning and Filtering Training Data

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

What a Model “Sees” During Training

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

Learning as Pattern Detection, Not Reading

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

The Scale of Modern Training Datasets

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

Data Quality, Biases, and What Data Shapes the Model

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

Chapter 5: Words Become Numbers: Tokens and Embeddings

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

How Computers Understand Text: The Necessity of Numbers

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

Tokenization: Splitting Words into Pieces

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

Why Tokens Matter for Context Limits and Cost

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

Embeddings: Turning Tokens into Vectors of Meaning

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

Similarity in Vector Space: Why Related Ideas Cluster Together

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

What Embeddings Capture and What They Miss

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

Chapter 6: The Neural Network: Brains Made of Math

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

The Neuron Metaphor and Its Limits

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

Layers and Computation: How Data Flows Through

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

Parameters: The Knobs the Model Adjusts

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

Learning by Minimizing Errors

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

Why Deep Networks Are Powerful

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

The Meaning of Model Size and Scaling

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

Chapter 7: Attention and Transformers: The Architecture That Changed Everything

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

The Problem with Earlier Models

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

The Transformer Revolution

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

Attention Explained: Weighing What Matters

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

Self-Attention and Context

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

Why Transformers Handle Language So Well

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

The Recipe: Tokens Plus Embeddings Plus Attention Plus Layers

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

Chapter 8: From Pretraining to Polishing: Fine-Tuning and Alignment

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

Pretraining: Learning Language from Raw Text

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

What the Model Learns During Pretraining

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

Supervised Fine-Tuning: Teaching Helpful Behavior

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

Reinforcement Learning and Human Feedback

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

Reward Models and Preference Optimization

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

Why Raw Pretrained Models Are Not Chat-Ready

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

Chapter 9: When You Hit Send: How Inference and Generation Work

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

The Difference Between Training and Inference

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

Next-Token Prediction: The Basic Task

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

Autoregressive Generation: Building Responses Word by Word

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

Temperature and Sampling: Controlling Creativity

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

Why Generation Is Not Database Lookup

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

Speed, Cost, and the Context Window

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

Chapter 10: The Real Strengths of Modern AI

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

Language Understanding and Generation

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

Summarization and Information Synthesis

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

Translation and Cross-Lingual Capability

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

Code Generation and Technical Assistance

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

Creative Support and Idea Exploration

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

Speed, Scalability, and Accessibility

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

Putting It Together

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

Chapter 11: The Real Limitations of Modern AI

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

Hallucinations: What They Are and Why They Happen

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

Confidence Without Correctness

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

Arithmetic and Systematic Reasoning Failures

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

Biases and Incomplete Perspectives

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

The Knowledge Cutoff and Outdated Information

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

Prompt Injection and Security Concerns

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

Putting It Together

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

Chapter 12: What LLMs Do Not Know, Do Not Feel, and Do Not Understand

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

The Fluency Trap: Why Sounding Smart Is Not Being Smart

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

Pattern Matching Versus Understanding

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

The Chinese Room and Similar Arguments

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

Knowledge, Beliefs, and Truth

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

Consciousness, Feeling, and What We Know and Do Not Know

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

What We Can and Cannot Say About Model “Intelligence”

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

Chapter 13: How to Talk to AI: Prompting That Actually Works

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

What a Prompt Is and How the Model Interprets It

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

Clarity: The Most Important Rule

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

Providing Context and Constraints

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

Giving Examples: Few-Shot Prompting

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

Asking for Structured Outputs

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

Common Prompting Mistakes and How to Fix Them

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

Chapter 14: Working with AI on Complex Tasks

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

Decomposition: Breaking Big Tasks into Steps

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

Iterative Refinement: Prompting, Reviewing, Revising

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

Self-Critique: When It Helps and When It Does Not

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

Chaining and Multi-Step Workflows

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

Combining Multiple Tools and Techniques

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

When to Stop Delegating to AI

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

Chapter 15: Tools, Retrieval, and AI Agents

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

The Limits of Training-Time Knowledge

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

Tool Use: Giving AI Functions to Call

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

Retrieval-Augmented Generation (RAG)

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

AI Agents and Autonomous Action

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

The Trade-Offs: Flexibility Versus Reliability

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

Chapter 16: Building AI Into Your Workflow

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

Email, Communication, and Administrative Tasks

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

Writing and Content Creation

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

Research and Information Gathering

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

Programming and Technical Work

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

Learning and Education

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

When NOT to Use AI (or When to Be Extremely Cautious)

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

Chapter 17: How to Evaluate AI-Generated Information

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

Fluency, Plausibility, and Accuracy: Not the Same

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

Spotting Hallucinated Sources and Citations

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

Cross-Checking Key Facts

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

Using AI to Help Verify Itself

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

When AI Output Needs Expert Review

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

Building Verification Habits

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

Chapter 18: AI and Work: Transformation, Not Replacement

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

The Reality of Automation: Tasks, Not Jobs

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

What AI Can Realistically Replace

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

Augmentation: Working with AI Rather Than Against It

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

Who Benefits and Who Does Not

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

Skills That Matter More and Skills That Matter Less

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

Career Strategy in an AI-Shaped Economy

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

Chapter 19: Misinformation, Deepfakes, and the Future of Truth

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

The Scale Problem: AI-Generated Content at Industrial Volume

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

Deepfakes and Synthetic Media

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

Misinformation and Manipulation

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

Authentication and Provenance

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

The Trust Crisis

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

What Individuals and Society Can Do

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

Chapter 20: Power, Ethics, and the Questions We Have Not Solved

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

Who Controls AI and Why It Matters

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

Environmental Costs of Training and Running Models

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

Copyright, Data Rights, and Intellectual Property

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

Surveillance, Privacy, and Civil Liberties

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

Fairness, Bias, and Harm

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

Regulation and Governance

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

Chapter 21: The Next Steps: What Is Emerging Now

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

Scaling and Specialization Trends

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

Multimodal Systems

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

Reasoning Models and Chain-of-Thought

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

On-Device and Local AI

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

Robotics and Embodied AI

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

Reliability and Verifiability Research

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

Chapter 22: Living with AI: A Practical Conclusion

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

Your Mental Model of AI: The Key Takeaways

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

Using AI Wisely: Principles Over Tricks

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

Knowing When to Trust, When to Verify, When to Ignore

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

Continuing Your Education About AI

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

Being a Responsible User and Citizen

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

The Human Center: AI as a Tool, You as the Purpose

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

References

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