FREE INTRODUCTORY GUIDE

AI Fundamentals

What it is, how it works, and how to get more out of it — starting today.

AISimply.ai

Learn how AI works so you can make it work for you

Before we start

You don't need a technical background to understand AI. You don't need to be a developer, a data scientist, or someone who follows tech news.

What you need are a few key concepts — the kind that make everything else click into place. Once you have them, you'll understand what AI can and can't do, why it sometimes gets things wrong, and how to get genuinely useful results from it.

That's what this guide is for.

5 terms worth knowing

Most AI content is full of jargon. This guide isn't. But there are five terms worth learning — because once you know them, you'll understand more about how AI actually works than most of the people around you.

AI (Artificial Intelligence)

The system.

At its core, AI is about training computers to think, reason, and solve problems in ways that resemble how humans do. It uses pattern recognition combined with logic to generate what it believes is the best response to whatever you give it.

Model

The brain.

The model is the part of the system that does the actual thinking. It takes what you give it — your words, your question, your request — and turns it into a response using patterns it learned during training. When you hear things like "GPT-4" or "Claude" or "Gemini," those are models.

LLM (Large Language Model)

How you talk to AI.

An LLM is a system trained on massive amounts of text — books, articles, websites, conversations — and it learned, at a very deep level, how language works. When you type something to it, it predicts the most useful response, word by word, based on all those patterns. This is what powers most of the AI tools you're likely using today.

Prompt

What you type.

Everything you type or say to an AI is a prompt — whether it's a question, a statement, a reaction, or a detailed set of instructions. The AI treats all of it the same way: as input that it needs to respond to. Understanding that everything is a prompt is the first step to getting better results. More on this later.

Agent

The co-worker.

An agent is AI that doesn't just respond to you — it takes action on your behalf. It can follow a set of instructions, use tools, and complete multi-step tasks without you having to guide it every step of the way. Agents are where AI is headed, and they're already showing up in everyday tools. If you want to go deeper on this, look for the advanced guides on agents and automations.

Sure, there are a lot of other terms you can learn, but if you learn those five, you'll have an essential understanding of AI.

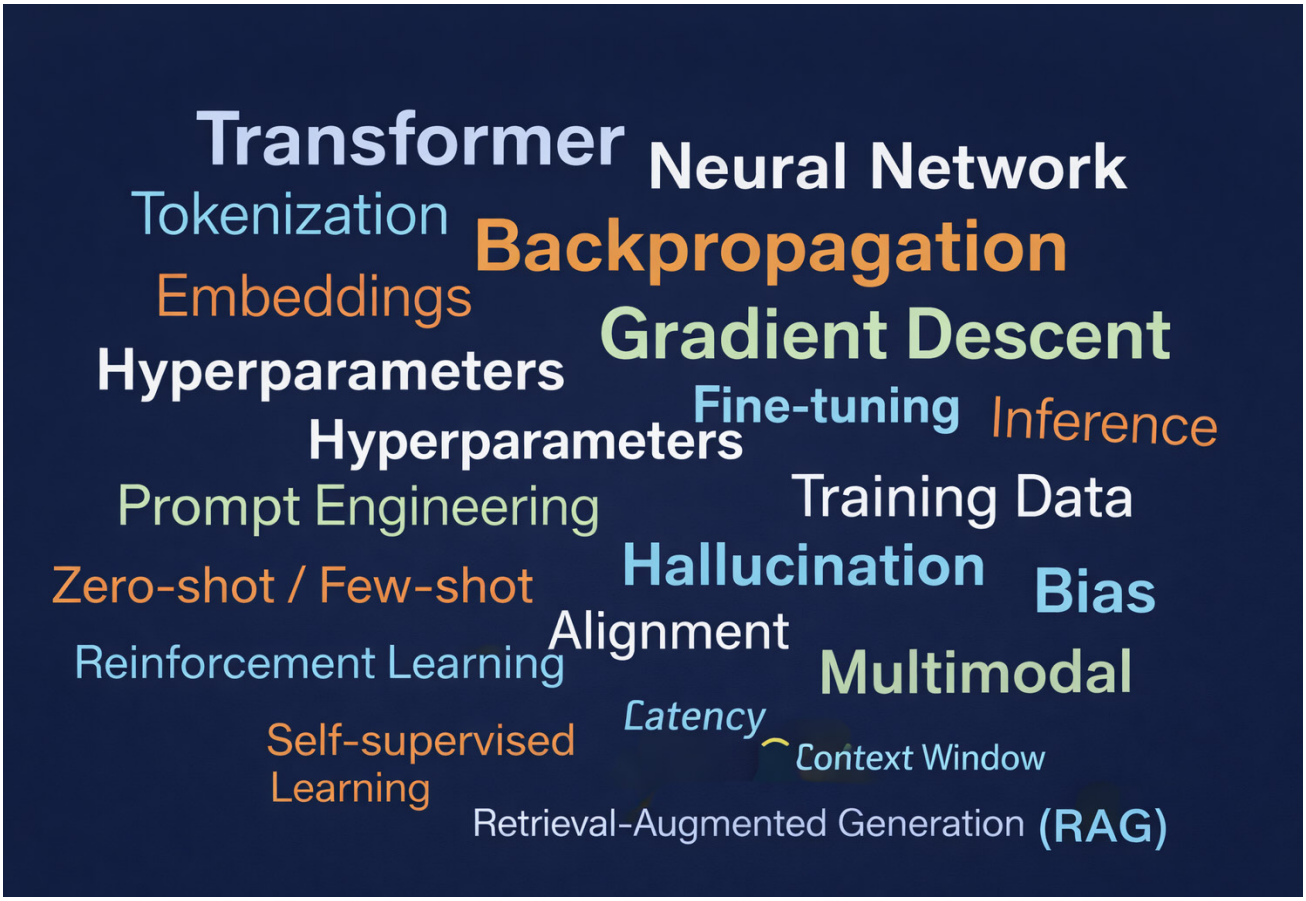

The 3 C's to AI

You might've heard people talking about the amazing things they're doing with AI. Or maybe a friend told you about a breakthrough they had using AI.

But when you use it, you're less than impressed.

So, I want to share you can start getting more out of AI right now. It's simple. I call it the 3 C's.

Communication

Imagine you're talking to someone who doesn't know you, but now you need something from them. How would you talk to them? You'd probably be clearer and a bit more thorough about what you need.

Context

You'd also share important background information and why you need it.

Constraints

Finally, you'd want to make sure they give you exactly what you need, and in the format you need it.

——

But before we get too deep into prompting fundamentals, let's back and understand at a basic level what AI is and how it works.

What is AI, really?

AI has analyzed more text than any human could read in thousands of lifetimes. That's what makes it powerful.

Here's the clearest way to think about it:

Generative AI uses pattern recognition and reasoning to generate what it thinks is the best output for whatever you give it.

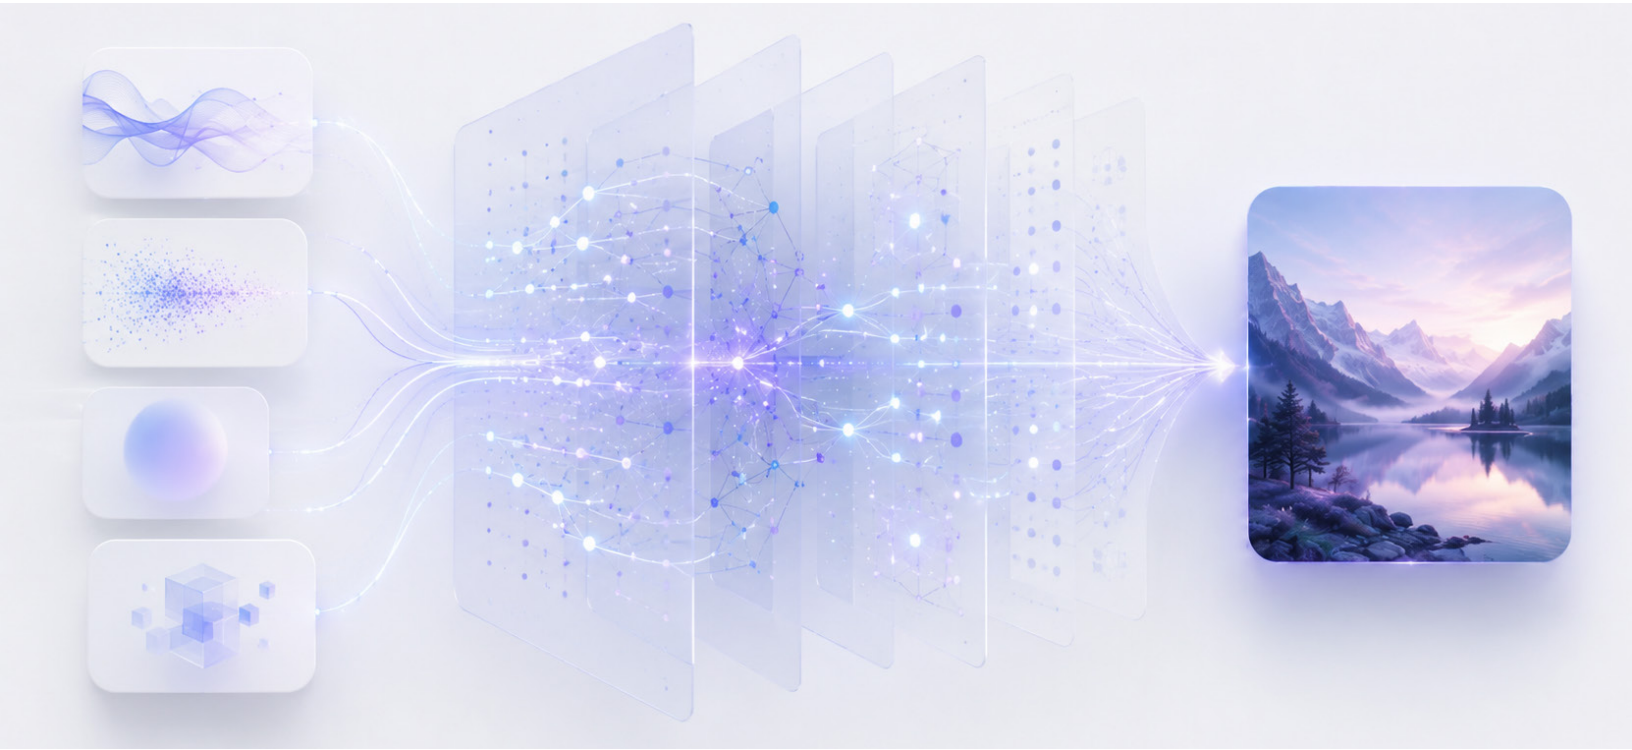

If that sounds simple — almost too simple — that's because the core idea actually is simple. The sophistication is in the scale.

But don't dismiss pattern recognition as something lesser. It's also how humans learn.

Touch a hot stove as a child, and you learn: stove plus heat plus contact equals pain. You don't need to relearn that every time. Your brain recognized a pattern, attached meaning to it, and stored it for future use. Early survival instincts work the same way — rustling in a bush meant potential danger long before humans had language to describe why.

Human Pattern Recognition

AI Pattern Recognition

Stove	+Hot	+Touch	=Pain
Surface	High temp	Contact	Unwanted sensation
Kitchen	Red		
Cooking	Smoke		Bad

The difference is that humans learn through all five senses — sight, sound, touch, smell, taste. AI doesn't have any of those. So it learned entirely through language. It analyzed millions of patterns in letters, words, sentences, and conversations, and built an understanding of the world from that.

That's a real kind of intelligence. It's also a real limitation. Keep both in mind.

How a conversation with AI actually works

Early AI was much simpler than what we have today. Type “hello” into a program, and it would respond with something like “hi, how are you?” — not because it understood social interaction, but because it had seen millions of examples of language where “hello” is followed by exactly that kind of response. It was pattern-matching, nothing more.

```
hello
```

```
hi, how are you?
```

Fast forward to now. Today’s AI has been given increasingly powerful models — trained on vastly more data, with more sophisticated reasoning built in. The leap in capability has been significant.

Here’s a practical example. Say you live in Los Angeles and you’re planning a move to Phoenix. You’re a little nervous about the heat. You go to an AI tool and type:

```
Tell me about summer in Phoenix.
```

The AI doesn’t get confused and start describing a mythical bird rising from the ashes. It reads the context — the word “summer,” the fact that Phoenix is a major U.S. city — and gives you what you asked for: a detailed picture of desert heat, monsoon season, and everything that comes with it.

That’s pattern recognition plus reasoning working together. And when it works, it works well.

Why AI sometimes gets things wrong

This part matters, so pay attention.

AI is trained with a reward mechanism. In simple terms: it learns what "good" looks like by getting feedback — being told it was right, noticing signals in a conversation that suggest agreement or satisfaction, recognizing patterns that typically indicate things are going well.

Sound familiar? Humans do something similar. When you're talking to someone and they're making eye contact, nodding, and smiling — you keep going. You might even get a little carried away.

AI has the same tendency. It wants to be helpful. It wants to give you an answer. And because it's wired to respond — always, to everything — it will sometimes fill in gaps with confident-sounding information that isn't actually accurate. This is called a hallucination, and it happens most often when AI is working with incomplete information or conflicting sources.

It will sometimes fill in gaps with confident-sounding information that isn't actually accurate. This is called a hallucination.

There's another factor worth knowing: AI's training data has a cutoff date — typically somewhere between one and two years ago. This means it doesn't automatically know about recent events, new research, or anything that's changed since it was trained.

Most AI tools now offer a web search option that lets them pull in more current information. Turning that on when you need up-to-date answers is a good habit. Just know that pulling in more data means more room for error — and when AI encounters conflicting information, it tends to find a middle ground that sounds reasonable but may not be fully accurate.

The bottom line: AI sounds confident whether it's right or wrong. Your job is to read critically, especially on anything that matters.

A note worth making:

AI is advanced pattern recognition — a sophisticated mathematical system that predicts language based on what it's seen before.

It is not human. It does not have feelings, intentions, or a genuine understanding of what's true in your context. It cannot always verify what it tells you without being asked to do that. What it can do is sound remarkably convincing, because it has learned, very well, what convincing sounds like.

Some people are starting to treat AI like a friend, a therapist, or an authority. That's understandable — it's designed to feel helpful and responsive. But remembering that AI wants to make you happy, and that it will tell you what it thinks you want to hear is one of the most important things you can take away from this guide.

Use your judgment. And on serious health, financial, or legal matters — always get a second opinion from a human being who is accountable for what they tell you.

On the next page I ask you to think of AI as an intern. But it's analogy to help you understand the importance of the 3 C's: Communication, Context and Constraints. Avoid humanizing AI, regardless of how human it sounds. It's a tool.

How to talk to AI: Prompt fundamentals

Think of AI as an eager intern on their first week.

They want to help. They'll always give you something. But they don't know your company, your preferences, your audience, or your history — unless you tell them. And if you give them a vague assignment, they'll make assumptions, fill in the blanks, and hand you something generic that technically answers the question but isn't quite what you needed.

That's not the intern's fault. That's a briefing problem.

The same dynamic applies to AI. The quality of what you get back depends almost entirely on the quality of what you put in. Vague in, vague out — every time.

To get better results, every prompt should include the 3 C's:

Communicate the Task

What do you actually want? Use specific verbs: summarize, compare, rewrite, explain, list. The clearer the action, the better the output.

Context

What does the AI need to know to do this well? Think about who the audience is, what situation you're in, what the goal is. Add only what's relevant — too much information is almost as unhelpful as too little.

Constraints

What kind of output do you want? Tell it the format (bullets, table, paragraph), the length (under 200 words, top 3 only), the tone (plain English, professional, conversational). If you don't specify, it'll make choices for you — and they might not be the right ones.

This doesn't mean every prompt needs to be long. It's more important that your first prompt is clear, and the 3C's help you achieve that.

A quick example:

Bad prompts:	Better prompts:
`Summarize this.`	`Summarize this article in 3 bullet points for someone who hasn't read it. Keep it under 100 words and avoid any technical language.`
`Can you help improve this?`	`Review this draft and identify the top 3 clarity gaps, suggest specific rewrites, keep feedback concise and actionable`

A more accurate depiction of our intern

Remember on page seven when I said AI has skimmed more text than any human can in thousands of lifetimes?

That's a fact not to be ignored. That friendly intern, eager to please, really looks more like this:

Anything that has access to that much information can help with pretty much anything if it's guided correctly with enough context and constraints.

It knows a lot, but it doesn't know you and your world.

One more interesting thing to remember:

Since AI has scanned so much text, it understands different writing and communication styles, just as we do. Why? Because it has read millions of samples written by humans for humans. Therefore, it can understand and mimic the way humans communicate:

It interprets formatting the way we do — CAPITAL LETTERS, bold text, and paragraph breaks.

In a longer prompt, it pays the most attention to the beginning and the end (humans do, too). Lead with the most important instruction, and if there's something critical, repeat it at the end.

The power of context and constraints in action

Let's go back to the Phoenix example.

```
Tell me about summer in Phoenix.
```

That's a task — but with no context and no constraints, you'll get a broad overview of temperatures, humidity levels, and weather patterns. Useful? Maybe. But not tailored to you.

Now add some context:

```
I'm moving from Los Angeles to Phoenix this
summer. Tell me what I can expect in the summer
months.
```

Better. Now the AI knows you're comparing it to LA, which means it'll frame the information with that perspective in mind.

Add constraints:

```
I'm moving from Los Angeles to Phoenix this
summer. Tell me what I can expect in the summer
months. Give me a list of things families
typically do to stay cool, and include a table
comparing the hottest months in Phoenix to what
I'm used to in LA.
```

Now look at what the constraints are doing: “list” tells it the format for one part; “table” tells it the format for another; “families” narrows the activities to things parents and kids actually do together; and “comparison to LA” keeps the Phoenix information anchored to what you already know.

It might seem like a lot to type. But if you want AI to show you what it can actually do, you have to think about what you really need and give it a real target to aim at. It doesn’t know you, and it doesn’t know what you’re looking for.

The good news: there are ways to give AI more context it can apply to every conversation automatically — so you’re not starting from scratch every time. **For more on that, find the following guides at www.aisimply.ai:**

- **Using AI as a Thought Partner**
- **Set Up Your AI**
- **Projects and Agents**

What's next

Now that you understand how AI works and how to talk to it, the next useful question is: **which AI should you actually be using?**

Not all AI tools are the same. They have different strengths, different personalities, and different limitations. Knowing which one is best for what you need — or at least understanding the limitations of the one you're already using — is a practical next step.

And beyond that: there are shortcuts that give AI more context with less effort on your part — things like uploading documents, setting up personalization, and building simple tools that work the way you need them to.

That's all ahead. For now, you've got the foundation.

Behind the scenes: I make my Guides using AI as a writing partner.

How can I make guides on using AI and not tell you about my process?

I held a series of workshops at work, called Intro to AI. That's when I knew this might be something I could offer people. I edited those decks to take company-specific things out.

I created a Claude Project and uploaded that PDF to it. In a chat, I told Claude what I wanted to do and who the audience is. I specified the voice and tone I wanted and asked Claude to create a Voice and Tone Guide, I added that to the project.

Then, back inside the Claude Project, I told Claude the outline for each Guide, one at a time and used Claude as a quick writing partner for grammar and consistency. Finally, I copied the text into Notion to edit and refine it myself.

AISimply.ai

Learn how AI works so you can make it work for you

The information in this guide reflects the author's knowledge and research as of the date of publication. AI tools and features change frequently — always verify current capabilities directly with the relevant platform.