

Fundamentals & Prompt Engineering

Learn How LLMs Work and Master the Art of Prompting

A foundational guide covering how ChatGPT and Claude work, the core prompting techniques, practical text-generation use cases, and the essential AI tool stack.

A Study Guide for AI Learners

Table of Contents

- 1. Introduction to LLMs — How ChatGPT/Claude Work
- 2. The Art of Prompting: Zero-shot, Few-shot, Chain of Thought
- 3. Text Generation: Marketing Copy, Blogs, Emails, Code
- 4. Tool Stack: ChatGPT Plus, Claude 3, Google Gemini
- 5. Practice Exercises
- 6. Quick-Reference Cheat Sheet

1. Introduction to LLMs

A Large Language Model (LLM) is an AI system trained on massive amounts of text data to understand and generate human-like language. Tools like ChatGPT and Claude are built on top of LLMs (GPT-4/GPT-4o for ChatGPT, and the Claude model family for Anthropic).

How do they actually work?

- **Training:** The model reads billions of words from books, websites, and articles to learn patterns in language.
- **Tokens:** Text is broken into small chunks called tokens; the model predicts the next token, one at a time.
- **Transformer Architecture:** A neural network design that lets the model pay 'attention' to relevant parts of the input, no matter how far apart they are in the text.
- **Fine-tuning & RLHF:** After initial training, models are refined using human feedback to make responses more helpful, accurate, and safe.

Important distinction:

LLMs don't 'think' or 'know' facts the way humans do — they predict the most statistically likely next word based on patterns. This is why they can sound confident even when incorrect (a phenomenon called **hallucination**). Always verify critical facts independently.

2. The Art of Prompting

Once you understand that LLMs are pattern-predictors, prompting becomes about giving the model the clearest possible pattern to follow. Below are the three foundational techniques.

Zero-shot Prompting

You ask the model to complete a task directly, with no examples. Works well for tasks the model has seen many times during training (translation, summarization, simple Q&A).

```
"Summarize this article in 3 bullet points: [paste article]"
```

Few-shot Prompting

You provide 2-5 examples of the input-output pattern before your actual request. This helps the model match a specific style, format, or logic that zero-shot alone might miss.

```
"Review: 'Great product, fast delivery' -> Sentiment: Positive\nReview: 'Broke after one day' -> Sentiment: Negative\nReview: 'It's okay, does the job' -> Sentiment: ?"
```

Chain of Thought (CoT) Prompting

You explicitly ask the model to reason step-by-step before answering. This significantly boosts accuracy on math, logic, and multi-step reasoning tasks by forcing the model to 'show its work'.

```
"A store had 120 apples. It sold 45% on Monday and 30 more on Tuesday. How many are left? Think step by step."
```

Combine techniques: a few-shot prompt that also demonstrates step-by-step reasoning in each example is one of the most powerful patterns for complex tasks.

3. Text Generation: Practical Use Cases

Marketing Copy

Give the AI the product, audience, tone, and platform. Ask for multiple variations to A/B test.

```
"Write 3 Instagram ad captions for a vegan protein powder targeting gym-goers aged 20-30. Tone: energetic and bold. Max 40 words each."
```

Blogs

Use prompt chaining: first generate an outline, then expand each section, then edit for tone and flow.

```
"Create a detailed outline for a 1000-word blog post titled '5 Morning Habits of Productive People'."
```

Emails

Specify the relationship with the recipient, the goal of the email, and desired length/tone.

```
"Write a polite follow-up email to a client who hasn't responded in 2 weeks about a project proposal. Keep it under 100 words, professional but warm."
```

Code

Always specify the programming language, libraries/versions, and what the code should do. Ask the model to explain its logic or add comments for learning purposes.

```
"Write a Python function using pandas that reads a CSV, removes duplicate rows, and saves the cleaned file. Add comments explaining each step."
```

4. Tool Stack: ChatGPT Plus, Claude 3, Google Gemini

Having access to multiple LLM tools lets you pick the right one for each task. Here's a quick overview:

Tool	Highlights	Great For
ChatGPT Plus	Custom GPTs, plugins, code interpreter, image generation	General productivity, coding, multi-tool workflows
Claude 3 (and newer)	Very long context window, strong writing & analysis, long document processing	Long-form content, research, nuanced writing
Google Gemini	Real-time web info, multimodal (image/video), Workspace integration	Multimodal tasks, Docs/Sheets work

Tip: Many professionals use 2-3 tools together — e.g., Claude for drafting long-form content, ChatGPT for quick coding help, and Gemini for real-time research.

5. Practice Exercises

- 1. Write a zero-shot prompt asking an LLM to explain a topic of your choice in simple terms.
- 2. Write a few-shot prompt teaching the model a custom classification task (e.g., categorizing emails as Urgent/Not Urgent).
- 3. Write a Chain-of-Thought prompt for a word math problem.
- 4. Draft a marketing email prompt for a product you use daily.
- 5. Try the same prompt in ChatGPT, Claude, and Gemini — compare the differences in tone and depth.

6. Quick-Reference Cheat Sheet

Concept	Key Point
LLM	AI trained on text to predict and generate human-like language
Token	Small chunk of text (word/part of word) the model processes
Zero-shot	No examples given — direct instruction
Few-shot	2-5 examples given to establish a pattern
Chain of Thought	Ask the model to reason step by step
Hallucination	When the model confidently states incorrect information
Prompt Chaining	Breaking a task into sequential prompts

Master the fundamentals first — every advanced prompting technique builds on these core ideas.