KSHIPRA Academy

MASTER AI SKILLS GUIDE

Your Practical Playbook for Working Smarter with AI

14 MODULES INSIDE

How AI Works · Prompt Methods · Context · Privacy · Frontier Models · AI Tools by Category · Self-Assessment · When AI vs. Manual

Table of Contents

Introduction: Why This Guide Exists

Every professional today is using AI. But there is a massive gap between someone who uses AI and someone who uses AI well. That gap is not technology — it is understanding.

This guide was built for people who want to cross that gap permanently. Whether you are running a business, working in finance, audit, or operations, or simply someone who wants to get more done in less time — this is your complete field manual.

You will not find theoretical jargon here. What you will find is how AI actually works (in plain language), what separates a basic prompt from a brilliant one, which AI tool to use for what job, and how to know exactly what level you are at — and how to move to the next one.

What Makes This Guide Different

Most AI guides tell you WHAT to do. This one tells you WHY it works — so you can adapt, not just copy.

Every module ends with an Immediate Action — something you can apply within 5 minutes of reading.

How to Read This Guide

Start from Module 1 if you are new to AI. Jump to Module 11 if you want to know which AI to use right now. Go to Module 13 if you want an honest assessment of where you stand. Every module stands alone — read what you need, when you need it.

01 How Generative AI Works — The Moment You Click Send

Most people treat AI like a magic box. You type something in, something comes out. What happens in between feels mysterious. It does not have to be. Understanding the steps is the difference between guessing and engineering your outputs.

The Six Steps From Your Prompt to the Response

Step 1 — Tokenization: Your words get chopped into tokens. A token is not a word; it is a word made into piece. Eg: "Prompt engineering" becomes three tokens: "Prompt", "engine", "ering". Numbers, punctuation, and rare words are also tokens. This matters because AI charges by token and thinks by token.

Step 2 — Embedding: Each token is converted into a long list of numbers (a vector). The AI has learned meaning through millions of examples, encoded as numbers [01,1000001]. Note: We have explained the Embedding meaning for surface level understanding.

Step 3 — Attention: The model checks how every token in your prompt relates to every other token. This is the core of the Transformer architecture that powers ChatGPT, Claude, and Gemini. Think of it as the AI re-reading your message with extreme focus, mapping relationships.

Step 4 — Prediction: Based on all the patterns in its training, the model predicts the most statistically likely next token. Then the next. Then the next. It is not searching a database. It is not looking up answers. It is completing a very sophisticated pattern.

Step 5 — Temperature (Creativity Control): A setting called temperature controls how much randomness is added. Low temperature = predictable, precise, safe. High temperature = creative, varied, sometimes surprising. Most AI interfaces handle this for you, but you can ask for "more creative" or "more precise" in your prompt.

Step 6 — Output Decoding: The predicted tokens are converted back into text and streamed to your screen word by word. The "thinking" display you see in some AI tools shows the model working through its reasoning before generating the final answer.

The Most Important Insight in This Module

AI does not "know" things the way you know your name. It has learned statistical patterns from massive amounts of text. When it answers your question confidently, it is generating the most likely continuation of your prompt — not retrieving a fact from a database.

This is why AI can be wrong with full confidence. It is completing a pattern, not verifying truth. Your job as a user is to provide enough context that the pattern it completes is the right one.

Real-World Analogy

Imagine the world's most well-read person who has read 90% of the internet, every book ever published, and most human conversations. But they have a condition: they cannot look anything up after the conversation starts, and their memory wipes clean after each session. That is your AI. Brilliant pattern recognition, no live access to truth, no memory by default.

Immediate Action

Open ChatGPT or Claude. Type this prompt:
"Explain what tokens are using the example sentence: I love prompt engineering. How many tokens is that?"
This will make token counting tangible, not theoretical.

02 How AI Thinks — The Mind Behind the Response

Now you know AI generates text token by token. But how does it reason? How does it solve problems? And why does it sometimes produce confident nonsense?

Pattern Completion at Massive Scale

AI does not "think" the way humans think. It does not have beliefs, feelings, or intentions. What it does have is extraordinarily refined pattern completion. It has seen so much human reasoning, debate, problem-solving, and writing that it has encoded the shape of good thinking — without the consciousness behind it.

This is why it can write a compelling legal argument it has never seen before, or produce a marketing plan for a product that did not exist when it was trained. It is recombining patterns, not inventing from scratch.

Chain of Thought: When AI Reasons Step by Step

One of the most powerful discoveries in AI research came from a simple phrase: "Think step by step." When you ask AI to reason through a problem rather than just answer it, the quality of the output improves dramatically. Researchers at Google Brain found this in 2022 — and it changed how professional prompt engineers work.

What happens: instead of jumping to a conclusion, the model generates its reasoning process as text. That reasoning then becomes context for its final answer. It is using its own output to improve its output.

Varun Maaya, one of India's leading AI educators, consistently emphasises this principle: "Never ask AI for the answer. Ask AI to think through the problem first, then give the answer." This single shift will improve 70% of your AI interactions.

Why AI Hallucinates (And What to Do About It)

Hallucination is when AI generates confident but incorrect information. It is not lying. It is completing a pattern based on what usually comes next in similar contexts — and sometimes that pattern is wrong.

Situation	What AI Does (And How to Handle It)
Fact you can verify	AI may be right or wrong. Always verify statistics, dates, names, and legal facts independently.
Niche or recent information	AI training has a cutoff. Use Perplexity AI or Claude with web search for anything after 2024.
Mathematical calculations	AI is weak at complex arithmetic. Ask it to show all steps, or use it alongside a calculator.
Citations and references	AI often fabricates references. Never use AI-generated citations without verifying each one.
Technical specs	Versions, APIs, and software features change. Always cross-check against official documentation.

The Three Thinking Modes of Modern AI

Mode	What It Means & When to Use It
Fast Mode (Standard)	Immediate response, high confidence, used for most interactions. Great for drafting, summarising, brainstorming.
Reasoning Mode (Extended Think)	AI "thinks" before answering — visible in Claude's extended thinking and OpenAI o1/o3 models. Use for complex problems, logic, strategy.
Tool-Augmented Mode (Agents)	AI can use tools: search the web, run code, read files. Use for research tasks requiring live data.

Immediate Action

Next time you ask AI for advice on a real decision, add these words: "Before you give me your answer, list the key factors you're weighing and any assumptions you're making. Then give your recommendation."
Watch the quality difference.

03 How Prompts Work — The Language of AI

A prompt is not just your question. A prompt is the entire input the AI receives — everything it has to work with. The quality of your prompt is the single biggest variable in the quality of your output. Not the AI model. Not the settings. The prompt.

Anatomy of a Prompt

Prompt Component	What It Is & Why It Matters
System Prompt	Hidden instructions that define the AI's role, tone, and constraints before you even start. Used by companies building AI products. In ChatGPT Custom GPTs and Claude Projects, you can set these yourself.
Context	Background information you provide. The richer the context, the more precise the output. Your conversation history also counts as context.
Instruction	What you actually want done. Be specific about the task, the format, the length, and the audience.
Examples (optional)	Sample inputs and outputs that show AI exactly what you want. The most powerful way to get consistent formatting.
Constraints	What to avoid, what format to follow, what length to target. Without constraints, AI fills gaps with defaults — which may not be what you want.

Why Two Similar Prompts Produce Wildly Different Outputs

Vibhu Vashisth and the 100x Engineers community have documented this extensively: the gap between a 1% prompt and a 99% prompt is often just 20 additional words. Those words give the model the right context, constraints, and role.

The 1% vs 99% Prompt Comparison

❌ 1% Prompt: "Write an email about the project delay."

☑ 99% Prompt: "Write a professional email to a senior client informing them that our software project is delayed by 2 weeks due to a third-party API issue. Tone: apologetic but confident. Length: under 150 words. Include: reason for delay, revised timeline, what we are doing to prevent future delays. Do not use the phrase 'I apologise for any inconvenience.'"

Same topic. Completely different output quality.

The Three Layers of Prompt Quality

Layer 1 — Basic (What most people do): "Write me a report on X." You get a generic, surface-level output.

Layer 2 — Contextual (What good users do): "You are a senior finance analyst. Write a report on X for a board audience. Use data where possible. Keep it under 500 words. Focus on risk, not just opportunity." You get something specific and useful.

Layer 3 — Engineered (What power users do): You provide the role, the audience, the format, examples of what good looks like, constraints, and then iterate based on the output. You get something that rivals expert human work.

⚡ Immediate Action

Take a prompt you use regularly — one you've used 5+ times. Rewrite it using the Anatomy framework: Role + Context + Task + Format + Constraints.

Compare the old output to the new one. Save the upgraded prompt somewhere you can reuse it.

04 Prompt Methods & Frameworks — Your Complete Toolkit

There is not one way to prompt AI. There are specific methods for specific jobs. Using the right method for the right task is what separates power users from average ones. Here are the 8 methods that matter — with practical examples for each.

Method 1: Zero-Shot Prompting

What it is: A direct request with no examples. The AI relies entirely on its training to complete the task.

Best for: Simple, clear tasks where format does not matter much. Quick drafts, summaries, simple Q&A.

Zero-Shot Example

"Summarise this paragraph in 3 bullet points: [paste paragraph]"

Method 2: Few-Shot Prompting

What it is: You provide 2–5 examples of what good output looks like, then give the actual task. The AI learns your preferred format and style from the examples.

Best for: Consistent output format, brand voice matching, structured data extraction, classification tasks.

Few-Shot Example

Here is how I summarise meeting notes: Input: "We discussed the Q3 budget. Ramesh proposed 15% cut to travel. Team agreed. Next step: Ramesh to send revised budget by Friday." Output: "Decision: 15% travel budget cut. Owner: Ramesh. Deadline: Friday." Input: "Marketing shared campaign results — 3.2% CTR, 800 leads. Cost per lead came in at ₹340, below ₹400 target." Output: "Result: Campaign succeeded. CTR 3.2%, CPL ₹340 (target ₹400). 800 leads generated." Now summarise this note: [your meeting note here]

Method 3: Chain of Thought (CoT)

What it is: You ask AI to reason through its thinking before giving you the answer. The phrase "think step by step" or "reason through this first" activates this mode.

Best for: Complex analysis, strategy, problem-solving, anything where you want to see the reasoning, not just the conclusion.

Why it works: The generated reasoning becomes part of the AI's context. It literally helps itself arrive at better conclusions by thinking out loud.

CoT Example

"I need to decide whether to hire a full-time marketing person or use a freelancer for our growing business. Before giving me your recommendation, think through the key factors from both a financial and operational angle. Then give your recommendation with the top 3 reasons."

Method 4: Role Prompting

What it is: You assign AI a specific professional identity before asking your question. This activates relevant knowledge patterns in the model.

Best for: Getting expert-level perspectives, generating content for a specific domain, changing the tone and depth of response.

Role Prompting Example

"You are a senior internal auditor with 15 years of experience at Big 4 firms, specialising in process risk. Review the following process description and identify the top 5 control weaknesses: [paste process]"

Method 5: CO-STAR Framework

Developed and popularised by Sheila Teo. CO-STAR gives AI every piece of context it needs to produce a targeted, high-quality output. Use this when quality is non-negotiable.

Component	What to Fill In
Context (C)	What is the background situation?
Objective (O)	What is the exact goal of this task?
Style (S)	What writing style should be used? (formal, casual, storytelling, technical)
Tone (T)	What emotional register? (confident, empathetic, urgent, neutral)
Audience (A)	Who is reading this? What do they know? What do they care about?
Response (R)	What format should the output take? (table, bullet list, email, report, etc.)

Method 6: RTF — Role, Task, Format

The fastest structured framework for everyday use. When you do not have time for CO-STAR but still want a structured prompt.

RTF Example

Role: You are a business consultant specialising in small business growth. Task: Create a 90-day sales action plan for a training academy that sells certification courses to working professionals. Format: 3 phases, each with 3 specific weekly actions, presented as a numbered list.

Method 7: Constraint-Based Prompting

What it is: You define what NOT to do as clearly as what TO do. Constraints eliminate AI's tendency to pad, generalise, or drift.

Constraint-Based Example

"Write a LinkedIn post about why professionals should get a CIA certification. Constraints: Under 150 words. No hashtags. No emojis. Do not mention the word 'journey'. Do not use passive voice. Start with a question. End with a specific CTA."

Method 8: Meta-Prompting (Using AI to Improve Your Prompts)

What it is: You ask AI to review and improve your own prompt before you use it. The 100x Engineers community uses this routinely — it is one of their highest-leverage habits.

Meta-Prompting Example

"Here is a prompt I want to use. Before I run it, identify 3 ways it could be made more specific or effective, then give me the improved version: [paste your draft prompt here]"

Immediate Action

Pick ONE method from this module that you have never used before. Use it on your very next AI interaction today.

Note the difference in output quality. That difference compounds over time.

05 Context — Your Secret Weapon

If you take one concept from this guide and apply it permanently, make it this: context is not just helpful — it is the primary determinant of AI output quality. Prompt method matters. Context matters more.

What Is a Context Window?

The context window is the total amount of text the AI can "see" at one time during a conversation. Think of it as the AI's working memory — short-term, not permanent. It includes your system prompt (if any), the full conversation history, and your current message.

When a conversation exceeds the context window, the AI starts forgetting the beginning. This is why long conversations can seem to go off track — the AI is working without the full picture.

AI Model	Context Window Size
GPT-4o	128,000 tokens (~96,000 words)
Claude 3.5 Sonnet / Claude 3.7	200,000 tokens (~150,000 words)
Gemini 1.5 Pro	1,000,000 tokens (~750,000 words)
Gemini 2.0 Flash	1,000,000 tokens (~750,000 words)
Claude 3 Opus	200,000 tokens (~150,000 words)
Mistral Large	128,000 tokens (~96,000 words)
GPT-4o mini	128,000 tokens (~96,000 words)

Five Types of Context That Transform Output Quality

1. WHO YOU ARE: Tell AI your role, industry, experience level, and what you already know. "I am a finance manager at a mid-size manufacturing company" is more valuable than just starting your question.

2. WHO THE AUDIENCE IS: Who is reading this? A board of directors and a junior employee need the same information presented differently. Specify.

3. WHAT ALREADY EXISTS: If you are improving an existing document, give it to the AI. If you have a previous draft, share it. AI performs dramatically better when it has something concrete to work with.

4. WHAT HAS ALREADY BEEN DECIDED: If certain decisions are off the table, say so. "We have already decided to use Zoho CRM — do not recommend switching platforms" saves you three rounds of back-and-forth.

5. WHAT GOOD LOOKS LIKE: Share an example of the output you want. "Similar to the format used in McKinsey reports" or pasting an actual sample gives AI a quality benchmark.

Context Injection: The Power Move

Context injection means deliberately adding reference material to your prompt — documents, data, meeting notes, or previous outputs — so AI can work with your specific situation rather than generic knowledge.

Example: Instead of "How should I price my training programme?", try: "Here is our current pricing [paste], here is our competitor's pricing [paste], here is our target audience [describe]. Given this context, analyse whether our pricing is optimal and suggest 3 alternatives with rationale."

This is what Vibhu Vashisth means when he says "give AI a complete brief, not a job title." The brief is the context.

Context Management: Keeping Long Conversations on Track

When working on a complex project across multiple AI sessions: Start each new session by pasting a "context summary" — 5–10 lines covering what you are doing, what has been decided, and what the current task is. This replaces the memory the AI does not have.

Context Summary Template (Copy and Adapt)

"Context: I am [role] at [company type]. We are working on [project]. Key decisions already made: [list]. The audience for our work is [describe]. Format preferences: [describe]. Today's task: [specific task]."

Immediate Action

Create a "Context Card" for yourself: a 10-line description of who you are, your role, your company, and your most common AI tasks.

Save it somewhere accessible. Paste it at the start of every important AI conversation. The quality difference will be immediate.