

Prompt Engineering Basics

“ **Plain English:** Prompt engineering is the skill of giving AI clear, structured instructions that get useful results. It's not about memorizing "magic prompts" — it's about clear thinking. If you can write a good brief for a colleague, you can write a good prompt.

Why this matters

Here's a pattern you may recognize: you type something into an AI tool, get a mediocre response, and think "AI isn't that useful." But the problem usually isn't the AI — it's the prompt.

A vague prompt gives AI very little to work with. "Write something about marketing" could go in a thousand directions. "Write a 200-word email to my team explaining why we're shifting our Q2 campaign focus from brand awareness to lead generation, using a direct but supportive tone" gives the model enough context to produce something genuinely useful.

The difference between mediocre and excellent AI output is almost always in the prompt. And the good news is that the underlying techniques are simple to learn.

The core techniques

Be specific about what you want

This is the most impactful improvement most people can make. Compare:

Vague prompt	Specific prompt
"Summarize this document"	"Summarize this document in 3 bullet points, focusing on budget implications for our team"
"Write a response to this email"	"Write a professional but warm reply declining the meeting request, suggesting an async alternative"
"Help me with my presentation"	"Give me 5 compelling opening lines for a presentation about remote work to an audience of skeptical middle managers"

You're not writing code. You're writing a brief — the same way you'd brief a colleague who's smart but doesn't know your context.

Give AI a role (system prompts)

Telling AI *who it is* changes the quality of its output dramatically. This is the foundation of the [Your First AI Team Meeting](#) exercise.

```
You are a senior financial analyst reviewing a startup's pitch deck.  
Focus on: revenue model assumptions, burn rate, and market size claims.  
Be skeptical but constructive. Flag anything that seems unrealistic.
```

Why this works: it constrains the model's vast knowledge to a specific perspective and expertise level. Without a role, AI defaults to generic helpfulness. With a role, it brings a point of view.

Show examples (few-shot prompting)

If you want AI to produce output in a specific format or style, show it what good looks like:

```
Classify each customer comment as POSITIVE, NEGATIVE, or NEUTRAL.  
  
Comment: "Love the new dashboard, it's so much faster"  
Classification: POSITIVE  
  
Comment: "The update broke my saved filters"  
Classification: NEGATIVE  
  
Comment: "I noticed you changed the sidebar layout"  
Classification: NEUTRAL  
  
Comment: "Can't believe you removed the export feature, this is useless now"  
Classification:
```

Two or three examples are usually enough. The AI picks up the pattern — format, tone, judgment criteria — from your examples rather than having to guess what you mean.

Ask for step-by-step reasoning

For complex questions, adding "think step by step" or "show your reasoning" dramatically improves accuracy. This is called chain-of-thought prompting:

A company has 150 employees. They're cutting 12% of staff across three departments: Engineering (80 people), Sales (45 people), and Operations (25 people). The cuts should be proportional to department size.

Think step by step: how many people are cut from each department?

Without the step-by-step instruction, AI often jumps to a wrong answer on multi-step problems. With it, the model works through the math visibly, and you can check each step.

Structure your prompt clearly

When your prompt has multiple parts — context, instructions, content to process — use clear structure to separate them:

```
## Context
I'm a project manager preparing for a stakeholder review meeting tomorrow.

## The document
[paste document here]

## What I need
1. Three key risks I should be prepared to discuss
2. One piece of good news I can lead with
3. Any numbers that seem inconsistent and should be double-checked
```

This matters because AI processes your entire prompt as one stream of text. Without clear separation, it might confuse your instructions with the content you're asking it to analyze. Delimiters (like headers, brackets, or triple backticks) prevent this.

Common mistakes

Mistake	Why it fails	Better approach
"Write something good about X"	Too vague — "good" means nothing to a model	Specify length, audience, tone, and purpose
Putting instructions <i>after</i> the content	AI may lose track of late instructions in long prompts	Instructions first, content second
Asking for everything at once	Quality drops when the task is too broad	Break complex tasks into smaller prompts, chain the outputs

Mistake	Why it fails	Better approach
No examples for ambiguous tasks	The model interprets differently than you expect	Add 2-3 examples of what you want
Accepting the first output	First drafts are rarely best	Ask for revisions: "Make this more concise" or "Rewrite focusing on X"

Temperature: creativity vs. precision

Most AI tools let you adjust "temperature" — how creative vs. predictable the output is. You may not always have direct access to this setting, but understanding it helps:

Setting	Behavior	Good for
Low (precise)	Picks the most predictable words	Factual Q&A, classification, code, data extraction
Medium	Balanced	General writing, summarization, analysis
High (creative)	More varied and surprising	Brainstorming, creative writing, generating diverse options

Temperature doesn't change what the AI knows — it changes how it samples from possibilities. Low temperature = safe, predictable choices. High temperature = more variety, more risk of weirdness.

How this connects to the playbook

Every exercise in this playbook is, at some level, a prompt engineering exercise. Here's how the techniques map:

Technique	You'll practice it in
Giving AI a role	Your First AI Team Meeting — dual expert roles
Being specific	The Reusable Prompt — building prompts worth keeping
Showing examples	The Prompt Chain — structured multi-step outputs
Step-by-step reasoning	The Signal in the Noise — extracting structured insight
Clear structure	The Framework Transplant — complex reframing prompts
Iterating on output	The Fact-Check Habit — pushing back on AI responses

The hierarchy of AI improvement

Before you look for a fancier tool, try a better prompt. This is the cheapest, fastest way to improve your AI output:

Most effort: Train a custom model
Fine-tune an existing model
RAG (give AI access to your documents)
Better prompts ← start here
Least effort: Use defaults and hope for the best

Most people are somewhere between "use defaults" and "better prompts." Moving up just one level — from vague prompts to structured, specific prompts with roles and examples — is where the biggest improvement happens.

Where to go next

- [The Reusable Prompt](#) — build a prompt you'll actually use again
- [Your First AI Team Meeting](#) — practice role-based prompting

Revision #1

Created 2026-03-16 15:47:56 UTC by Admin

Updated 2026-03-16 15:47:56 UTC by Admin