




Prompting Guide 101

The basics of prompting



REVISION HISTORY

DATE	AUTHOR	NOTES
2025-04-23	R. Zaso	Initial Publication
2025-05-14	R. Zaso	Added Memory and Workspaces
2025-06-16	R. Zaso	Added Using GenAI Tools Safely

Disclaimer: AI technologies are evolving rapidly. While care has been taken in content creation, AI-generated responses may contain inaccuracies. You are encouraged to use this content as a foundation for independent research and verification. As with any publicly available AI tool, care must be taken not to share personal, proprietary, or sensitive content. Never input confidential data (e.g., client names, strategies, proprietary code) into third-party AI platforms unless there is an approved enterprise license with data controls. Assume all inputs and outputs may be logged by the provider unless stated otherwise. For secure AI deployment, use enterprise-grade or self-hosted models with proper governance.

This document was produced through human-AI collaboration, with human review and final approval.

Table of Contents

Introduction	1
Using GenAI Tools Safely.....	2
Public Doman Usage (non-enterprise license).....	2
Using GenAI Tools with an Enterprise License	3
Using Secure Self-Hosted AI Models	3
Key Precautions	4
Structuring Your Prompts: The Power of Clarity and Context	5
Unleashing the Power of Chain-of-Thought (CoT)	7
Optimizing Outputs: Iteration and Refinement	9
Make the Most of Memory and Workspaces	10
Managing Memory.....	10
Managing Workspaces	10
Avoiding Common Prompting Pitfalls.....	12
Wrap Up: Prompting with Purpose and Confidence	14
Resources and Contact Information	15

Introduction

Chances are you've already tried out GenAI tools like ChatGPT, Claude, Gemini, Grok, Microsoft Copilot, or Perplexity. Whether you've been amazed or annoyed, you're not alone and that's part of the learning curve!

Welcome to Effective Prompting 101. This guide introduces the core principles of prompt engineering for AI models, emphasizing the importance of clarity and structure in requests. Effective prompts are well-crafted instructions to AI, where precision directly influences output quality. We provide a structured breakdown of prompt components—role definition, task, context, constraints, and format—to guide users in eliciting relevant AI responses.

Crafting an effective prompt is both an art and a science. Even clear prompts can be significantly improved by specifying the audience, tone, format, and the level of depth expected in the response. Small adjustments create significant improvements in the quality, relevance, and professionalism of AI-generated outputs (figure 1).




Figure 1 – Framework

The ability to write effective prompts is crucial because they are the primary interface for communicating with AI. Prompts bridge human language and machine understanding; their efficacy determines the relevance, accuracy, and utility of AI-generated content.

Let's dig in!

Using GenAI Tools Safely

Generative AI (GenAI) tools are transforming how organizations operate, offering powerful capabilities for ideation, drafting, and analysis. However, like any technology, their safe and responsible use is important to protect sensitive information and maintain compliance. Whether you're leveraging public AI tools with a free or paid subscription, utilizing enterprise-grade solutions, or deploying secure self-hosted models, understanding best practices is essential.

This section provides clear, actionable guidance for Finance professionals to navigate GenAI usage across different environments; public domain, enterprise licenses, and self-hosted models. By adopting these strategies, you can harness the potential of GenAI while safeguarding your organization's data and aligning with governance standards. For deeper insights into strategic adoption, refer to our companion guide, *Getting Started with GenAI Tools (May 2025)*.

Public Domain Usage (non-enterprise license)

If your organization does not have an enterprise license for generative AI (GenAI) tools like ChatGPT, you can still use these tools responsibly with free or paid subscriptions (e.g., ChatGPT Free or Plus plans). However, strict precautions are necessary to protect sensitive information:

- **Avoid inputting sensitive data:** Never enter personally identifiable information (PII), employee data, compensation details, financial information, or confidential company data into public AI tools, whether using a free or paid subscription. Assume all inputs and outputs may be logged or stored by the provider unless explicitly stated otherwise in their terms.
- **Redact or anonymize information:** Paraphrase real scenarios when testing prompts. For example, use generic terms like “a regional manager” or “a customer” instead of specific names, roles, or titles. Replace sensitive data with placeholders (e.g., “Employee X” or “Company Y”).
- **Disable data retention features:** If the tool allows, turn off chat history or training features to prevent your inputs from being used to train the model. For instance, in ChatGPT, disable “Chat history & training” in the settings for both free and paid accounts.
- **Review terms of service:** Understand how the provider stores, shares, or retains data. Free and paid subscription plans often lack the robust data controls found in enterprise licenses or self-hosted models, so exercise caution.
- **Limit use to non-sensitive tasks:** Use public AI tools only for general tasks like content ideation, drafting templates, or brainstorming non-confidential ideas. Avoid using them for regulatory-sensitive, financial, or legal use cases.
- **Conduct output reviews:** Before integrating AI-generated outputs into internal systems, review and edit them outside the platform to ensure no sensitive information is inadvertently included.

For tasks involving employee data, compensation, financial modeling, or legal policies, switch to tools with enterprise-grade data controls or secure self-hosted AI models. When testing prompts, use generic, anonymized inputs and review outputs in a secure environment before transferring to internal systems.

Using GenAI Tools with an Enterprise License

If your organization has an enterprise license for a GenAI tool (e.g., ChatGPT Enterprise or similar), you benefit from enhanced data controls and governance features designed for secure usage. However, responsible use remains critical:

- **Leverage enterprise data protections:** Enterprise licenses typically include features like data encryption, restricted data sharing, and options to prevent inputs from being used for model training. Confirm these settings are enabled through your organization's IT or compliance team.
- **Follow organizational policies:** Adhere to your company's guidelines for using enterprise AI tools, including approved use cases, data handling protocols, and access controls. Enterprise licenses often integrate with internal IT systems for monitoring and compliance.
- **Redact or anonymize sensitive data when necessary:** Even with enterprise-grade protections, avoid inputting sensitive information unless explicitly permitted by your organization's policies. For example, anonymize PII or confidential data (e.g., "Employee X" or "Company Y" instead of real names) to minimize risk.
- **Use for sensitive but approved tasks:** Enterprise tools are suitable for tasks like drafting internal reports, analyzing anonymized datasets, or generating templates for regulatory compliance, provided they align with your organization's governance framework.
- **Monitor and audit usage:** Enterprise licenses often provide audit logs and usage reports. Work with your IT or compliance team to ensure that usage complies with internal policies and external regulations.
- **Secure prompt testing:** When testing prompts for sensitive use cases, use generic or anonymized inputs and validate outputs in a secure environment before integrating them into workflows.

For maximum security, enterprise-grade tools should be configured to align with your organization's IT policies, including data residency requirements and access controls. Consult your IT team to ensure proper setup and ongoing compliance.

Using Secure Self-Hosted AI Models

For organizations requiring the highest level of data security and control, secure self-hosted AI models offer a robust alternative to public or enterprise cloud-based tools. These models are deployed on your organization's infrastructure, ensuring complete control over data and usage:

- **Full data control:** Self-hosted models keep all inputs, outputs, and model operations within your organization's secure environment, eliminating risks associated with third-party providers. This is ideal for handling PII, employee data, financial modeling, or legal policies.
- **Customize to IT policies:** Configure self-hosted models to comply with your organization's security standards, including data encryption, access controls, and data residency requirements. Work with your IT team to align the deployment with regulatory and compliance needs.

- **Redact or anonymize sensitive data as a precaution:** Even in a self-hosted environment, anonymize sensitive information (e.g., use “Employee X” instead of real names) during testing or development phase of implementation to reduce the risk of internal exposure.
- **Use for highly sensitive tasks:** Self-hosted models are well-suited for confidential tasks such as analyzing employee data, generating financial reports, or drafting legal documents, as they operate within your secure infrastructure.
- **Implement governance and monitoring:** Establish clear policies for model access, usage, and auditing. Use internal logging and monitoring tools to track interactions and ensure compliance with organizational standards.
- **Secure development and testing:** When developing or testing prompts, use anonymized or synthetic data to prevent accidental exposure of sensitive information. Validate outputs in a controlled environment before integrating them into workflows.

Deploying self-hosted AI models requires technical expertise and infrastructure investment. Collaborate with your IT and data security teams to select, configure, and maintain models that meet your organization’s needs.

Key Precautions

Whether using a free, paid, enterprise, or self-hosted AI model, always take these precautions:

- **Thoroughly review inputs:** Ensure all names, PII, sensitive company data, and confidential information are redacted or anonymized before entering prompts. A single oversight could lead to unauthorized disclosure, even in secure environments.
- **Treat public tools as external services:** Free and paid subscriptions to public AI platforms lack the governance of enterprise or self-hosted solutions. Assume all interactions are visible to the provider unless otherwise specified.
- **Opt for secure alternatives when needed:** For confidential or regulatory-sensitive tasks, prioritize enterprise-grade tools or self-hosted AI models with robust data controls.
- **Stay informed:** Regularly review the provider’s terms of service (for public and enterprise tools), your organization’s AI usage policies, and self-hosted model configurations to stay compliant with evolving standards.

By prioritizing safe and responsible use of GenAI tools whether in the public domain, through enterprise licenses, or via secure self-hosted models your organization can unlock transformative benefits while safeguarding sensitive data and ensuring compliance.

Structuring Your Prompts: The Power of Clarity and Context

Think of your prompt as a well-crafted request to a highly intelligent, but sometimes literal, assistant. The clearer and more structured your request, the better the output. Here's a breakdown of a powerful structure:

1. Role Definition (Optional but Highly Recommended):

Why it works: Explicitly assigning a role helps the AI adopt a specific persona, accessing relevant knowledge and tailoring its tone and style.

How to use it: Start by stating the desired persona.

- a. **Example:** *You are a seasoned marketing strategist...*
- b. **Example:** *As a knowledgeable historian specializing in the French Revolution...*
- c. **Example:** *Imagine you are a helpful coding tutor...*

2. Task/Instruction:

Why it works: This is the core of your prompt, clearly stating what you want the AI to do. Use action verbs.

How to use it: Be precise and unambiguous.

- a. **Instead of:** *Tell me about the benefits of exercise.*
- b. **Try:** *Explain three key benefits of regular aerobic exercise, detailing the physiological mechanisms behind each.*

3. Context/Background Information:

Why it works: Providing relevant context helps the AI understand the scope and nuances of your request, leading to more focused and relevant responses.

How to use it: Include any necessary background information the AI needs to fulfill the task effectively.

- a. **Example:** *Given that our target audience is environmentally conscious millennials...*
- b. **Example:** *Considering the following research paper on renewable energy sources...*
- c. **Example:** *Building upon our previous conversation about blockchain technology...*

4. Constraints/Parameters:

Why it works: Setting clear boundaries helps the AI stay within the desired scope, format, and style. This prevents rambling or irrelevant outputs.

How to use it: Specify limitations and desired characteristics.

- a. **Format:** *Present your answer as bullet points. -or-*
- b. **Format:** *Write a three-paragraph essay.*

- c. *Length: Keep your response under 200 words.*
- d. *Style/Tone: Write in a formal and objective tone. -or-*
- e. *Style/Tone: Adopt a friendly and encouraging tone.*
- f. *Specific Elements: Include at least three examples. -or-*
- g. *Specific Elements: Do not mention specific brand names.*

5. Output Format (If Applicable):

Why it works: Guiding the AI on how to present the information ensures it meets your specific needs for readability and usability.

How to use it: Specify the desired output structure.

- a. **Example:** *Create a table with columns for 'Feature,' 'Benefit,' and 'Implementation.'*
- b. **Example:** *Generate a JSON object with the keys: 'title,' 'summary,' and 'keywords.'*

Examples of Structured Prompts:

You are a creative content writer specializing in travel. Your task is to generate three catchy Instagram caption options for a photo of the Eiffel Tower at sunset. The captions should be under 30 words each and evoke feelings of romance and wonder. Include at least one relevant hashtag in each caption.

You are a professional business negotiator. I'm negotiating with [describe the other party] on [insert topic or proposal] where I want the outcome to be [describe your dream outcome]. Generate a list of 10 questions I should ask for me to better understand their non-negotiables. Then make a list of the possible BATNAs (best alternatives to a non-agreement) that I should consider. And finally, help me decide [key decision you want to make, for example your walkaway price].

You are a professional networker. I'm about to meet [outline person and their role] who is [describe their position or interests] and I want to build rapport. Generate a list of 15 unique questions tailored to that person that will help me stand out and demonstrate my curiosity in their work.

Unleashing the Power of Chain-of-Thought (CoT)

Sometimes, getting the AI to "think step-by-step" can dramatically improve its reasoning and the quality of its answers, especially for complex or multi-stage problems. This is where Chain-of-Thought prompting comes in.

1. The Basic Idea:

Instead of directly asking for the final answer, you prompt the AI to explicitly show its reasoning process. You guide it to break down the problem into smaller, manageable steps.

2. How to Implement:

Explicitly ask for reasoning: Encourage the AI to explain its thought process.

- a. **Example Before CoT:** What is 15% of 80?
- b. **Example with CoT:** To calculate 15% of 80, what are the steps involved?

This phrase often signals to the AI that you want a detailed thought process. **“Lead with let's think step-by-step.”**

a. Example:

Let's think step-by-step. If a train travels at 60 mph for 3 hours and then at 70 mph for 2 hours, what is the total distance traveled?

3. Few-Shot Prompting:

Provide examples of the desired thought process by showing the AI how to break down similar problems. This is particularly powerful for more nuanced tasks.

a. Example to include in your prompt:

Question: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they now have?

Let's think step-by-step:

The cafeteria started with 23 apples.

They used 20 apples, so they had $23 - 20 = 3$ apples left.

They then bought 6 more apples, so they have $3 + 6 = 9$ apples.

Answer: 9

b. Example to include in your prompt:

Question: Olivia has \$23. She bought five bagels for \$3 each. How much money does she have left?

Let's think step-by-step:

Olivia spent 5 bagels * \$3/bagel = \$15 on bagels.

She started with \$23 and spent \$15, so she has $23 - 15 = 8$ left.

Answer: \$8

c. **Example to include in your prompt:**

Question: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have in total?
Let's think step-by-step

+ 🛠️ Tools  

Question: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have in total?
Let's think step-by-step:

Sure! Let's break this down step by step:

1. Starting amount:

Roger starts with **5 tennis balls**.

2. Additional purchase:

He buys **2 cans**, and **each can has 3 tennis balls**.

So, $2 \text{ cans} \times 3 \text{ balls per can} = \mathbf{6 \text{ tennis balls}}$.

3. Total tennis balls:

Initial 5 balls + 6 new balls = **11 tennis balls**.

✅ **Final Answer: 11 tennis balls.**



4. When to Use CoT:

- Complex reasoning tasks (math problems, logical puzzles).
- Multi-step problem-solving.
- Situations where you want to understand the AI's reasoning, not just the final answer.

Optimizing Outputs: Iteration and Refinement

Prompt engineering is often an iterative process. Your first attempt might not be perfect, and that's okay! Here's how to optimize your outputs:

1. Analyze the Output:

Carefully review the AI's response. Does it fully address your prompt? Is it accurate, relevant, and in the desired format?

2. Refine Your Prompt:

Based on the shortcomings of the initial output, adjust your prompt. Be specific about what you want to change.

- a. **If the output is too vague:** Add more detail or context to your initial prompt.
- b. **If the output is off-topic:** Re-emphasize the core task or provide negative constraints (what *not* to include).
- c. **If the format is incorrect:** Explicitly state the desired output format.
- d. **If the tone is wrong:** Specify the desired tone (e.g., "Make it sound more professional," "Write in a humorous tone").

3. Experiment with Parameters (If Applicable):

Many AI models have adjustable parameters (like temperature, which controls randomness). Experimenting can influence the creativity and predictability of the output. Lower temperatures often lead to more focused and deterministic responses, while higher temperatures can introduce more creativity.

4. Break Down Complex Tasks:

If you're struggling to get the desired output for a complex task, try breaking it down into smaller, sequential prompts. You can guide the AI through the process step-by-step.

5. Use Few-Shot Prompting for Style and Format:

Provide a few examples of the desired output style or format directly in your prompt. This can be very effective in guiding the AI.

6. Example of Prompt Refinement:

- a. **Initial Prompt:** "Write a short story about a robot."
- b. **Output:** A generic story about a robot exploring a city.
- c. **Refined Prompt:** "Write a short science fiction story (under 300 words) told from the perspective of a sentient cleaning robot on a deserted space station. Focus on its feelings of loneliness and its daily routines. End with a hint of discovery."

Make the Most of Memory and Workspaces

If the AI platform you are using supports memory and/or workspaces, you can streamline your interactions and avoid repeating background details every time you start a new chat. Note: This feature may require a paid subscription.

Managing Memory

During conversations, explicitly tell ChatGPT to "remember this" to save specific information to your memories. The AI can remember key information about **you**, your **company**, and your **goals**, so you don't need to reintroduce yourself each session. For example, you might say "Remember that I prefer using a professional tone in all my correspondence." This direct command instructs the AI to store this information in its long-term memory rather than just within the current conversation. Another example might be "Remember my..."

Name, role, and industry,
Company's mission or strategic priorities,
Long-term projects or preferences."

To manage the Memories ChatGPT has about you follow these steps:

1. Navigate to the "Personalization" section.
2. In "Personalization," view all the memories ChatGPT has stored based on your conversations.
3. Add new memories to customize how ChatGPT interacts with you, such as:
 - a. Specify how you want ChatGPT to write or word responses.
 - b. Including details about your customer persona or business.
4. Explore and manage these memories to personalize ChatGPT's responses. You can delete memories that may be inaccurate or out of date.

Tip

- Use memory to keep your AI aligned with how you like to work and the context of your business. You can also update or delete memory at any time.

Managing Workspaces

Workspaces let you group related conversations, files, and projects in one place. Ideal for:

- Managing client accounts
- Running internal projects (e.g., AI strategy, content creation)
- Organizing prompts and outputs by topic

To manage Workspaces in ChatGPT follow these steps:

1. Navigate to chat.openai.com and sign in to your account.
2. Initiate Workspace Creation:

- a. Click on your profile icon located in the top-right corner.
 - b. From the dropdown menu, select "ChatGPT Team".
 - c. Follow the prompts to set up your Team workspace, which includes providing billing information and confirming your subscription.
3. Configure Workspace Settings:
 - a. Once your Team workspace is active, click on the workspace name in the lower-left corner beneath your chat history.
 - b. Select "Manage workspace" to access settings where you can adjust preferences, manage members, and oversee billing. OpenAI Help Center

To manage Team workspaces in ChatGPT follow these steps:

1. Adding Members:
 - a. Within the "Manage workspace" section, navigate to the "Members" tab.
 - b. Click "Invite member" and enter the email addresses of those you'd like to add.
 - c. Assign roles such as Member, Admin, or Owner to define access levels.
2. Modifying Roles:
 - a. To change a member's role, click on their current role and select the new role from the dropdown menu.
3. Removing Members:
 - a. Click the three dots next to a member's name and select "Remove member" to revoke access.

If you're part of multiple workspaces (e.g., personal and Team), you can switch between them:

1. Click on your profile icon in the top-right corner.
2. In the dropdown menu, you'll see a list of available workspaces.
3. Select the workspace you wish to access.

Tips

- Use workspaces to collaborate more effectively or to keep multi-step initiatives organized without losing context.
- Multiple Workspaces: It's possible to be a member of multiple ChatGPT Team workspaces. Each workspace operates independently, allowing for distinct projects or departments within an organization.
- Billing Considerations: Each Team workspace requires its own subscription. Ensure that billing information is correctly set up for each workspace to avoid interruptions.

With memory and workspaces your AI becomes more like a true assistant, one that remembers the big picture, understands your priorities, and helps you move faster without backtracking.

Avoiding Common Prompting Pitfalls

Even with the best tools, poor prompting can lead to irrelevant, vague, or unusable results. This section outlines the most frequent mistakes people make when working with AI—and how to fix them. Learning to avoid these common pitfalls is just as important as learning to write great prompts.

1. Being Too Vague or Open-Ended

Example (weak): “Tell me about climate change.”

Improved version: “You are an environmental science teacher. Write a 150-word summary of climate change impacts tailored for high school students, using simple language and one real-world example.”

- Be specific about topic, audience, tone, and output structure.

2. Combining Multiple Requests in One Prompt

Example (weak): “Summarize this article, create a headline, and generate a list of keywords.”

Improved version: “First, summarize the article in one paragraph. Then create a headline. After that, list 5 relevant SEO keywords.”

- Think step-by-step. Guide the AI through stages.

3. Forgetting to Set the Tone or Style

Example (weak): “Write a product description for a water bottle.”

Improved version: “You are a playful copywriter for a fitness brand. Write a fun, engaging 2-sentence product description for a stainless-steel water bottle targeting health-conscious millennials.”

- Match tone to the purpose and audience.

4. Providing No Context

Example (weak): “Help me write a sales pitch.”

Improved version: “You are a SaaS sales expert. Draft a short pitch for a new task management app designed for remote teams. Highlight ease of use, mobile access, and integrations with Slack and Zoom.”

- Context allows the AI to generate more relevant, aligned content.

5. Skipping Output Format Guidance

Example (weak): “What are some onboarding best practices?”

Improved version: “List five onboarding best practices for hybrid workforces in a bulleted format. Keep each point under 25 words.”

- Structure your request to match how you want to consume the output.

6. Failing to Define What Not to Include

Example (weak): “Write a blog intro about AI tools.”

Improved version: “Write a 3-sentence blog intro about AI tools for business productivity. Avoid technical jargon and don’t mention ChatGPT or OpenAI by name.”

- Define boundaries as clearly as you define goals.

7. Treating the First Output as Final

Improve the AI’s response by asking follow-up questions like:

- “Make it shorter.”
- “Rewrite this for a skeptical executive audience.”
- “Now add a case study example.”

- Iterate, refine, and clarify. Prompting is a process, not a one-and-done step.

8. Prompting Recap: How to Stay on Track

Mistake	Fix It With...
Vague prompt	Specific task, audience, context
Multi-task prompt	Break it into steps or sections
Missing tone/style	Define the voice and intent
No background context	Add relevant setup or user info
No format guidance	Specify list, table, summary, etc.
No constraints	Use 'don't include...' where relevant
Stopping at one try	Refine, clarify, follow up

Wrap Up: Prompting with Purpose and Confidence

Prompting is more than a technical skill—it's a new form of communication and collaboration with AI. Whether you're drafting content, solving problems, exploring ideas, or analyzing data, your ability to guide the model shapes the quality of the results.

This guide introduced foundational prompting techniques that apply across tools like ChatGPT, Claude, Gemini, and beyond. By structuring your prompts with clarity, providing context, and experimenting with techniques like Chain-of-Thought or Few-Shot prompting, you empower AI to act more like a strategic partner than a search engine.

Key Reminders:

- **Be specific** with your instructions. Clear prompts produce clear outputs.
- **Set the tone and format**—what you ask for is what you get.
- **Use iteration** as a strength, not a weakness. Every prompt is a chance to improve.
- **Think like a coach**—give direction, feedback, and framing.
- **Practice ethically and responsibly.** Never include sensitive data in unsecured tools.

The more intentional you are with your prompts, the more value you'll extract from generative AI—regardless of your role or industry.

As the tools evolve, so should your skills. Prompting is not just a technical exercise—it's a creative edge. Keep experimenting, keep refining, and let your curiosity lead the way.

Mastering these techniques takes practice, but by focusing on clarity, structure, and iterative refinement, you'll be well on your way to crafting prompts that unlock the true potential of AI. What specific prompting challenges are you facing right now? If you're ready to elevate your prompting skills today, contact our team at The Zaso Group. We're here to help you solve these challenges together!

Happy Prompting!

Resources and Contact Information



Curious about what AI can do for your business? Visit us at www.thezasogroup.com for strategies, resources, and expert insights.



Get inspired by transformative conversations on *Technology Reimagined with Ray Zaso*, now streaming on Spotify Podcasts.



Get inspired by transformative conversations on *Technology Reimagined with Ray Zaso*, now streaming on Apple Podcasts.

CONTACT

Ray Zaso, Strategic AI Advisor
The Zaso Group
+1 609-451-7026 office
rzaso@thezasogroup.com
www.thezasogroup.com