

AI THINKING / A DOWNLOADABLE GUIDE

AI Thinking Skill Auditor

What is your AI doing when you're not looking? :)



Run your first AI Thinking audit.

A few minutes, one skill, and a fresh look at what your AI is doing - and what all that work means for your consumption.

ai-thinking.ai

Let's see what your skills are up to :)

You've created a skill. Maybe two. Maybe you've built yourself an entire collection.

Wonderful! :)

Now let's see if we can optimize their way of thinking, so they have to think less about why they're doing what they're doing. :)

Because every time your AI stretches its brain, there go a few more tokens! :)

Well, let's find out. :)

Let's take a look at your skills

Pick one custom skill from the list your AI will show you. You can start with whichever one you use most.

Only custom skills are worth investigating for this audit - not general built-in capabilities.

Before we begin

- An AI assistant that can read your custom skills or their instruction files.
- At least one saved skill you want to examine.
- A few minutes to review the recommendations.

Using ChatGPT? Start in your usual environment. If your AI cannot access your saved skills, upload the instruction file or paste its contents. No special model or mode is required for this document.

One important thing: this is an audit, not permission to change anything. Your skills and files should stay exactly as they are.

Ready? Let's investigate! :)

Copy the complete prompt on the following pages into your AI conversation. It will show you a list of available skills, let you choose one, then provide recommendations - without asking you to complete a questionnaire.

The audit prompt

Copy everything from "You are conducting" through "Do not make any changes."

You are conducting an AI Thinking Skill Audit.

Your objective is to examine an existing AI skill and provide specific, actionable recommendations for improving its execution, output quality, and token consumption.

This is a strictly read-only audit.

Safety and execution boundaries

- Do not execute the selected skill.
- Do not edit, delete, overwrite, or create skills or files.
- Do not modify settings, permissions, workflows, or external systems.
- Do not run tests or commands that change data.
- Do not install software or access external services for the audit.
- Do not expose confidential information, credentials, or the full contents of private skill files in the report.
- Use only information already available in the current environment or supplied directly by the user.
- If an action requires additional permissions or involves modifying anything, do not perform it.

Phase 1 - Skill selection

Identify the custom skills available to you in the current environment.

Present a numbered list with these columns: Number | Skill name | What it does.

Use only skills you can verify are available. If the environment does not expose a skill inventory, explain that you cannot access the list. Ask the user to provide a skill instruction file or paste its instructions.

Stop and wait for the user to select a skill. Do not begin the audit before a selection is made.

Phase 2 - Examine the selected skill

After the user selects a skill, review its instructions and any supporting information already accessible. Do not execute the skill.

Identify:

- The goal the skill is designed to achieve.
- The thought process AI is instructed to carry out.
- The information required to complete the task.
- The sequence of execution steps.
- The expected output and how the skill determines whether the result is acceptable.

The audit prompt / continued

Look for opportunities to improve:

- Unnecessary or repeated operations.
- Repeated searches or retrieval of the same information.
- Instructions that cause excessive output.
- Missing information that leads to additional work.
- Steps that could be simplified or combined.
- Processes that could be separated into lighter and more detailed execution modes.
- Instructions that introduce unnecessary uncertainty or repeated corrections.

Consider the full process required to produce an acceptable result. Do not assume that fewer instructions or fewer steps automatically mean lower token consumption.

Phase 3 - Present recommendations

Produce an audit report immediately after examining the available information. Do not ask preliminary questions.

Present up to five recommendations, ranked by expected practical value, in this table: Priority | What we found | Recommended improvement | Expected benefit.

Each recommendation must be specific to the selected skill. Explain why you recommend the change. Distinguish between observations supported by instructions and hypotheses requiring testing.

If information is missing, mention it briefly and continue with the audit. Do not invent findings to fill the table.

Phase 4 - Summarize the results

End with three short sections:

- What already works: Identify the parts worth keeping.
- Where to start: Recommend the single improvement you would prioritize first and why.
- How to measure improvement: Suggest a simple comparison of current and revised skill, including output quality, execution steps, and consumption where measurable.

Do not claim token savings without evidence.

Final instruction

This audit is complete once the recommendations have been presented. Do not make any changes.

And? What did your AI discover? :)

Maybe your skill is already doing a great job.

Maybe it found a few unnecessary steps.

Or maybe you discovered that something you thought was obvious was never included in the instructions.

That's exactly why we're doing this.

What should you do with the recommendations?

Start with the first one. Look at what AI found, why it recommended a change, and whether that makes sense for your work.

The goal is not simply to make your skill shorter. It is to understand what AI needs to complete the task and whether the process you gave it helps achieve that goal.

If you decide to make changes, preserve the original and test the revised version before replacing it.

Not sure what a recommendation means? Ask your AI to explain it. :)

Want to go a little deeper?

Take one recommendation and ask:

"Show me which part of the skill led you to this recommendation, and explain the thought process behind your suggested change."

Now you're examining how your AI arrived at its recommendation. That's where things start getting interesting. :)

A little note from AI Thinking

This document helps you examine existing AI skills. Recommendations depend on the information available to your AI, the skill, and your working environment. Improvements in performance or token consumption are not guaranteed.

The audit is intended to be read-only and does not authorize modifications to skills, files, or systems. Always review and test recommendations before implementing them.

Need our help? [Let's talk. :\)](#)