Make AI Work for You

Use AI WorkFlows to Save Hours Each Week

By: Steven L. Miller
https://stevenlmiller.substack.com/

Disclaimer

This book is presented solely for educational and entertainment purposes although every effort has been made to ensure the information in this book was correct at press time, the author does not assume and hereby disclaims any liability to any party for any loss, damage, or disruption in any matter whatsoever caused by any errors or omissions whether such errors or omissions result from negligence, accident, or any other cause. The views expressed are those of the author alone, and should not be taken as expert instruction or commands. The author is not offering this book and the information contained therein as legal, financial, tax, or other professional services advice. While all attempts have been made to verify the information provided in this publication, neither the author nor the publisher assumes any responsibility for errors, omissions, or contrary interpretations of the subject matter herein. Neither the author nor the publisher assumes any responsibility or liability whatsoever on the behalf of the purchaser or reader of these materials. Every circumstance is different and the advice and strategies contained herein may not be suitable for your situation. You are responsible for your own choices, actions, and results. Please seek the services of competent professionals as needed. Certain links within the book lead to resources maintained by third parties over whom the author and publisher have no control, and the author and publisher make no representations or warranties as to the accuracy of those resources. Any perceived slight of any individual or organization is purely unintentional.

Table of Contents

stevenlmiller.substack.com

Introduction

Understanding the Opportunity

The Current Reality

Right now, while you're reading this, AI is reshaping how work gets done:

Content creators are using AI to research, outline, write, and edit 15-20 high quality articles weekly, rather than 4-5 monthly, without sacrificing quality.

Small business owners are streamlining customer service, creating marketing copy, and looking at market data without hiring more staff. Shifting $3,000 to $5,000 in monthly labor costs to investments in growth initiatives.

Freelancers are taking on more complex projects and delivering them faster by leveraging AI technologies to handle work that used to require full teams.

Entrepreneurs can quickly validate business ideas by creating minimum viable products for market testing in days rather than months.

PRO TIP: The AI advantage isn't about replacing human creativity, it's about amplifying human capabilities and eliminating the tedious work that prevents people from focusing on their highest value tasks.

Is Something Holding You Back?

The three most significant barriers to getting started with AI:

1: Tool Overwhelm

New AI tools launch weekly. Most people either stick with whatever they tried first (usually ChatGPT) or randomly jump between tools, never developing fundamental proficiency with any of them.

2: Prompt Poverty

You ask AI a vague question and receive a generic answer. Then you think, "AI doesn't work for my job." The issue isn't the AI, it's that most people use these advanced tools like Google search bars. They expect the AI to read their minds rather than provide clear instructions.

3: Single-Tool Thinking

No single AI tool is good at every task. ChatGPT is helpful for brainstorming, but it often loses consistency in longer writing. Claude produces detailed content, but it usually lacks creativity. Gemini is dependable for research and fact-checking, but it can be too cautious with creative work.

Relying on one AI tool restricts your results, just like trying to build a house with only a hammer limits what you can do. You'll become frustrated, blame the tool, and overlook the fantastic results possible when using the right AI tool for the job.

Self-Assessment: What's Your Biggest Barrier?

Review the three barriers above. Which one affects you most right now?

- **Tool Overwhelm:** You've tried multiple tools but feel scattered and unfocused
- **Prompt Poverty:** You're using AI but getting inconsistent or generic results.
- **Single-Tool Thinking:** You're committed to one tool and frustrated by its limitations

Identifying your main barrier helps you focus on the most relevant strategies in the chapters ahead.

What Becomes Possible

Visualize your transformed workdays and weeks:

Monday morning: Instead of staring at a blank page, you collaborate with AI to outline three different approaches to your new project, choosing the best elements from each.

Tuesday afternoon: What used to be a two day research project becomes a focused two hour session where AI helps you gather, synthesize, and organize information from dozens of sources, now that you've narrowed down your project approach.

Wednesday evening. Instead of working late on your project proposal after completing a two day research project. You have your project proposal completed and ready to go before lunch.

Thursday. With extra time, you can add compelling visuals to your proposal to present your data more effectively.

Friday: Instead of feeling behind, with your proposal complete, you're already planning how to move things forward next week with the time you've reclaimed.

This isn't a fantasy. It's reality for professionals who have learned to work with AI strategically.

How to Approach This Book

Every technique in this book includes specific examples you can implement immediately. The strategies are organized to build your skills progressively.

Chapters are designed to stand alone. If you have a pressing challenge that a future chapter addresses, jump ahead.

As you read, focus on application over perfection. Start with one technique, practice until it becomes natural, then add another.

The professionals pulling ahead didn't do everything at once. They laid the foundation to consistently apply a few powerful strategies with compounding benefits over time.

ACTION STEP: Before moving to the next chapter, identify your biggest work frustration right now. Is it time-consuming research? Inconsistent content quality? Repetitive administrative tasks? Keep this challenge in mind as you read. You'll discover multiple ways to address it using the strategies ahead.

Chapter 1

The Multi-AI Workflow Mindset

Most people use AI incorrectly. They pick one model, usually ChatGPT, because it was first, and use it for every task like writing emails, analyzing data, brainstorming, and code debugging.

Instead you should be looking to find the right AI tool for the task at hand.

Why Single Tool Thinking Fails

Think about your physical toolkit. You have a hammer, a screwdriver, a saw, and a drill. When a task comes up, you choose the tool that fits the job, not just the one you know best. The right tool depends on the job at hand.

AI models are no different. Each one is optimized for a specific set of tasks.

ChatGPT is useful for brainstorming ideas and drafting conversational text.

Claude is better suited for organizing information and keeping long documents consistent.

Perplexity helps when you need to pull together research and check sources.

Gemini integrates deeply with Google's ecosystem.

Using ChatGPT for everything is like using a hammer to tighten screws. You'll eventually get results, but you're working harder than you need to.

The Real Cost of Model Loyalty

Sticking with one AI because you learned it first leads to suboptimal outcomes.

ChatGPT produces engaging, conversational prose. Perfect for blog posts while working poorly for technical documentation that needs precision and structure.

Claude maintains a consistent voice and logical flow across large documents. Perfect for research papers and detailed reports. It works poorly for rapid brainstorming, where you want creative leaps.

Using the wrong model doesn't just take longer; it produces inferior results.

Decision Framework: Identifying the Right AI for the Task at Hand

ChatGPT is most useful for tasks that require quick idea generation or conversational writing.

For projects that require careful analysis or longer pieces where consistency matters, **Claude** is a better fit.

You can break down complex topics step by step to make them easier to understand.

If you need to write something longer, such as a piece over 3,000 words, you can maintain a steady tone throughout.

Perplexity is helpful when you need up-to-date information with sources you can check yourself.

Fact-checking specific claims
Research with citations you can verify
Market intelligence gathering
Competitive analysis with recent data

Gemini is designed for integration with Google tools and real-time data from your own files.

Integration with Google Workspace
Search across your Google Drive files
Calendar and email context

Real time web data synthesis

Three AI Workflows That Save Time

Each of these approaches are adaptable to a range of professional situations.

Sequence 1: Research Projects

Start by using Perplexity to collect relevant information and sources. Takes 15 to 20 minutes.

Next, use Claude to review what you've found and organize the material into a clear structure. Plan for 30 to 40 minutes.

Summarize your findings. Consider using ChatGPT to edit if you'd like to strike a more conversational tone.

Sequence 2: Strategy Documents

ChatGPT: Brainstorm 15-20 strategic options for the challenge at hand (10 minutes)

Claude - Create a decision framework to evaluate each option systematically (45 minutes)

Ask Gemini to identify the three strongest options to present based on the decision framework. (20 minutes)

Sequence 3: Content Creation

ChatGPT: Create first draft with multiple angle options (20 minutes)

Claude: Refine Structure, ensure logical flow, maintain consistent voice (30 minutes)

Perplexity: Verify facts, update statistics, add current sources (15 minutes)

Writers following this workflow can double their output without sacrificing quality. The key difference is using multiple AI tools rather than relying on just one.

Getting Started: A Practical First Week

You do not need to master everything at once. Here is a weekly plan you can use to get started.

Monday: **Set Up Access**

Begin by creating accounts for ChatGPT, Claude, Perplexity, and Gemini. The free versions are enough to learn how each tool works. If you have the budget, the paid versions offer additional features, but you don't need them to get started.

Tuesday: Test Creative Tasks

Pick a project where you need to generate ideas. Use ChatGPT to list ten possible approaches. Right now, the goal is to gather options, not judge them. Timing yourself can help keep the process efficient.

Wednesday: Test Analytical Tasks

With your list of 10 options, you can work through each one using a simple framework:

- Strengths
- Weaknesses
- How difficult would it be to implement?
- What resources would it take?
- How long do you expect it to take to complete?

Once you have this information, compare the options and select the three most attractive ones. Document your reasoning for the choices. This approach will help you clarify trade offs between them.

ChatGPT is efficient at generating a wide range of options. Claude is better at analyzing and comparing those options in detail. Using both in sequence covers both needs.

Thursday: Test Research Tasks

Use Perplexity to research a topic you need current information about. Note the source citations. Compare this to Gemini's output

to see which one provides the highest quality current information for your area of focus.

Friday: Execute Your First Sequence

Choose a real project. Execute using the decision framework you've developed throughout the week above. Track time spent and quality of output to determine if the workflow is enhancing productivity.

Document what works. You're building your personal AI playbook.

Common Objections and Responses

Objection 1: "Learning multiple AIs takes too much time."

Switching between tools does take some effort.

If you already use ChatGPT, it takes time to get comfortable with Claude, Perplexity, and Gemini.

If you are frequently engaged in analytical or research work, investing a few hours to increase your effectiveness will be recovered within weeks.

Objection 2: "I can't afford multiple subscriptions."

It makes sense to begin with the free tiers. ChatGPT, Claude, Perplexity, and Gemini each provide access at no cost, with usage limits. Try them out for 30 days before considering paid options.

If the multi-AI approach saves you even a few hours per month, though, the $20-80 per month subscription cost pays for itself quickly.

Objection 3: "Switching between tools disrupts my flow."

You're not switching mid-task. You're switching between distinct phases of work. The same as switching from brainstorming (whiteboard) to drafting (word processor) to finalizing (presentation software).

The switching happens at natural transition points. Brainstorm with ChatGPT, then switch to Claude when you're ready to structure and analyze. The transition takes 30 seconds, copy your work and open a new browser tab.

When to Use Which Model:

ChatGPT

- Speed over depth
- Multiple creative options
- Conversational tone
- Rapid iteration
- Brainstorming

Claude

- Depth and thorough analysis rather than speed.
- Formal documentation style, Claude tends to produce more structured and polished language.

Perplexity

- Fact-checking, Perplexity can help you quickly confirm details with linked references. If your work requires citations, Perplexity is designed to provide them directly in the output.
- Information from the past six months that may not be available elsewhere.
- When credibility is a concern, Perplexity's ability to show sources can help you assess reliability.

It's easy to end up using the wrong tool for the job. Here are some signs you might want to switch:

- If your task is taking twice as long as you expected, it's a good indication that you may be using the wrong tool.
- When you find yourself spending a lot of time editing the output, consider whether another tool might be a better fit.
- If the output lacks depth or Structure and you need those qualities, it's worth trying a different approach.
- Not getting the sources you need is a clear sign that the tool you're using may not be the right one for the job.
- If you keep getting generic responses when you need specific details, it's time to switch.

The 30 Day Challenge

Commit to this 30 day experiment:

Week 1: Baseline

Track your current AI usage. Time for every task. Note which AI you use. Record output quality (1-5 scale).

Week 2: Learn Differences

Test each AI on identical tasks—a research task in ChatGPT, then Claude, then Perplexity, the Gemini. Note the differences in depth, quality, sources, and time.

Week 3: Implement Sequences

Implement the three sequences previously discussed in this Chapter (create link), track time savings and quality improvements.

Week 4: Optimize

Review which sequences worked best for your specific work. Start documenting what works best.

Documenting Your Experiments

Create a simple template:

Task Type: [Research, Analysis, Writing, Brainstorming]
Best AI: [ChatGPT, Claude, Perplexity, Gemini]
Prompt Structure: [Your tested Prompt]
Average Time: [Minutes]
Quality Rating: [1-5 scale]
Notes: [What works, what doesn't]

After 30 days, you should have 10-15 documented patterns. After 90 days, you'll have a comprehensive playbook for your job specific tasks.

This playbook is your competitive advantage. While others debate which AI is "best," you'll have systematic approaches that consistently produce better results in less time.

What You've Learned

The Multi-AI Workflow Mindset isn't about using more tools. It's about using the right tool for each job.

Your physical toolkit contains specialized tools. Your AI toolkit should too. ChatGPT for creative work. Claude for analytical work. Perplexity for research, Gemini, when you want to leverage Google apps. Each excels at specific tasks.

Switching costs are minimal. 2 hours to learn a new AI, 30 seconds to toggle between them.

Start this week. Set up accounts. Test the differences. Execute your first sequence. Track the results.

The professionals pulling ahead aren't using better AI. They're using AI better.

Next Steps:

1. Create accounts for ChatGPT, Claude, Perplexity, and Gemini today
2. Complete the "First AI Handoff" exercise this week
3. Track your time and quality for 30 days
4. Build your personal AI playbook

The multi-AI approach works. The question is whether you'll implement it or keep using a single tool for everything.

Chapter 2

Prompting 101: From Random Questions to Precision Commands

In this Chapter, we'll explore effective prompting. Those who rely on a structured approach to prompting consistently see better results and finish their work more efficiently.

Since AI language models can now handle tasks that previously required human expertise. If you know how to communicate with them. Approaching LLMs with unstructured prompts minimizes their usefulness. Leading to the conclusion AI is overhyped.

The difference between someone who knows how to prompt effectively and someone who doesn't is stark. It's the difference between having a powerful assistant who anticipates your needs and delivers exactly what you want, versus having a confused intern.

The Four Pillars of Effective Prompting

Perfect prompting does not come from a single formula. AI systems are complex, unpredictable at times, and constantly changing. Still, four core principles tend to yield better results regardless of which model you use.

1. Role: Who Are You Talking To?

The first mistake: treating AI like Google

Don’t ask questions as if you're querying a search engine: "What is machine learning?" AI works differently. It's not retrieving information. It's generating responses based on patterns it learned during training. Those patterns include understanding context, perspective, and expertise.

Role assignment activates specific knowledge patterns from the AI's training data. When you ask LLMs to take on the role of a senior marketing executive, they pull language and examples from the marketing world. If you tell it to act as an elementary school teacher, the vocabulary and explanations shift to match that environment. The underlying system is the same. The only difference is the instructions you give it.

So what does this actually look like when you use it?

Weak Prompt: "Explain blockchain to me."

Strong Prompt: "You are a technology journalist who specializes in making complex topics accessible to business executives. Explain

blockchain technology to a CEO who has heard the term but doesn't understand how it might apply to their business."

The difference is immediate. The weak prompt gets you a generic explanation that could come from Wikipedia. A strong prompt yields a targeted explanation that considers the audience's background, interests, and specific use case.

Role assignment goes deeper than just picking a profession. You can ask AI to embody specific thinking styles, cultural perspectives, or historical figures.

Want to understand how different economists might approach inflation? Ask the AI to respond as Keynes, then as Friedman.

Need to see a problem from multiple angles? Have it play devil's advocate, then switch to being your biggest supporter.

2. Context: What Does the AI Need to Know?

Humans fill in gaps without prompting. If you give a colleague a document and ask them to review it, they will understand you are seeking feedback on content, clarity, and accuracy. They are already familiar with your standards, your audience, and your timeline. AI doesn't have this background knowledge about your situation.

The AI needs to understand not just what you want, but why you want it, who it's for, and what constraints you're operating under.

Without context, even the most sophisticated AI is just making educated guesses about your intent.

Consider these examples:

Without context: "Write a proposal for a new product."

With context: "I'm a product manager at a mid sized SaaS company. Our current customer base consists mainly of small to medium businesses using our project management software. We've identified through user research that these customers struggle with time tracking across projects. Write a product proposal for a time tracking feature integration that I can present to our engineering team. Include technical considerations, user stories, and implementation timeline."

The second example provides the essentials: your role, the company, the audience, the problem, and the required Format. With this information, you can focus on solving the real issue instead of guessing at the basics. This approach makes it far more likely the AI will produce a result you can actually use. The AI can now make informed decisions about tone, technical depth, and content.

Context also includes examples when relevant. If you want the AI to write in a specific style, show it examples of that style. If you need output in a particular format, provide templates. If you're asking for analysis, share the data or describe the methodology you prefer. The more context you provide, the more precisely the AI can tailor its response.

3. Constraints: What Are the Boundaries?

Constraints are the most underutilized prompting technique. They aren't limitations. They're catalysts forcing the AI to be more thoughtful and purposeful in its responses.

Without constraints, AI tends toward verbosity and generality. With intelligent constraints, it becomes focused and actionable.

Constraints can be about Format:

Write this in exactly three paragraphs.

They can be about scope:

Focus only on solutions implementable within 30 days and under $10,000.

They can be about perspective: Don't mention any competitors by name.

They can even be about the thinking process: Show your reasoning step by step, but keep your final recommendation to one sentence.

Here's an example

Unconstrained: Help me design a curriculum for teaching AI to business students.

Constrained: Design a 10 week curriculum aimed at MBA students who do not have a technical background but need to understand how AI affects business decisions. Structure each week around a single 90 minute session and include a practical exercise that reinforces the central concept. The course should gradually build the skills needed for a final project: developing an AI implementation plan for an actual company.

The goal is to give students enough practical knowledge to evaluate and apply AI in real business settings without requiring them to become technical experts. Focus on practical applications, not technical details, and assume students have access to ChatGPT, Claude, Perplexity, and Gemini but no coding experience.

The constrained version gives you a curriculum you can use right away with MBA students who have no technical background. It focuses on what they need to know to apply AI in real business situations.

The unconstrained version would likely produce a generic overview that wouldn't fit any particular situation well.

Constraints also help with quality control. Want more creative responses? Add constraints that force creative thinking: "Provide three solutions, but the third must be something no one in the industry is currently doing." Want more practical advice? Constrain the response to only include strategies that have been successfully implemented elsewhere.

4. Format: How Should the Output Look?

Format specification separates amateur from expert prompting. Experienced prompters don't just ask for information. They specify exactly how they want that information structured and presented.

The Format you use depends on what you are trying to accomplish. If you are presenting to executives, bullet points and clear headers will quickly convey your message.

For documenting a process, numbered steps or a decision tree make it easier to follow.

If you are running a brainstorming session, grouping ideas by feasibility and impact can help you decide what to tackle first.

Compare these approaches:

If you want a list of ideas, you might ask: "Give me ideas for improving our customer service."

If you need a more structured plan, ask: "Create a prioritized action plan for improving our customer service. Format the output as a table with columns for Action Item, Expected Impact (High/Medium/Low), Implementation Difficulty (1-5 scale), Timeline, and Owner. Include 8 to 10 specific actions we can take immediately."

The format specific Prompt doesn't just give you ideas. It delivers a decision ready document for review at your next team meeting. The

AI understands that you need practical, implementable suggestions organized in a way that facilitates planning and accountability.

Format specifications can be detailed. If you need an executive summary with a set word count, a JSON file for a technical project, or something formatted in a company specific way say so. The important part is to be clear about what you want to see in the output.

.

When to Use Which Pillar

You don't need to include every element in every Prompt. Here's how to decide what to use:

Quick Questions (30 seconds or less)
Format only: "List the top 5..."
No role, minimal context needed

Standard Work Tasks (5-15 minutes)
Role + Format: "As a [role], create a [Format]..."
Add context if it's not clear what you're asking about.
Include any limits or requirements if they affect the scope.

Complex Projects (30+ minutes)
All four pillars
Multiple iterations expected
Document your process so you can use it again later

High Stakes Deliverables (anything customer facing or executive level)

- All four pillars with extensive detail
- Multiple rounds of refinement
- Consider testing with multiple AI models

The Anatomy of Poor Prompts

Here's what bad prompting looks like. Recognizing these patterns in your own work is the first step to improvement.

The Vague Request: "Make this better."

This gives the AI no guidance about what "better" means in your context. Better how? For whom? According to what criteria?

The Assumption Trap: "Write a marketing email for our new product."

Don’t assume the AI knows your product, your audience, your brand voice, your goals, and your constraints. It knows none of these things.

The Kitchen Sink: "Write a comprehensive business plan including market analysis, competitive landscape, financial projections, marketing strategy, operational plan, risk assessment, management team overview, and implementation timeline for a tech startup."

While this is specific about what to include, it's asking for too much in a single prompt. Each section needs its own detailed prompt to be done well.

The False Precision: "Write exactly 500 words about artificial intelligence."

This constrains the wrong thing (word count) while leaving the essential stuff (audience, purpose, angle) unconstrained.

The Anatomy of Powerful Prompts

Now let's examine what excellent prompting looks like:

The Strategic Brief:

"You are a consultant who specializes in helping traditional retailers adapt to digital transformation. I'm the CEO of a 50-store regional furniture chain that has seen online sales plateau at 15% of total revenue while competitors are achieving 40-50% online sales ratios.

Most of our customers are between 45 and 65. They prefer personal service and usually take their time to research before making a purchase.

Look at why our digital transformation may be stalled. Then outline three practical strategies we can use in the next six months to boost online sales. For each, break down the main idea, the steps to put it in place, the resources needed, how long it should take, and what might get in the way.

Write your response as a strategy memo I can share with the executive team. Keep technical terms to a minimum, but include enough detail so we can start planning.

The Practical Challenge:

Imagine you are a marketing creative director with a track record of connecting with Gen Z. Our startup makes jeans from recycled ocean plastic, but focusing on environmental benefits hasn't worked with 18-25 year olds who care more about style and social status.

The task is to develop three campaign ideas that make these jeans appealing to young, fashion-motivated adults.

Do not focus mainly on environmental messaging. For each idea, include the central creative concept, the emotion you want to trigger, the key message, how you might use influencers, and one specific idea for TikTok.or TikTok.

Present these as campaign briefs I can give to our creative team to develop further.

Advanced Techniques: Beyond the Basics

Once you master the four pillars, several advanced techniques can improve your results. Chain of Thought Prompting means asking the AI to show its reasoning, not just give a final answer. Add phrases like "Think through this step by step" or "Before giving your recommendation, analyze the pros and cons of each option."

Perspective Taking: Ask the AI to consider multiple viewpoints before responding. "First, analyze this from a customer perspective, then from an employee perspective, then from a shareholder perspective. Finally, synthesize these into a balanced recommendation."

Iterative Refinement: Use follow-up prompts to improve outputs. "That's a good start, but make it more actionable," or "Rewrite that for a less technical audience," or "Add specific examples for each point."

Negative Constraints: Tell the AI what not to do. "Don't use industry jargon," "Avoid generic advice that could apply to any company," or "Don't suggest anything that would require more than two people to implement."AI systems can do a lot, but most people treat them like search engines. They type in broad questions and get back broad answers. Through extensive testing and observation, five Prompt patterns stand out as the difference between generic results and work you can actually use.

Practice Exercise: Prompt Analysis

1. Find three prompts you've used in the last week.
2. Score each on the four pillars (Role, Context, Constraints, Format) - 0 to 5 points each
3. Rewrite your lowest-scoring Prompt using all four pillars
4. Test both versions and compare the outputs

The Future of Human-AI Collaboration

The best results come from treating AI as a partner that needs clear communication to do its best work.

The most successful AI users think of prompting as a form of conversation design. They're not just asking questions. They're creating structured dialogues that bring out the AI's strengths while compensating for its weaknesses. They iterate, experiment, and continuously refine their approach based on results.

Because AI is becoming central to how we work, learn, and solve problems. Individuals who develop strong prompting capabilities will have a significant advantage over those who don't.

Good prompting is good thinking made explicit. When you craft effective prompts, you're forced to clarify your goals, consider your audience, define your constraints, and Structure your thoughts. These are valuable skills whether you're working with AI or with other humans.

Begin with the four pillars. Focus on applying them to actual work problems rather than hypothetical scenarios. Pay attention to which techniques lead to better outcomes. Use what proves effective and set aside what does not.

Position yourself to get the most value from AI. The tools are available. What sets you apart is how systematically you put them to use.

Chapter 3

Prompt Patterns That Always Deliver

Chain of Thought: Make AI Show Its Work

Chain of thought prompts require the AI to lay out its reasoning in a precise sequence. Step by step reasoning lets you see each part of the process. Making it easier to follow and understand how the answer was built

Traditional Prompt: Should we expand into the European market?

Leads to a quick answer with little explanation. You are left without a clear sense of how the answer was reached.

Step by step prompt:

Step 1: Analyze our current market position and capabilities.

Step 2: Assess European market conditions and competition.

Step 3: Identify specific expansion options.

Step 4: Evaluate risks and resource requirements for each option.

Step 5: Make a recommendation with a clear rationale.

The advantage is you can see how conclusions are reached. If something is unclear, you can ask for clarification or go back to a previous step. You are not left guessing how the answer came about.

Use the same approach when learning something new. Rather than asking for a broad explanation like "Explain blockchain," break it down into steps:

Step 1: Explain the fundamental problem blockchain solves.

Step 2: Describe how traditional solutions fall short.

Step 3: Walk through how blockchain addresses these shortcomings.

Step 4: Provide concrete examples.

Step 5: Discuss current limitations and future potential.

Using a step by step structure helps ensure you cover all the essential points and keep your reasoning organized.

Socratic Dialogue: Let AI Ask the Questions

You can also reverse the process and have the AI ask you questions instead of just providing answers.

The setup: "I want to think through [specific problem]. Instead of giving advice, ask me probing questions that will help me think more clearly. Ask one question at a time and build on my responses."

Think of this as working with a colleague who asks the right questions and doesn't rush you. The AI prompts you to consider things you might have missed, pushes you to examine your assumptions, and helps you clarify ideas you already have.

When I was unsure how to start an article on AI and creativity, I didn't just ask for an outline. I started a back and forth prompting the AI, asking it to question my thinking.

The AI asked:

- What do you mean by creativity?
- Are you focusing on individual or collective creativity?
- How are you defining impact?
- What evidence would be most convincing to your intended audience?

After about 20 minutes of this exchange, it became clear that my main point was democratization rather than replacement. The AI didn't give me that insight. It asked the questions that helped me find it myself.

Most decisions are more complex than a simple yes or no. Consider evaluating whether to accept a job offer. Instead of asking should I take the job. Dive deeper to clarify your priorities and evaluate the offer based on your own criteria.

The questions that emerge about career trajectory, personal values, risk tolerance, family considerations, and learning opportunities often reveal decision factors you hadn't consciously considered.

Persona Emulation: Channel Expert Perspectives

AI can mimic the thinking patterns of experts in almost any field. You can use it to see problems the way a seasoned professional would.

A vague Prompt like "Act like a consultant" rarely produces valuable results.

Instead, specify the background and focus: "Respond as a McKinsey partner with 15 years in retail transformation, known for rigorous data analysis and practical implementation."

The more specific you are, the better the reasoning and recommendations you get.

Here are a few examples from different fields:

Design feedback: "Respond as a senior UX designer at a major tech company who prioritizes accessibility and user research."
Writing advice: "Channel an experienced magazine editor who is known for clear, engaging prose and a no-nonsense approach to editing."

Technical review: "Think like a principal engineer with two decades of experience building systems that scale reliably."

Each persona brings its own way of thinking. For example, a venture capitalist looks at scalability and market size. A teacher pays attention to learning goals and how to keep students engaged. A project manager is focused on timelines and risks.

Self Critique and Iteration: Build Better Through Reflection

Ask the AI to review and improve its own work using a clear set of questions.

Start by getting an initial response. Then ask the AI to critique and revise it.

Effective version: "Critique your previous response on these dimensions: 1) What important perspectives or stakeholders did I miss?2) What assumptions am I making that might not hold? 3) Where is my reasoning weakest? 4) What additional information would strengthen this analysis? Then provide a revised response addressing these issues."

AI typically produces initial responses that reflect the most common patterns in its training data. Prompting it to critique its

own answers encourages consideration of less obvious perspectives and potential exceptions.

You can also direct the AI to apply domain expertise. Ask it to critique a business plan from the perspective of a skeptical investor, focusing on market assumptions, competitive threats, and execution risks.

AI can often identify gaps in its own reasoning when prompted to do so. It may highlight overgeneralizations or overlooked details.

Make your self critique prompts specific so the AI knows exactly what to examine.

Combine Patterns for Maximum Impact

You get the most value by combining several of these patterns at once.

If you want to use AI as a thinking partner to enter a new market, start by breaking the problem down step by step. Ask questions at each stage to clarify your assumptions and make your reasoning explicit.

As you work through the process, note how your thinking evolves and where your initial ideas might need adjustment. This approach turns AI from a tool that gives answers into one that helps you structure your analysis and develop your own solutions.

7 Day Implementation Plan

Day 1: Identify a real problem you are facing at work. Walk through your reasoning step by step and pay attention to how your approach changes as you make your thinking more explicit.

Day 2: Try step by step reasoning on a dataset you need to analyze. See if you get more precise results than usual.

Day 3: Start a Socratic dialogue with the AI about a decision you need to make. Let it ask questions for 15 minutes.

Day 4: Pick three expert personas relevant to your work. Test each on the same problem. Compare their differences.

Day 5: Pick an AI output from earlier in the week. Use self critique to make it better. See how much the quality improves.

Day 6: Apply three of these patterns to one ongoing project your involved with.

Day 7: Teach these patterns to someone else. Explaining them will help you understand them better yourself.

In a world where AI capabilities are rapidly democratizing, your competitive advantage lies in knowing how to use these tools effectively. Master these five patterns to transform AI into a genuine thinking partner.

Chapter 4

The Prompt Library Approach

The AI advantage gap is widening every day. The question isn't whether artificial intelligence will transform how you work and earn. It's whether you'll position yourself on the winning side of that transformation.

Large Language Models break down every word, punctuation mark, and space into tokens. They work within strict memory limits called context windows. They predict responses based on statistical patterns, not proper understanding.

This chapter focuses on taking the prompting pillars outlined in Chapter 3 to build a high quality prompt library.

Why Prompt Quality Determines Output Quality

LLMs predict the next most likely word based on patterns in their training data. Give them vague inputs, and they generate vague

outputs. Give them precise inputs, and they generate precise outputs.

Let's review a couple of important principles before starting our library.

The Garbage In, Garbage Out Principle

Vague Prompt: "Write a marketing email about our product."

What the AI Lacks:

- Product details
- Target audience
- Desired action
- Tone preferences
- Length constraints
- Key selling points

Result: Generic marketing copy that could apply to any product in any industry.

Precise Prompt: "Write a 150-word marketing email for CFOs at mid-sized manufacturing companies (50-500 employees). The product is inventory management software that reduces carrying costs by 15-20%. Emphasize ROI and ease of integration with existing ERP systems. Tone: professional but conversational. Call to action: schedule a 15-minute demo."

Result: Targeted copy that addresses specific pain points, includes concrete benefits, and provides a clear next step.

Common Prompt Failures and Fixes

Failure 1: Assuming Shared Context

Poor: "Analyze our Q3 results."

Problem: The AI doesn't know your company, industry, or what metrics matter to you.

Fixed: "Analyze Q3 results for our SaaS company (20-person team, B2B focus). Revenue was $450K (up from $380K in Q2). Customer acquisition cost increased from $800 to $950. Churn rate stayed at 3%. Identify trends and recommend three specific actions to improve unit economics."

Failure 2: No Success Criteria

Poor: "Make this better."

Problem: "Better" is subjective. The AI can't optimize toward an undefined goal.

Fixed: "Revise this proposal to: (1) reduce total length by 30%, (2) front-load the ROI analysis in the first two paragraphs, (3) remove

technical jargon, and (4) add three specific case study examples from similar-sized companies."

Failure 3: Single Shot Expectations

Poor: Accepting the first output without iteration.

Problem: AI responses are starting points, not final products.

Fixed: Review the output, identify specific weaknesses, and request targeted improvements. "The financial analysis is strong, but the implementation timeline is too vague. Add specific week by week milestones for the first 90 days."

Failure 4: Ignoring Output Verification

Poor: Using AI-generated content without fact-checking.
Problem: LLMs hallucinate information, inventing statistics, citations, or details that sound plausible but are false.

Fixed: "Provide citations for all statistics. If you cannot find a verified source for a claim, flag it as [VERIFY NEEDED] rather than generating a plausible-sounding number."

Practice Exercise: Prompt Comparison

Select a task you regularly use AI for (email drafting, data analysis, content creation, etc…

Step 1: Write your typical Prompt.

Step 2: Rewrite it, including:

- Specific context about your situation
- Explicit constraints (length, format, tone)
- An example of desired output
- Verification steps for the AI to follow

Step 3: Run both prompts with the same model.

Step 4: Compare outputs on these dimensions:

- Relevance to your actual needs
- Accuracy of information
- Usefulness without additional editing
- Time saved

Now it's time to start building your library.

Library Structure

Category 1: Communication Templates

- Client emails (various tones and purposes)
- Internal reports (weekly updates, project status)
- Meeting agendas and follow-ups

Category 2: Analysis Templates

- Data interpretation (sales, marketing, operations)
- Competitive analysis
- Risk assessment

Category 3: Content Creation Templates

- Blog posts (various formats and lengths)
- Social media (platform specific)
- Marketing copy (email, web, print)

Category 4: Problem-Solving Templates

- Decision frameworks
- Scenario planning
- Root cause analysis

Template Creation Process

Step 1: Document Your Best Result. When an AI interaction produces excellent output, save the exact prompt that generated it.

Step 2: Extract the Pattern. Identify which elements made it successful:

- What context did you provide?
- What constraints did you set?
- What examples did you include?

- What verification steps did you specify?

Step 3: Create Variable Fields. Replace specific details with brackets: [client name], [industry], [metric], [timeframe]

Step 4: Test and Refine. Use the template for three similar tasks. Adjust based on results.

Step 5: Document Use Cases. Note when this template works well and when it doesn't.

Example Template: Client Report Generation

CONTEXT:
You are a [role] preparing a report for [client name], a [client description: industry, size, specific situation]. This report covers [time period] and focuses on [primary metrics/outcomes].

BACKGROUND:
[Key context about the project/engagement]
Previous period performance: [relevant baseline data]
Client's stated goals: [specific objectives]

TASK:

Create a [length] client report with the following structure:

1. Executive Summary (200 words max)

2. Key Metrics ([number] metrics, each with YoY or QoQ comparison)
3. Narrative Analysis (300-500 words explaining trends)
4. Recommendations ([number] specific, actionable items)
5. Next Steps (bulleted list with owners and deadlines)

CONSTRAINTS:

- Professional but conversational tone
- Specific data points back all claims
- No jargon without definitions
- Maximum total length: [word count]

VERIFICATION

Before finalizing:

- Confirm percentages are calculated correctly
- Verify that recommendations connect directly to identified issues
- Check that the following steps include specific deadlines
- Flag any data points you're uncertain about with [VERIFY]

EXAMPLE OF DESIRED STYLE:

[Paste 2-3 paragraphs from a previous successful report]

Template Maintenance

Monthly Review:

- Which templates got used most frequently?
- Which produced outputs requiring minimal editing?
- Which needs refinement based on recent failures?

Quarterly Update:

- Add new templates for emerging task types
- Retire templates that are no longer relevant
- Consolidate similar templates to reduce library size

Version Control:

- Date each template update
- Note what changed and why
- Archive prior versions

Practice Exercise: Build Your First Template

- Review last month's AI usage
- Identify your most frequent AI task
- Collect three successful prompts for this task
- Identify instances where the AI output required minimal editing
- Note what made each Prompt effective
- Extract common elements
    - What context appears in all three?

    - What constraints are consistent?
    - What verification steps helped?
- Create template structure
- Write standardized sections
- Add variable fields in brackets
- Include example output
- Test with three new instances
- Note necessary refinements
- Adjust template based on results

The best part of creating templates is that you can use them as the foundation for project files. Circling back to the beginning of this chapter. You can make a project for communications and add your communication templates as project files.

Instead of having to upload the template every time you want to draft an email. You can go to the project folder and shorten your prompt to

Using the project files, draft an email to go out to clients with our latest product updates.

Key Takeaways

Context windows set the limit for how much information your AI can process at once.

Building a template library means you do not have to reinvent the wheel every time. Instead of rewriting prompts, you save what works and reuse it. Investing 20 hours to build your core templates can save you 50-100 hours each year.

Using your templates as foundations for project files will make you even more efficient by allowing you to reference them within project files instead of having to upload them for every use.

Chapter 5

Automating Your Admin: The AI Revolution in Daily Office Work

If you are a professional without administrative support, the administrative workload can quickly become overwhelming. Managing email, keeping track of your calendar, and reviewing documents often take up hours better spent on more critical assignments..

AI can reduce this workload by sorting your email, automating your scheduling, and summarizing documents for efficient review.

The main obstacle is present bias. It is common to focus on the short-term benefit of clearing your inbox instead of spending a few hours learning AI tools that could save you significant time over the longer term. That dopamine hit from achieving inbox zero today feels better than the abstract promise of reclaimed time tomorrow.

The professionals who overcome this bias and invest in learning AI tools to reduce their administrative burden will quickly separate themselves from. those who don't.

The Email Time Trap

Modern AI assistants can analyze incoming messages with remarkable sophistication. They identify urgent communications, categorize messages by project or priority, and draft contextually appropriate responses.

The pattern recognition is exceptional. AI understands that emails from your CEO need immediate attention while promotional messages from conferences can wait.

AI also recognizes when someone is requesting a meeting versus providing an update, and can even detect emotional tone that might signal urgency or frustration.

But addressing the elephant in the room: yes, AI makes errors. The practical solution is to set up review steps before any mistakes can cause problems.

For routine messages like status updates or scheduling, let AI create the first draft. You can then review and edit in about a minute instead of spending fifteen minutes writing from scratch.

For more complex or sensitive topics, outline your main points and use AI to help organize or refine the language. You're still the decision maker, still the quality controller. The AI handles scaffolding, you fill in the blanks with your expertise.

Calendar Chaos and the Hidden Costs

Scheduling appears simple on the surface, find a time that works and book it. In reality, calendar management involves complex negotiations between competing priorities, time zones, and availability constraints.

AI scheduling automation handles this complexity, but only if you approach it strategically. The most effective implementations don't just automate scheduling. They codify your time management decision making principles.

The AI becomes an extension of your judgment about how to prioritize and protect your time.

Properly configured AI scheduling assistants can handle roughly 80% of meeting coordination while enforcing boundaries you previously struggled to maintain. They'll block time for deep work, prevent back to back meetings without breaks, factor in travel time between locations, and add buffer time between different types of meetings.

The key is to be explicit about your implicit rules

- How much notice do you need for various kinds of meetings?
- What's your policy on evening or weekend scheduling?

- How do you prioritize between competing requests for your time?

The fundamental transformation happens when scheduling automation integrates with broader workflows. AI assistants that don't just schedule meetings but prepare for them.

By pulling together relevant documents, summarizing recent communications with attendees, and suggesting agenda items based on ongoing projects. The meeting becomes a preconfigured workspace, not just a calendar entry.

Document Overwhelm to Strategic Reading

The average knowledge worker encounters hundreds of pages of text daily across emails, reports, articles, and presentations. Traditional approaches to managing this information load, such as speed reading, executive summaries, and delegation, are increasingly inadequate for the pace and complexity of modern work.

AI document summarization changes not just the pace of reading, but the foundation of reading itself. When AI can process a 50-page report in seconds and extract key insights, it fundamentally alters what kinds of documents you can engage with and how deeply you can understand complex topics across multiple domains.

Instead of choosing between reading everything superficially or a few things deeply, you can now maintain both breadth and depth.

The AI provides the broad overview that helps you identify which documents deserve deeper attention, while also surfacing connections and patterns that would be impossible to notice when reading documents in isolation.

The real power emerges when you move beyond simple summarization to synthesis. AI can show you what a document says, point out how it connects to your other work, highlight inconsistencies or gaps in reasoning, and surface questions the document leaves unanswered.

AI gives you the first pass with analysis and summary. You use the summary to spot areas deserving of more attention, then focus your effort on digging deeper to uncover the insights relevant to your work.

Expertise isn't about knowing facts or understanding concepts. It's about knowing which questions to ask and how to synthesize information across different sources and contexts.

Your 4 Step Implementation Framework

A practical approach to automating administrative work with AI looks like this.

1. Start by figuring out which administrative tasks take up the most time. Track your activities for a week. Write down each time you handle email, coordinate calendars, or process

documents. This shows you how much time you're spending on admin work.

2. Try using AI on tasks where mistakes are unlikely to cause problems. Routine meeting confirmations, holding time on the calendar for internal meetings, or drafting short summaries of news articles to share are good places to start. Find out what the AI can and cannot do without risking anything important.

3. Set up a process to check AI generated work before sharing it. For email drafts, check that the facts are correct and that the tone fits. For calendar invites, confirm the times and attendees. For document summaries, compare key points to the original. Catch mistakes early and learn where the AI is likely to go wrong.

4. As you get more comfortable, try the AI on more complicated tasks. Move from simple email replies to handling messages with several people involved. Go from basic calendar management to blocking out time for your most important work. Shift from summarizing single documents to comparing and analyzing information from several sources.

Tools and Resources

Email:

Superhuman (superhuman.com): Designed for speed, with AI that sorts and categorizes email. At $30 a month it makes the most sense for people who get a lot of email.

SaneBox (sanebox.com): Filters email based on your habits and priorities. It works with any email provider and costs between $7 and $36 a month.

If you subscribe to Gemini Pro, $20 a month, and use Gmail, you can plug into Gemini's AI features for Gmail management, including:

Semantic Filtering: It identifies "VIP" contacts based on your interaction frequency and highlights high-stakes items, such as upcoming bills or time-sensitive legal/financial documents.

Contextual Replies: Suggested replies are now multi sentence and match your specific writing style rather than using generic "Sounds good!" buttons.

Proofread Mode: A premium feature that checks not just grammar, but also the **tone and professional impact** of your draft before you hit send.

Calendar Automation:

- Reclaim.ai (reclaim.ai): Schedules habits, tasks, and breaks around your meetings. Helpful in building regular routines and protecting time to focus. Free version available, with pricing tiers above it.

- Motion (usemotion.com) The most all in one approach available. Combines your tasks, projects, and calendar into a single AI managed workflow. $30 per month.

Document Analysis:

- Claude (claude.ai) is built for nuanced analysis and can handle long, complex documents. If you need to work through material that requires careful reasoning, it is a strong option. For basic summaries, though, it may be more than you need. A pro option is available for $20 per month.

- ChatGPT (openai.com) is versatile, but you may need to be specific in your prompts to avoid overly long responses. A pro option is available for $20 per month.

- Notion AI (Notion.so) is built into the Notion workspace, so if you already use Notion for organizing information, it can streamline your workflow. For more complex analysis standalone AI tools tend to be more capable.

The Integration Imperative

The value of any tool depends on how well it connects with the others you use. Managing your email, calendar, and documents is not an isolated task. They are part of the same workflow, and it's suboptimal to handle them separately.

Picture an email assistant that flags deadlines directly on your calendar. Making it easier to keep track of what needs your attention. Complimented by document summaries to help you prepare for meetings and respond to messages.

This kind of integration requires more than technical setup. It needs an explicit design that centers on how you organize and prioritize your work. The AI forces you to clarify rules and relationships that were previously implicit and intuitive. In doing so, it often reveals inefficiencies and inconsistencies in your current approaches that you may not have noticed.

Most importantly, effective administrative automation requires ongoing iteration and refinement. These aren't set it and forget it tools. They're collaborative systems that learn and improve through interaction. The AI gets better at understanding your priorities and preferences over time, but only if you actively engage with and refine its suggestions.

The Future Is Already Here

The productivity divide between AI users and non users is already forming. Professionals who learn these tools are completing more work, in less time, and with higher quality.

Next step: Pick one tool from the resources above. Choose one administrative task you do today. Spend 30 minutes setting up the AI and testing it on that specific task. Not next week. Not when you have more time. Today.

Chapter 6

Writing & Editing Like a Pro: An AI-Assisted Approach

Most knowledge workers spend 8 to 12 hours each week on writing. There are practical ways to cut that time in half while improving output quality as well..

This chapter walks you through how. You'll build a systematic approach to leverage AI writing tools to get your ideas into polished form faster.

Why Most Professionals Struggle with Writing

Before getting into AI, it helps to be clear about the real challenge. Professional writing is hard because it forces you to juggle four distinct mental tasks at once.

1. Develop and refine our ideas
2. Organize those ideas logically
3. Find precise language to express them
4. Maintain consistent tone and flow

Trying to do all of this at once overloads your working memory. Most people find it challenging to keep track of so many moving

parts. The result is writer's block, which isn't a creativity problem. It's a bandwidth problem.

The traditional writing process asks you to be both architect and carpenter at the same time. That's inefficient. It's also unnecessary.

A Better Approach: Separate Thinking from Execution

AI writing tools work best when you use them to separate the thinking phase from the execution phase. You handle the thinking. The AI handles the first-pass execution. Then you refine.

Here's what this looks like in practice:

When I wrote my last book chapter, I spent 30 minutes talking through my argument with Claude. I laid out my reasoning as I would for a colleague. The AI responded with clarifying questions, highlighted gaps, and flagged assumptions that needed attention.

In 30 minutes, the conversation resulted in a 3,000-word draft. It was not perfect, but it was coherent, organized, and reflected my thinking. I then spent 90 minutes editing and refining the draft. Time Spent: 2 hours for a chapter that previously would have taken 4.

The numbers are straightforward. If you reduce your writing time from 10 hours a week to 5 without sacrificing output, you gain 260 hours over the course of a year.

The Drafting Process: Getting Ideas Down Fast

Most people start with a blank page and immediately feel paralyzed. Experiencing immediate discomfort of that blank page overrides your knowledge writing gets easier once you start.

AI eliminates this problem. You start with a conversation instead of a blank page.

In a business report, the task is to determine whether expanding into the Seattle market makes sense. Start by examining population growth, the competitive landscape, and capacity constraints. The aim is to make the reasoning clear enough anyone can follow the logic and reach their own conclusion.

Use the AI to organize the report, backing each argument with relevant data. This way, the report remains efficient and serves as a practical tool for decision making.

When writing blog posts, focus on a central insight, use concrete examples, and lay out practical steps readers can follow. The AI helps maintain a logical structure, making it flow for readers.

For scripts and presentations, use voice mode, describe key points and transitions out loud. The AI helps find the rhythm of spoken language:

- shorter sentences
- clearer transitions
- stronger conclusions.

The key insight: AI drafting is a thinking tool. It helps you clarify your ideas by forcing you to articulate them clearly. The draft that results is a byproduct of that thinking process.

The Editing Process: Strategic Improvement

Most people equate editing with fixing grammar or tightening up sentences. That is proofreading. Editing is about making deliberate choices regarding what to include, what to cut, and how to organize your ideas so the final piece serves its purpose.

AI helps spot structural issues precisely because it has no emotional investment in your writing. It will notice problems you might overlook simply because you are too close to your own material.

Here is a systematic approach you can use to edit your own work:

Step 1: Ask AI to summarize your main argument and supporting points. If the summary does not reflect what you intended, that signals a clarity issue. Address this first, since no amount of editing will help if your core message is unclear.

Step 2: Have AI identify which sections are most and least effective. Not in terms of writing style, but in how well they support your main goal. Strengthen the weaker sections or remove them if they do not add value.

Step 3: For business writing, use AI to simulate different audience perspectives. For example, what objections might a skeptical board member raise? What concerns would a budget focused CFO have that you have not addressed? This process helps you uncover gaps in your reasoning before your actual audience does.

Step 4: Check for repetition. If you find yourself making the same point in more than one place, either remove the duplicate or make sure each instance serves a different purpose.

Step 5: Check that examples are concrete and specific. Replace any abstract concepts with tangible details.

This five step process takes 20 30 minutes for a typical business document. It catches issues traditional editing often misses because AI can maintain consistent focus across thousands of words.

Style Transfer: Maintaining Your Voice Consistently

Every professional needs multiple voices. The tone for client facing reports differs from that of internal strategy memos or blog posts. Leverage AI to maintain the context needed to use your voices appropriately .

Build a library of your best work across formats. When starting a new piece provide 2 3 examples to the AI and ask it to help maintain that voice for similar types of pieces.

For Academic writing use longer sentences and provide more background. In business writing, get to the point quickly, use shorter paragraphs, and include concrete examples. If you are writing for a general audience, using stories or analogies to make complicated topics easier to follow.

AI can help you match your writing style to each audience within a single document by adjusting the tone, structure, and examples as needed. More importantly, it catches when you accidentally default to the wrong voice for an audience.

Practical Implementation: Your Next Steps

Start with one type of writing with clear success criteria. For most professionals, that's either business reports or client communications.

For your first project:

Start by picking a document type you write often. A weekly report, a client update, or something similar.

Pull together three examples of your strongest work in that Format.

Before you start your next document, spend 15 to 20 minutes talking through your main points with an AI tool. Clarify your thinking before you write.

Let the AI generate a first draft from your discussion.

Work through the five step editing process highlighted earlier. Compare how much time this approach takes versus your usual process.

Keep track of the time you save. For example, if you usually spend two hours on this kind of document and the AI assisted method takes one hour with similar or better results, you have cut your time in half. If you write these documents regularly, the savings add up quickly.

When This Approach Doesn't Work

This process is most useful for informational or analytical writing, like reports, proposals, articles, or documentation. In other cases, it may not be as effective.

- Creative or artistic writing where your own style is the principal value
- Technical writing that relies on specialized terms the AI might not use correctly
- When your audience expects writing that feels handcrafted
- Political or sensitive communications where precise word choice has significant implications

In these situations, you can still use AI to get feedback on structure or for editing suggestions. Just don't rely on AI for your first draft.

The Bottom Line

AI writing tools don't make you a better writer by doing the writing for you. They make you a better writer by removing the mechanical barriers that prevent you from expressing what you already know clearly and efficiently.

The question isn't whether these tools will change how professional writing gets done. That's already happening. The question is whether you'll learn to use them effectively.

Everything described in this chapter is something you can implement yourself, starting today. No special training required. Just systematic practice with the frameworks outlined above.

The time savings are real. The quality improvements are measurable. The only question is whether you're willing to

experiment with a new process for 5-10 hours to discover if it works for you.

Chapter 7

AI Enhanced Presentations and Reports: A Practical Guide

Most people try using AI for presentations by asking something like, "ChatGPT, create a presentation about our Q3 sales performance." The result is usually a generic outline with little substance. It's not surprising that many walk away thinking AI can't help with presentations. But is that really the whole story?

This failure stems from a cognitive bias called **anchoring,** the tendency to rely too heavily on the first piece of information encountered. Users anchor on the idea that AI is a content generator, so they approach it as a magic content machine. When it fails to read their mind and produce exactly what they need, they give up.

The actual value of AI for presentations isn't content generation. It's structured thinking support. AI has analyzed millions of presentations across every industry and format. It understands what makes presentations effective and can help you systematically think through your specific communication challenge. But you need to engage with it as a thinking partner, not a vending machine.

Here's the difference in practice:

Ineffective approach: "Create a 10-slide presentation on our new product strategy."

Practical approach: I need to present a new product strategy to our executive team. They're cost conscious after last quarter's budget overruns. They're skeptical about expanding our product line. They need to approve $500K in development funding. Help me map out three different structural approaches: one leading with market opportunity, one starting with competitive threats, and one beginning with customer pain points.

The second approach provides the AI with the context it needs to offer helpful guidance. You're still making the strategic decisions about structure and content while leveraging AI's pattern recognition to explore options you might not have considered.

A Practical Framework for AI Enhanced Presentation Development

Rather than treating AI as a content generator, use this five step process to develop presentations efficiently.

Step 1: Context Dump (10 minutes)

Giving the AI clear context before you create content will make the process more efficient. The following questions help you clarify what you need.

Identify who you want to reach. If you know their names or roles, include that information.

Consider what they already know about the topic.

Think through any concerns or objections they might raise.

Be clear about what decision or action you want from them.

Define any limits or constraints you are working with. Time, budget, resources, etc…

Step 2: Structural Exploration (15 minutes)

Have the AI generate three different ways to structure your presentation. The goal is not to pick one as is, but to see which approach aligns best with what you want to accomplish.

For example, presenting a new customer onboarding process to an operations team.

Prompt for three structures

- One focused on efficiency
- One on customer satisfaction

- one on reducing errors

With three options, you can pick the one that resonates best with your audience. For example, if the primary concern is processing time, a focus on efficiency might be the most suitable approach.

Step 3: Content Development with AI as Research Assistant (45 minutes)

Once you've identified your structure, use AI to:

Find supporting statistics and research
Identify relevant case studies or examples
Draft initial talking points
Anticipate questions and develop responses

The key is to treat AI generated content as raw material, not a finished product. You'll refine everything based on your specific knowledge and understanding of your audience.

Step 4: Visual Strategy Development (30 minutes)

Most people stick to the same five chart types: bar charts, line graphs, pie charts, scatter plots, and simple tables. These are fine for routine updates, but they are rarely the best way to get your point across.

AI can help you get past these defaults by matching visualization options to your data and your goal. It only works if you give it enough context, though.

When asking AI for visualization recommendations, provide:

1. Data Structure: I have quarterly revenue data across five product lines over three years.
2. Communication goal: I want to show that Product C's growth trajectory changed dramatically after our pricing change in Q2 2023.
3. Audience context: Senior executives who need to make resource allocation decisions and prefer simple, scannable visuals.
4. Specific question: What visualization approach would most effectively show this inflection point while maintaining context about overall performance?
5. This method uses AI to quickly spot patterns in large sets of data visualizations. It helps you identify key changes, like shifts in growth, while leaving the actual decision making in your hands.

AI can help you find the right visualization among options you don't use regularly.

Step 5: Iterative Refinement (20 minutes)

After developing your draft presentation, ask the AI to:

Review the logical flow and identify weak transitions

Spot unsupported claims or missing evidence
Identify where you might lose audience attention
Suggest where to add concrete examples
Point out jargon or complexity that needs simplification

This review helps you catch problems that usually only show up when you practice your presentation or, even worse, when you are already presenting.

Total time investment: 2 hours

Practice Exercise: Presentation Structure

Take a presentation you need to create in the next few weeks.

1. Write a context brief following the Step 1 framework above. Include everything you know about your audience, their concerns, and your objectives.
2. Using an AI tool (ChatGPT, Claude, or Gemini), ask for three distinct structural approaches. Paste your context brief and request: "Based on this context, suggest three different structural approaches for this presentation. For each approach, explain what it emphasizes and what type of decision maker it appeals to."
3. Examine the three structures and decide which one best fits your goals. You do not need to use the AI generated outline word for word. Instead, focus on which parts could actually make your presentation stronger.

Practice Exercise: Data Visualization

Think about a dataset you will need to present in the near future. It could be sales numbers, a project timeline, a budget breakdown, or customer data. The key is to choose something quantitative.

Lay out how your data is organized, what you want your audience to take away, and who will be in the room. Use the framework above to keep your explanation focused.

Review the suggestions. Even if you don't implement every recommendation, identify which aspects of your current visualization approach could be improved.

What AI Can't Do (And Why That Matters)

The limitations of AI in presentation development are as important as its capabilities. Understanding these constraints prevents you from wasting time on approaches that won't work. Helping you focus on where AI adds genuine value.

AI cannot:

1. **Understand your organizational dynamics.** It doesn't know that your CFO hates pie charts because the previous CEO used them in a disastrous acquisition pitch. It doesn't

know that your head of sales will reject any presentation that doesn't lead with customer testimonials. This contextual knowledge comes only from experience.

2. **Make strategic choices about message emphasis.** Should you lead with the good news or the bad news? Should you emphasize short term wins or long term vision? These decisions depend on your strategic goals and your reading of the current organizational moment. AI can show you the implications of different choices, but it can't choose for you.

3. **Check every number and claim for yourself**. AI will often give you statistics that look convincing but are sometimes just plain wrong. If you are using this information in a presentation, especially one where people will question your sources, double check everything.

4. **Make sure your own voice comes through.** AI tends to produce talking points that are smooth but generic, and they rarely sound like something you would actually say. The Audience isn't just receiving information. They're assessing your credibility and expertise. Generic AI content undermines that.

These limitations aren't defects. They're features. They force you to maintain ownership of the strategic and creative aspects of your presentation while leveraging AI for the analytical and exploratory aspects where it excels.

Implementation: Your Next Steps

The difference between understanding these concepts and actually improving your presentation development process comes down to implementation. Here's what to do next:

Tomorrow:

1. Suppose you haven't signed up for ChatGPT Plus, Claude Pro, or Gemini Pro yet. (All are $20/month). Consider doing so. The free versions work, but the paid versions are significantly more capable for this use case.

2. Identify a presentation you need to create in the next few weeks as a test case.

Within 3 days:

1. Complete the context dump exercise described in Step 1 of the framework above. Spend a full 10 minutes writing out everything relevant about your audience, objectives, and constraints.
2. Ask the AI for three structural approaches using the exact framework provided.
3. Review the structures and select one (or synthesize elements from multiple options).

Within 1 week:

1. Develop your presentation following Steps 2-5 of the framework.
2. Track your time investment and compare it to your typical time for presentation development.
3. Note specifically what worked well and what didn't.

After your presentation:

1. Assess whether the AI enhanced process produced a better outcome than your traditional approach.
2. Identify which aspects of the AI workflow added the most value.
3. Determine which elements to incorporate into your process going forward.

The Competitive Reality

While you're deciding whether to experiment with AI enhanced presentation development, your colleagues and competitors are already doing it. The quality bar for presentations is rising, and the time investment required to meet that bar is falling, but only for those who've developed effective AI workflows.

Consider what this means practically. If you're competing for resources, attention, or buy-in, you're increasingly competing against people who can iterate through multiple presentation

approaches, test different data visualizations, and refine their messaging more thoroughly than you can using traditional methods.

The good news is that the barrier to entry is low. For $20 per month and a few hours of learning, you can match or exceed the presentation development capabilities of people who previously had access to expensive design and strategy consultants. The playing field is leveling, but only for those who actually get on the field.

Clear communication is now one of the main factors that decides who receives resources, whose projects move forward, and which ideas are put into practice. The presentations you put together often have more influence on these results than most people realize.

What This Really Means

AI hasn't eliminated the need for human judgment in presentation development. It has, however, eliminated the excuse that you don't have time to do it well. The question isn't whether you can afford to invest in learning these tools. It's whether you can afford not to.

The professionals who will thrive over the next decade aren't those with access to the best tools. They're those who understand how to use readily available tools to amplify their expertise and judgment.

All you need is a $20/month AI subscription, the frameworks provided in this chapter, and a willingness to experiment.

Chapter 8

Multi-Turn Conversation Management: Building Context Without Starting Over

Most people waste 15-20 minutes per day re-explaining context to AI systems. The cause isn't AI limitations. It's the absence of systematic conversation management skills that preserve context across sessions and break complex projects into manageable pieces while maintaining continuity.

Learning to manage multi-turn conversations systematically eliminates this dependency and captures the full value of AI collaboration.

The Context Preservation Framework

Context isn't just information. It's the accumulated understanding, decisions, and nuanced insights that emerge through extended collaboration. Preserving this context requires systematic handoffs between sessions.

The Four Layer Context Stack:

1. Start by clarifying your primary goal, what might get in the way, and how you will know if you succeeded. For example, you need to build a Q4 content calendar for an industrial automation company, focused on reaching purchasing managers to increase qualified leads by 15 percent.

2. Keep track of what you have already discussed, what you decided to do, and what you chose not to pursue. For instance, you have settled on content pillars like case studies, technical troubleshooting, and industry trends. If you decided product focused content was too sales heavy or ruled out video because you do not have the resources, make a note of that.

3. Document any preferences or adjustments you have made along the way. If you have set a practical, data driven tone and are avoiding jargon that purchasing managers might not know, write that down. If you are including cost savings or efficiency numbers in every piece, make that explicit.

4. Note where you stopped and what comes next. If you have finished defining your content pillars and monthly themes, your next move might be to draft post ideas for October, starting with a case study.

Time investment: 3-5 minutes to create a context stack at the end of the session. Time saved: 15-20 minutes at next session start. Net savings per session: 10-17 minutes.

The Strategic Chunking System

The second systematic error is attempting too much in single conversations. People try to tackle entire projects in one session, overload the system, and end up with generic, surface level output. The solution is strategic chunking. Breaking projects into interconnected pieces that each conversation can handle effectively while maintaining coherence.

The Three Layer Chunking Framework:

Layer 1: Strategic conversations set the foundation for a project. They focus on big picture direction, structure, and the most critical decisions. These discussions don't happen often, but when they do, you need to bring a thorough understanding of the project. For example, an initial strategy session might define your content goals, target audience, success measures, and the limits of your approach.

Layer 2: Tactical: Specific implementations, detailed content creation, problem solving for individual components. Break each section into pieces that stand on their own and provide something useful. At the same time, make sure they fit together so the overall message stays consistent. For instance, you might write a single social media post, build out an email sequence, or set up a framework for a case study.

Layer 3: Integration accounts for a small portion of conversations but a much larger share of quality assurance. The work involves connecting the pieces, identifying gaps, and ensuring everything fits together. For example, you might review a content calendar to ensure consistent messaging, identify missing topics, and ensure each content pillar gets enough attention.

Most people get trapped in Layer 2, producing isolated pieces without stepping back to Layers 1 or 3. Result, a collection of content that lacks coherence and strategic alignment.

Four Practical Continuity Techniques

1. Before making new requests, have the AI summarize its understanding of your project using the context you provide. Ensuring both you and the AI are working from the same set of assumptions.

A simple way to do this is to begin with: "Based on the context I'm providing, summarize what you understand about this project's goals, constraints, and current status before we proceed." Why use this: It helps make sure the AI understands the context and continues the conversation smoothly. It only takes about 30 seconds and can save you 5-10 minutes of fixing mistakes later.

2. Instead of overwhelming yourself with background information, bring in each piece of context when it becomes necessary. This

approach keeps things manageable and lets you focus on what matters right now.

For example, you might start by getting a sense of the big picture in Session 1. In Session 2, you add the specific details that shape your options. By Session 3, you are ready to see how everything fits together.

The reason it works is straightforward. Your brain can only keep about 4 to 7 things in mind at once. If you try to take in too much, you lose track of what's essential. By spreading out and absorbing the information, you can focus on what matters at each step and remember it when needed for the task at hand.

3. A simple way to avoid mistakes is to occasionally ask the AI to look for inconsistencies or missing pieces in what you have worked on so far.

Example Prompt: Review the three case studies we've drafted and identify any inconsistencies in messaging, tone, or factual claims between them.

Catching errors early saves time. Spending a couple of minutes to check for problems now can prevent much longer fixes later.

Pay attention to how your ideas change as you work. If the AI suggests a new approach, make sure you update earlier sections so everything stays consistent.

For example, if you change how you explain ROI in a later section, go back and update earlier parts that used the old explanation. Otherwise, you end up with mixed messages that confuse readers.

The Learning Curve Reality

This process does not come naturally. In everyday conversation, we rely on shared context and memory to fill in the gaps. Working with AI requires a different set of habits that take some getting used to.

Most people need 8 to 12 sessions before these techniques start to feel automatic. Managing the conversation feels awkward at first and needs to be adjusted to.

Expected progression:

Sessions 1-3: Feels slower than your old approach. You're consciously applying unfamiliar frameworks.

Sessions 4-7: Speed neutrality. New approaches take as long as old approaches but produce better outputs.

Sessions 8-12: Efficiency gains emerge. You're spending less time on context reconstruction and more on actual work.

Sessions 13+: Compound returns. You can tackle more complex projects and maintain higher quality across larger intellectual territories.

Start with a project that matters to you but is not essential to your day to day work. Use it to deliberately practice these techniques. Pay attention to where conversations flow smoothly and where they

break down. Develop a sense for how the LLM's cognitive patterns work.

Your Next Action

Pick one current project involving multiple AI interactions. Before your next session with AI:

1. Create a four-layer context stack for the project
2. Identify which layer (Strategic, Tactical, Integration) your following conversation belongs to.
3. Start the session with a Session Bridge by having the AI summarize its understanding.

stevenlmiller.substack.com

Chapter 9

Project Tools: Storing Context Systematically

Last chapter we discussed how people often waste 15-20 minutes per session re-explaining context to AI systems and how to overcome it.

This chapter we'll walk thru leveraging the tools major AI platforms have built to store your context once and have the AI remember it. Claude has Projects. ChatGPT has custom GPTs and project folders. Gemini has Gems.

You'll learn what each offers and how to configure them for your work to optimize your AI setup.

Note, access to project tools requires a Plus or Pro subscription.

If you are on the free tier, you can keep a 'master context' document on your computer. Follow this link to skip ahead to that section.

What Project Tools Actually Do

Project tools solve the fundamental limitation of AI assistants discussed in the last chapter. No memory between conversations. Every chat starts fresh. Knowing how these tools work can eliminate the need for repetitive explanations.

Project tools automatically inject stored context into every conversation. When you open a project-linked chat, the AI receives your background information before you type anything. From the AI's perspective, you've already explained everything it needs to know.

Three components make project tools useful:

Persistent instructions tell the AI how to behave. These might specify your preferred tone, formatting requirements, or domain expertise you want the AI to assume.

You can upload documents like brand guidelines, technical specifications, or research notes. The AI will use these as reference points during your conversations, so you do not have to repeat yourself.

Conversation history tracks your previous discussions within a project. The AI picks up where you left off at the end of the prior session.

Instead of spending the first several minutes of each session re-explaining your project, you can get started right away.

Claude Projects: Persistent Context for Ongoing Work

Use Claude Projects when you need to maintain continuity across longer or more complex projects. It keeps your materials and conversation history together so you do not have to start from scratch each time.

What you get:

With Claude Projects, you have a workspace where you can upload documents, set your own instructions, and keep your conversation history in one place. Each new conversation automatically uses this context.

The feature that distinguishes Claude Projects from competitors is that you can upload substantial reference materials. Upload a 50-page technical specification, and Claude can answer detailed questions about section 4.3 without you quoting it verbatim. Upload your company's style guide, and Claude applies those standards to every piece of content it creates.

Setting up a Claude Project:

1. Open Claude and click Projects in the left sidebar.

2. Click Create Project and name it something specific (e.g., Q4 Marketing Campaign rather than Marketing Stuff)

3. Add project instructions in the Project Instructions field. Write these as direct statements: “You are helping develop marketing content for an industrial automation company.

Our target audience is purchasing managers at manufacturing firms. Maintain a practical, data-driven tone. Avoid jargon."

4. Upload relevant documents using the paperclip icon. These might include brand guidelines, past successful content, competitor analysis, or technical documentation.

Start a conversation within the project by clicking New Chat while the project is selected.

Subscription requirements:

To use Claude Projects, you need a Claude Pro subscription ($20/month). Free-tier users do not have access to Projects.

Claude Projects is great for long-form writing, technical documentation, and any task that needs regular reference to uploaded files. Its large context window makes Claude the top choice for working with large files.

Writers use Projects to maintain character consistency throughout a novel. Consultants upload client research and create analyses that point to specific findings. Technical writers store API documentation and create guides that align with the systems they document

.

Limitations:

Claude Projects does not sync across devices automatically. You can access the same projects, but your conversation history might look different on mobile and desktop. The search tool in uploaded documents works well for direct questions, but it can miss less obvious links.

ChatGPT: Two Approaches to Persistent Context

OpenAI has two ways to keep context: Custom GPTs and Projects. Both require a $20 monthly Plus subscription.

Custom GPTs:

With Custom GPTs, you can create specialized versions of ChatGPT that retain your instructions and uploaded content. Set them up once, and then use them like separate apps.

The difference from Claude Projects: Custom GPTs are reusable tools rather than workspaces for ongoing projects. A Custom GPT for "Email Drafting" applies the same instructions to every email you write. The Claude Project for "Client Proposal - Acme Corp" contains context specific to a single engagement.

Setting up a Custom GPT:

1. Click your profile icon, then My GPTs.
2. Click Create a GPT
3. In the Configure tab, set:
    - Name and description

- Instructions (what the GPT should do and how it should behave)
- Conversation starters (optional prompts to help users begin)
- Knowledge files (upload reference documents)
- Capabilities (web browsing, image generation, code interpreter)

4. Save and access your GPT from the sidebar

ChatGPT Projects

ChatGPT Projects functions more like Claude Projects. A workspace for ongoing work rather than a reusable tool. You can group conversations, upload files that persist across chats, and set project-specific instructions.

What works well:

Custom GPTs excel at repetitive tasks with consistent requirements. If you write the same type of content regularly (weekly reports, customer emails, social media posts), a Custom GPT remembers your format preferences, tone requirements, and any boilerplate language you want included.

The ability to share Custom GPTs with teams is valuable for standardizing AI usage across an organization. You can set up a GPT to follow your company's communication guidelines, then share the link with your team. This way, everyone can generate output that stays consistent.

Limitations:

Custom GPTs tend to be rigid. The instructions you set will apply to every conversation. Useful for standardized tasks, but less effective if your project needs to adapt over time.

The bigger issue: ChatGPT sometimes ignores Custom GPT instructions, particularly for complex or multi-step tasks. You may need to reinforce key instructions within individual conversations.

Gemini Gems: Google Workspace Integration

Google's Gemini offers Gems. Custom versions of Gemini with persistent instructions. The primary advantage: integration with Google Workspace.

What you get:

Gems let you set custom instructions that persist across conversations. More significantly, Gemini can access your Google Drive, Gmail, and Google Docs directly. If your work lives in Google's ecosystem, this integration eliminates the upload and update cycle that other platforms require.

Setting up a Gem:

1. Open Gemini and click Gem manager in the sidebar
2. Click New Gem

3. Name your Gem and write instructions describing what it should do
4. Save and access your Gem from the sidebar

Subscription requirements:
Creating Gems requires Gemini Advanced ($20/month). Free-tier Gemini users cannot create or access Gems.

The Google Workspace integration (accessing Drive, Docs, Gmail) requires a Workspace account with Gemini enabled, typically through an organizational subscription.

What works well:

If your documents already live in Google Drive, Gemini's ability to access them directly is fantastic. You don't upload files to a project; you tell Gemini where to look. When those documents are updated, Gemini automatically sees the current version.

Making Gemini Gems particularly effective for:

Analyzing data in Google Sheets
Summarizing email threads from Gmail
Referencing policies or procedures stored in shared Drives
Working with documents that multiple people edit

Limitations:
Gemini's instruction following is less consistent than Claude's or ChatGPT's. Complex multi-step instructions often get simplified or partially ignored. The integration with Google Workspace, while

valuable, can create confusion about what information the AI is accessing. Particularly for privacy conscious users.

Gems also lack the document upload flexibility of Claude Projects. You can't upload a PDF from your desktop; the document needs to live in Drive first.

Choosing the Right Tool for Your Work

Each platform serves different use cases best.

Choose Claude Projects when:

You're working on extended projects requiring consistent context over weeks or months

Your work involves substantial reference materials (documentation, research, specifications)

Output quality matters more than speed

You're doing long form writing where tone and style consistency are critical

Choose ChatGPT Custom GPTs when:

You find yourself handling the same types of tasks over and over, especially when the requirements rarely change. It's worth considering how automation or AI tools might help.

You are looking for a way to make AI tools available to your team so that everyone can leverage the same documentation.

If you need more than just text generation. If your work involves browsing the web for information, creating images, or running code, having these capabilities in one place can save time and reduce friction.

Your work involves many short interactions rather than extended sessions.

Choose Gemini Gems when:

Your documents and data live in Google Workspace

You need the AI to access current versions of frequently updated files

Integration with Gmail for email analysis matters

You're working within an organization that uses Google Workspace

For Free Tier Users

Without a paid subscription, your best option is to maintain context manually. Create a text file containing your standard context: your role, your project goals, your preferences, and any reference information the AI needs. Upload the document when starting a new chat and prompt the LLM to reference it when responding.

This approach requires more effort but costs nothing. Many professionals used this method before project tools existed and produced good results. The paid features save time; they don't unlock fundamentally different capabilities.

Setting Up Your First Project

Rather than describing project setup abstractly, here's a concrete example. You're a marketing manager creating content for a B2B software company.

Step 1: Define what the project contains

Before you start using any AI tool, take a moment to write down a few key details.

Like what the project is. For example, building a Q4 content calendar aimed at enterprise customers.

Identify your audience. In this case, that could mean IT directors and CIOs at mid-market companies.

Decide on the tone and style you need. Professional but not stiff, technically credible, focused on business outcomes, for example.

List out the reference materials you intend to upload for the project.

For example

- brand guidelines
- content samples
- product documentation
- competitor analysis

Step 2: Write your project instructions

Now, turn what you wrote into clear instructions the AI can follow.

You are helping create marketing content for [Company Name], a B2B software company selling to IT directors and CIOs at mid-market companies (500-5000 employees).

Set out your content requirements:

Professional tone that demonstrates technical credibility
Focus on business outcomes (cost savings, efficiency gains, risk reduction) rather than features. Avoid jargon unless you know your audience uses those terms every day.
Reference the uploaded brand guidelines for voice and terminology standards.
Reference the competitive analysis when positioning against alternatives.

Step 3: Upload the list of reference documents you made

Ex:

Brand/style guidelines
3-5 examples of successful past content (to demonstrate tone)
Product documentation (for accurate technical claims)
Competitive analysis (for positioning)
Any research on your target audience

Step 4: Test with a real task

Don't assume your setup is ready. Test it with a real task:

Draft a 300-word LinkedIn post about [specific topic]. It should appeal to IT directors considering [your product category].

Check the output against your requirements:

- Did it maintain the right tone?
- Did it reference your materials appropriately?
- Did it follow your formatting preferences?

If not, adjust your instructions. Some common fixes:

Add examples of what you do and don't want
Be more specific about terminology preferences
Clarify which uploaded documents to prioritize

Step 5: Iterate based on actual use

Your first setup won't be perfect. After 3-5 real work sessions, review what's working and what isn't. Add instructions addressing

recurring issues. Remove instructions that don't affect the output. Upload additional reference materials if gaps become apparent.

Your Next Action

Pick one current project where you've been re-explaining context. Begin by setting up a project workspace using the platform of your choice. The specific tool matters less than your ability to organize your work in a way that makes sense to you.

Write 3-5 sentences of project instructions
Run one real task through the project
Note what worked and what needs adjustment

After your next three sessions using this project, evaluate whether the setup is saving you time. If you're still spending significant effort on context-setting, your instructions need refinement. If conversations start productively and produce usable output faster than before, systematize the approach for other recurring work.

Chapter 10

Cross-Model Collaboration

Why do people stick with single models? Once they get acceptable results from ChatGPT, Claude, or Gemini, they stop looking for better approaches. The default model becomes the solution rather than one tool in a system.

The 3-Model Framework

Cross-model collaboration follows a specific sequence. Models serve a distinct function and are used in different phases.

1. Claude for systematic development (60% of work time)
2. ChatGPT for creative restructuring (25% of work time)
3. Gemini for fact verification (15% of work time)

Phase 1: Claude for Systematic Development (60% of work time)

Claude excels at methodical reasoning and maintaining consistency across long documents. Use Claude to build the foundation of your project.

Practical application: For a 25 page strategic analysis, present Claude with the core question and let it break down the problem systematically. Claude builds arguments step by step and maintains internal consistency that other models struggle with in longer contexts. This foundation work typically takes 60% of the total project time, but prevents 80% of structural revision later.

Claude's weakness: Conservative thinking. It rarely suggests unconventional approaches.

Phase 2: ChatGPT for Creative Restructuring (25% of work time)

Take Claude's systematic draft and present it to ChatGPT with this prompt:

Here's a draft analysis. First, identify the three weakest arguments or most conservative assumptions. Second, suggest alternative structures that would make the argument more compelling. Third, identify gaps in reasoning that need stronger evidence.

ChatGPT approaches this differently than Claude. It suggests alternatives to make arguments more compelling and identifies ways conservative reasoning could be strengthened.

Leverage ChatGPT's suggestions to improve Claude's initial output. Take the most compelling ideas from ChatGPT back to Claude for further development.

Both models make factual errors, which is where Gemini proves essential.

Phase 3: Gemini for Fact Verification (15% of work time)

After Claude and ChatGPT have collaborated on structure and content, present the result to Gemini with explicit verification instructions.

Review this document for factual accuracy. Identify every specific claim, statistic, or factual assertion. Flag anything that seems questionable, outdated, or incompatible with other claims in the document. List areas needing better evidence.

Gemini catches errors that both Claude and ChatGPT miss. It identifies outdated statistics, claims that need more nuanced treatment, and logical inconsistencies across document sections.

When This Doesn't Work

The 3-model framework has specific limitations you need to know before investing time in it.

Simple tasks don't benefit from multiple models. If you're writing a single-paragraph email or editing a short document, use one model.

Subject matter expertise is still required. Cross-model collaboration improves work quality, but it doesn't replace domain knowledge. If you don't understand the subject matter well enough to evaluate the models' suggestions, the framework won't help you. You'll have three different versions of possibly incorrect information.

Real-time collaboration isn't possible. This framework requires passing content between models sequentially. If you need immediate answers or real-time ouput, stick with one model. The benefit comes from thoughtful handoffs between different reasoning approaches.

The Bottom Line

Start a new project. Use Claude for systematic development, ChatGPT for creative restructuring, and Gemini for fact verification. The quality difference will be readily apparent.

stevenlmiller.substack.com

Chapter 11

Troubleshooting Bad Outputs: A Practical Guide to AI Quality Control

Most people treat AI troubleshooting incorrectly. They think hallucinations are random bugs when they're predictable responses to specific prompt patterns. The difference between someone who gets frustrated with AI and someone who saves 20 hours per week comes down to understanding that every bad output is feedback about how to structure your prompts better.

Why Professionals Waste Time and Money on Bad AI Outputs

The pattern is consistent across industries. Someone asks AI to do something complex, gets a confidently wrong result, tries again with minor prompt tweaks, gets another flawed output, then either gives up or spends hours fact-checking and rewriting everything the AI produced.

The core issue is confirmation bias. People want AI to work, so when they get comprehensive looking outputs, they don't apply the

same scrutiny they'd use with a junior employee's work. They take a detailed competitor analysis with specific market data at face value, when they should be asking, How would the AI actually know our competitor's internal sales figures?

C**reating a predictable failure pattern:**

Pattern 1: **Authority Gap**. Happens when you ask AI to give a confident answer about something it cannot actually know.

For example, if you ask for a prediction about a future event, the AI will often sound sure of itself even though it is just guessing.

Pattern 2: **Specificity Creep**. Happens when you ask for general information but phrase your question so it suggests you want a specific example. The AI will often respond with a made up detail to fill the gap.

Pattern 3: Temporal Confusion. If you ask about something that happened after the AI's last update, it may still try to answer as if it has current information, even though it does not.

The Five Step Hallucination Diagnostic Framework

Most people try to fix hallucinations by regenerating outputs with slightly different wording. Instead, treat causes instead of symptoms using a five-step diagnostic process.

Step 1: Verify Knowledge Boundaries

Before you trust an AI response, consider whether the AI could realistically know this information. AI cannot access recent events, private company data, specialized new research, statistics from closed systems, or real-time updates. Always mark these for extra checking.

Example: What's the current market share for our top three competitors? The AI cannot know this unless you've provided the data. Any specific numbers it generates are fabricated.

Step 2: Check for Suspiciously Perfect Details

Real information is messy. Fundamental research has conflicting results, methodological limitations, and uncertainty ranges. Genuine expert opinions usually include some uncertainty and even disagreements.

If AI provides exact dates for future events, flawless statistics without any confidence intervals, quotes that seem overly polished, or claims that all experts agree on controversial topics, it is likely making things up.

Example: Studies show that 73% of customers prefer Feature A over Feature B, with researchers at MIT, Stanford, and Oxford all agreeing that this represents a fundamental shift in user behavior.

You might notice a few things here. A specific percentage is given, without a cited study or source. Three well known universities are

listed, but no individual researchers. The claim is presented as if everyone agrees, yet it is described as a significant change. These are all signs you should look for when deciding whether to trust a statement.

Step 3: Use Verification Prompts

You can ask follow up questions to help the AI show where it is less specific.

Example: What is your confidence level for each part of this answer, on a scale from 1 to 5? Which parts should I double check before relying on them? What information would make you more confident in your response? Are there parts of this answer going beyond your training?

Well designed AI systems will often flag their uncertainty when prompted directly. Spend 30 seconds to save hours of downstream correction.

Step 4: Apply the Reverse Prompt Technique

Ask the AI to critique its own response:

What are the three weakest points in your answer?
What counterarguments exist to this position?
What assumptions are you making that might not be valid?
If you had to identify which parts might be wrong, where would you start?

The reason this approach works is simple. If you ask the AI to identify its own limitations, it simply runs the same process as before, only now it searches for gaps or weaknesses rather than aiming for the best answer. The practical takeaway is that you can often get the AI to reveal its blind spots by adjusting your question.

Step 5: Cross Reference with External Validation

For any claim with a material impact, verify through independent sources. This doesn't mean checking everything. It means having a systematic approach to what gets checked.

High priority validation: Financial figures, legal requirements, medical information, statistical claims, specific attributions (who said what), dates, and deadlines.

Low-priority validation: General business concepts, well-established best practices, basic technical explanations, and historical information from before 2020.

A diagnostic framework does not make AI flawless. What it does is surface problems early, so you can spot and address them before they lead to bigger issues.

Forcing Structured, Verifiable Outputs: The Template Method

The best way to avoid poor outputs is to design your prompts so that weak responses are not possible. For example, instead of asking a broad question like 'What should I know about competitor

analysis?', set precise requirements that force the AI to give you specifics.

The Explicit Constraint Template:

Provide a competitor analysis framework with exactly five categories.

For each category:

Give me 2-3 specific research questions

Rate how well-established this analysis method is (Scale: Experimental / Emerging / Standard / Foundational)

Note whether it's based on specific methodologies or general business principles

If referencing established frameworks, name them.

Notice what this accomplishes:

1. Spells out the structure you want.

2. Forces the AI to rate its confidence, so it has to admit when it isn't sure.

3. Separates specific knowledge from general knowledge.

4. Asks for names, so you can check if they actually exist.

This approach works because it stops the AI from giving you vague answers you can't check.

If you want answers you can actually check for yourself, these five techniques will help.

Technique 1: Ask the AI to show its work

Ask for the reasoning or sources behind each answer, even if the AI cannot give you a formal citation. This helps you see how the answer was reached and decide for yourself if it holds up.

Poor Prompt: What factors influence customer retention in SaaS?

Better Prompt: What factors influence customer retention in SaaS? For each one, explain the evidence or logic behind it. If it's based on specific research, cite it. If this is a general principle, label it as such.

The AI is forced to be more careful about unsupportable claims. Even when it can't cite real research, it will be more honest about what's established knowledge versus general principles.

Technique 2: Confidence Bracketing

Require confidence levels on different components.

Rate your confidence (1-5) on each recommendation:

1 = Speculative / based on general principles

2 = Based on limited examples
3 = Based on established patterns
4 = Based on well researched evidence
5 = Based on definitive data or requirements.

If you see an AI rating its recommendations as 2 or 3, that tells you exactly which areas need closer review.

Technique 3: Immediate Follow Up

After getting the initial response, immediately prompt:

What are the three weakest parts of this analysis?
What additional information would strengthen these recommendations?
What counterarguments should we consider?
If this analysis is wrong, what would be the most likely reason?

Technique 5: Structured Response Templates

Provide explicit format requirements:

Use this exact structure for each recommendation:

Recommendation: [One sentence]
Evidence Level: [Strong / Moderate / Weak / Speculative]
Key Assumption: [What must be true for this to work]
Failure Mode: [How this could go wrong]
Validation Method: [How to verify this is working].

This template forces the AI to think through limitations. More importantly, it gives you a structured document you can review quickly.

Progressive Verification: The Systematic Quality Control Process

The most effective AI users don't hope for good outputs. They've built processes that consistently produce them through layered verification.

Layer 1: Prompt Design (Prevention)

If you take a couple of minutes upfront to structure your initial request using the templates and techniques above, you can avoid most of the common issues that show up later. Investing two or three minutes at the start often saves you fifteen to thirty minutes in checking and fixing the output.

Here's an example of a poorly designed Prompt: Write a case study about how a manufacturing company successfully implemented AI to reduce costs.

A Prompt like this often leads the AI to invent details. You might get a fake company, made-up cost savings, and a story that sounds real but isn't true.

Here's an example of a better Prompt: Create a template for a manufacturing AI implementation case study.

Include: Data points needed (label as [ACTUAL DATA NEEDED])
Cost categories to analyze (list standard categories)
Success metrics to track (standard KPIs)
Common implementation challenges (general categories)
Don't make up specific companies or results. The template is for analysis, not a completed case study.

The result is a more reliable answer you can actually use, without any invented details getting in the way.

Layer 2: Immediate Assessment (Detection)

Have standard questions you ask to probe reliability:

What's your confidence level on different parts of this?
Which specific claims should I verify before using this?
What information gaps exist in this analysis?

Spending thirty seconds here can catch problems before they spread through the rest of your work.

Layer 3: Targeted Verification (Validation)

Focus your checking on the claims that matter most. Not every detail needs to be verified, but it is worth double checking specific statistics or financial numbers, statements about what competitors or customers are doing, recommendations that could cost you

money or create risk, and anything that would be hard to explain if it turned out to be wrong.

Layer 4: Iterative Refinement (Improvement)

Use what you learn from verification to improve future prompting:

Keep a "bad patterns" log of prompts that generated fabrications
Document which types of requests need what kinds of constraints
Build a library of effective Prompt templates for everyday tasks.

The Meta Skill: Developing AI Intuition Through Pattern Recognition

Current AI systems are trained to be helpful, harmless, and honest. But these objectives sometimes conflict. When you ask for something specific and the AI doesn't know the answer, it faces a choice: be less helpful by saying "I don't know," or be less honest by generating plausible content.

Most current systems err on the side of being helpful.

Understanding this trade off helps you design prompts that make it easier for the AI to be honest about its uncertainty.

Domain Sensitivity: Where Hallucinations Concentrate

Specific subjects tend to attract more misinformation than others. Recognizing which ones deserve a closer look can help you sidestep the usual traps.

Legal advice, regulatory rules, and medical topics are especially prone to errors. In these situations, double-checking the details yourself is essential.

Other areas warranting enhanced scrutiny include:

Recent scientific research (last 2-3 years)
Specific financial data or market statistics
Current events or news

Medium risk topics include business strategy, technical details, industry best practices, and recent historical events. Here, focus on verifying the central claims rather than every minor point.

Low Risk Domains include well established technical concepts, General business frameworks, Historical information, basic process explanations, fundamental theories, and principles.

Recognition Patterns: When Outputs Feel Wrong

It is helpful to be aware of certain warning signs.

Red Flag 1: Too Clean. Real expertise usually includes caveats, uncertainty, and an acknowledgment of complexity. If an AI

response has no qualifications, no "it depends," and no noted limitations, that's suspicious.

Red Flag 2: Suspiciously Comprehensive. If you asked a broad question and got a response that covers every angle perfectly with specific examples for each point, verify carefully. Real knowledge has gaps.

Red Flag 3: Perfect Consensus Real experts disagree. If AI presents a controversial topic with a unanimous expert opinion, that's a sign of fabrication.

Prompting: What are the competing perspectives on this? should reveal the disagreements.

Red Flag 4: Exact Figures Without Sources. For example**,** 73% of customers prefer… with no study citation is a fabricated result. Real statistics come with sources, methodologies, and error ranges.

Red Flag 5: Too Recent. Claims about very recent events and current trends are hallucinations unless you've provided the data.

Building a Sustainable AI Quality Control System

The companies and individuals who master AI troubleshooting develop standardized workflows that consistently produce high quality outputs.

Quality Control Implementation:

Start by tracking your current challenges. Keep a simple log of each time the AI produces a hallucination or a quality issue. Make a note of the type of Prompt that led to the problem.

Once you have a sense of where the problems are, try building templates for your most common AI tasks. Add explicit constraints to each template to give the AI less room to go off track.

For outputs that matter most, it's worth requiring a confidence rating and using a simple checklist to verify the results. After a month or so, you'll start to see which templates are the most effective. It's helpful to jot down patterns, such as which types of questions require additional constraints.

Over time, you'll find the quality of your outputs improves significantly.

The Systematic Troubleshooting Checklist

When you get a bad AI output, work through this diagnostic framework systematically:

Prompt Analysis
[ ]Am I asking the AI for information it has no way of knowing?
[ ] Did I focus on getting a specific answer when a broader framework would be more useful?

[ ]Am I asking about anything that happened after the AI's last update?
[ ] Did my Prompt leave too much room for vague or unverifiable answers?

Checking the Output
[ ] Do any details seem oddly specific or unlikely?
[ ] Are there statistics with no source given?
[ ] Is the answer presented as certain, with no mention of uncertainty or limits?
[ ] Does it present a single viewpoint on issues where people actually disagree?

How to Check the Answer
[ ] Which claims would cause problems if they turn out to be wrong?
[ ] Which claims can I check quickly and easily?
[ ] Is it riskier to trust the answer, or to spend time checking it?
[ ] Do I need to double check before using this, or can I monitor and adjust as I go?

Improving the Prompt
[ ] Are there clear instructions or limits I should add to get a better answer?
[ ] Does asking for a confidence rating help here?
[ ] Should I break this into smaller, more focused prompts?
[ ] Should I ask the AI to point out where its answer might fall short?

Process Improvement

[ ] Does this issue come up often enough to need a standard approach?
[ ] Would a template save time next time?
[ ] What did I learn here that could make future prompts more effective?
[ ] Should I add this to my list of things that always need a second look?

The Real Competitive Advantage

The future of AI isn't about systems that never make mistakes. It's about humans who know how to work with systems that make predictable types of errors. Learning to troubleshoot bad outputs isn't just a technical skill; it's becoming a core competency for anyone who wants to use AI effectively.

The individuals who figure this out first will have a significant advantage. Not because they have access to better AI systems, but because they know how to get better results from the AI systems everyone has access to.

In a world where AI capabilities are increasingly commoditized, the real competitive advantage lies in AI craftsmanship. The ability to consistently deliver high-quality outputs through better prompting, verification, and iteration.

Practice Exercise: Systematic Output Evaluation

Take an AI-generated output you're currently using or considering using. Work through this evaluation:

Step 1: Knowledge Boundary Check:

Which parts of this output require information that the AI couldn't have? List them specifically.

Step 2: Confidence Assessment

Re-prompt the AI: For each major claim in your response, rate your confidence 1-5 and explain why. What does this reveal?

Step 3: Adversarial Review

Re-prompt: What are the three weakest points in your previous response? What would you need to know to improve them? What gaps does this expose?

Step 4: Structured Verification

For claims with material impact, identify: (1) Specific verification method, (2) Time required to verify, (3) Cost of being wrong.

Step 5: Prompt Redesign How would you restructure your original [rompt using the techniques in this chapter to produce better initial output?

This exercise takes 10-15 minutes. Most people discover that 30-40% of what they accepted from AI needs verification or revision. More importantly, they identify specific prompt improvements that would have prevented these issues.

The value isn't in doing this once. It's in developing the habit of systematic evaluation that prevents bad outputs from entering your workflow.

Chapter 12

Building Your Personal AI Playbook

If you don't have a straightforward way to save and improve your prompts, you could lose 60-80% of what you learn with AI. This chapter shows you how to build a personal prompt library to compound your effectiveness over time.

Most knowledge workers are using AI systems that can write, analyze, and create with remarkable sophistication, and treat them like fancy search engines. They type new prompts from scratch every time, never capturing what works, never building on previous successes. It's the equivalent of a carpenter throwing away their tools after every project and buying new ones.

The difference between someone who uses AI occasionally and someone who has transformed their productivity isn't the AI itself. It's prompt library management.

The Real Cost of Starting From Scratch

Most people default to finishing the task at hand rather than taking a few extra minutes to save prompts for later. The reason is simple. The task in front of you always feels more urgent, even if taking a moment to save prompts now would make things easier later.

This same thinking leads people to skip documentation, avoid making templates, and miss out on building systems that could save a lot of time in the future.

The Hidden Infrastructure No One Talks About

Software developers don't start from scratch every time. They build libraries, use version control, and keep improving their code. The same idea works for AI.

Your personal AI playbook isn't just a collection of prompts. If you want to repeat what works and avoid reinventing the wheel, you need to break down your process into clear steps. A practical playbook has three parts:

1. Organize by function, not just by topic.
2. Document the context: when and where each step applies, and where it does not.
3. Include templates you can adapt to your own situation.

The organizational scheme you choose isn't just about structure; it helps you develop a sense of mastery. Building and maintaining your library encourages you to think critically about what makes prompts effective, fostering confidence in your problem solving abilities.

The Version Control You're Not Using

If you've ever ended up with files named something like "presentation_final_v2_FINAL_really_final.pptx," you're not alone.

This kind of ad hoc versioning is common, but it breaks down quickly when you're working with prompts.

Prompts are by nature iterative. You try something, see what happens, make a change, and repeat. Without a way to track what changed and when, it's almost impossible to know which tweaks improved things.

You don't need to learn Git or set up complicated tools. What matters is adopting a version control mindset. Treat your prompts as living documents that change over time. Keep a simple record of what you changed and why, so you can see what actually works.

You do not need any special tools to use this system. It is straightforward.

The 3-Component Version System

1. Version number and date (e.g., "v2.3 - 2025-01-15")
2. Change log entry ("Modified output Format from bullet points to numbered list - improved scannability")
3. Results comparison ("v2.3 outputs achieved desired outcome of improved scanability vs. v2.2")

At first glance, this seems like overkill. But compare the time: invest 90 seconds to jot down what changed versus 15 minutes trying to rebuild a prompt from memory.

The bigger benefit is confidence. You can try out changes knowing you can always go back if something doesn't work. And when you

share a Prompt with a colleague, you can include the context behind each change, so they aren't starting from scratch.

The 5 Step Continuous Learning Loop

Here's the loop:

1. Form a hypothesis about what will improve results
2. Test by running the modified Prompt
3. Evaluate outputs against specific criteria
4. Analyze what changed and why
5. Document findings and form a new hypothesis

The main benefit is in spotting patterns across different prompts. As you refine more prompts, you may notice that AI systems respond better to clear output Format instructions, perform better with explicit evaluation criteria, and handle complex problems more effectively when you ask for step by step thinking. Patterns become principles you can apply in new situations.

Building Your Prompt Library in 30 Days

There is no need to build a comprehensive system all at once. Start with a small set and expand as you go.

Week 1: Capture Phase

Set up a simple spreadsheet with five columns:

- Prompt text
- Use case
- Date created
- Version number
- Effectiveness notes

Add every Prompt you create during the week to the spreadsheet. No organization yet, capture.

Week 2-3: Pattern Recognition

Review your captured prompts and note:

Which prompts do you use most frequently?
Which consistently produce the best results?

Week 4: Template Development

Convert your three most-used prompts into templates with variable components.

Example:
Original fixed Prompt: Analyze this quarterly report for key financial trends.
Template version: Analyze this [document type] for [analytical focus] and provide results as [output Format] with particular attention to [specific focus areas].

The Practical Next Step

Pick one type of task you do at least weekly. This week, create a Prompt for that task and save it to a simple document with the date and context. Next week, use that saved Prompt instead of creating a new one from scratch.

Prompt capture

Do this for three different task types over the next month.
The meta-skills you develop, systematic experimentation, pattern recognition, and iterative refinement, will remain valuable regardless of how AI technology evolves. You're not just optimizing current productivity. You're developing the foundational capabilities for working effectively with AI systems as they continue to advance.

Chapter 13

Algorithms of Thought: How AI is Teaching Us to Think Better

Working with AI systems reveals an unexpected aspect of human cognition. Thinking follows predictable patterns. When you prompt an AI effectively, you are exposing the algorithms that govern human reasoning.

These algorithms can be learned, practiced, and improved. The techniques that produce better AI outputs: breaking problems into steps, making context explicit, and considering multiple perspectives. Are the same metacognitive strategies that separate expert decision makers from novices?

Consider the simple instruction to think step by step. Cognitive scientists have documented this approach in chess grandmasters and experienced diagnosticians for decades.

These experts do not possess mystical intuition. They have internalized systematic processes for decomposing complex problems. AI makes this normally invisible expertise visible and transferable.

Why We Don't Think This Way Naturally

The challenge is not ability. You already think algorithmically when the stakes are high enough and time permits. The problem is consistency.

Humans skip steps when tired, ignore contrary evidence when confident, and rationalize decisions after the fact. We are optimized for speed over accuracy because, historically, fast decisions mattered more than perfect ones.

The Five-Step Framework for Algorithmic Thinking

These steps work for decisions ranging from career changes to product launches.

Step 1: Define Constraints Explicitly

Do not start with brainstorming or discussion. Start by articulating every limitation and requirement. What cannot change? What must be preserved? What is the budget? What is the timeline?

Write these down. Most decision failures trace to conflicting implicit assumptions.

Step 2: Generate Perspectives Systematically

It's easy to assume your first take on a problem is enough. In reality, most people miss important details by sticking to a single viewpoint. The solution is to deliberately consider the situation from several angles before making a decision.

Try picking three to five distinct perspectives and running the situation through each one.

The optimist: What is the best reasonable outcome?
The pessimist: What could go wrong?
The financial analyst: What are the quantifiable costs and benefits?
The implementer: What is actually required to execute this?
The affected stakeholder: How does this look to people impacted by the decision?

Step 3: Make Second Order Effects Explicit

Direct consequences are easy to identify. The costly mistakes come from overlooking what happens next. The ripple effects and unintended consequences.

For any decision, ask yourself: And then what happens? Do this at least three times. This helps you see the chain of consequences instead of stopping at the first step.

Step 4: Articulate Your Uncertainty

Most bad decisions come from being too sure of yourself, not from missing facts. Before you decide, spell out:

- What are you certain about
- What are you fairly confident about
- What are you guessing at
- What you do not know but matters

Prevent yourself from turning guesses into facts just because you said them out loud with confidence.

Step 5: Design Decision Criteria in Advance

Before you start comparing options, decide what matters most.

What will tip the scale for you? What would make you pick one choice over another?

Don't keep moving the goalposts after the fact to justify a gut feeling.

When This Approach Does Not Work

Algorithmic thinking works best when you have time to think, the stakes are high enough to make it worth the effort, and you have the facts you need.

It falls short when you have to decide right away, when the situation is new, and there are no patterns to spot, or when emotions and relationships matter more than logic.

For everyday choices, your gut, shaped by experience, usually beats any checklist. When dealing with people, empathy and reading the room matter more than any framework. Save algorithmic thinking for big decisions where you can actually analyze the facts.

Implementing This Yourself

The competitive advantage no longer belongs to those with the most information, information is abundant and accessible. The advantage belongs to those who can structure problems and questions effectively.

You do not need AI to think algorithmically, but AI makes it more efficient. An LLM can execute each perspective in Step 2 in seconds, generate comprehensive second order effects, and force explicit articulation of uncertainty.

The tools serve the process, not the other way around. Master the framework first. Then use AI to accelerate and scale it.

The Next Step

Algorithmic thinking becomes powerful only through implementation. Reading about frameworks produces no results. Applying them to real decisions creates measurable improvements in outcomes.

The revolution is not artificial intelligence. It is the systematic intelligence these tools make visible and transferable. Every interaction with AI reveals more about how effective thinking actually works, and how to do it better yourself.

Chapter 14

Understanding Claude's System Prompt

Claude's Three Core Default Behaviors

1. Extreme Preference for Structured Output

Claude's system prompt clearly tells it to use bullet points, subheadings, and numbered lists. A built in preference that shapes every response unless you choose to override it.

This default helps save time when writing for business. You can feed Claude messy research notes and get back organized reports without specifying the structure. The model imposes hierarchy automatically.

Imagine uploading 40 pages of notes with the prompt: Create a structured analysis with an executive summary and three recommendation sections, and receiving a 12 page report with clear headings, numbered recommendations, and hierarchical organization.

But this default becomes a problem when you want unstructured creativity. To override defaults like structured output, give direct instructions: write as a single, flowing narrative without bullet points or subheadings. Understanding these override techniques helps you tailor prompts for a wider range of results.

2. Safety-First Response Pattern

Claude declines requests more often than ChatGPT or Gemini. Its instructions include detailed safety rules that make it cautious when a request is unclear, speculative, or involves sensitive subjects. For many users, this shows up as hesitation that can feel frustrating.

Instead of seeing this as a flaw, consider what it means in practice. When Claude hesitates, it shows you the limits of its knowledge, for work where mistakes carry real consequences. For compliance documents, risk analysis, and regulatory filings, you want a tool that points out assumptions and uncertainties rather than inventing answers.

3. Extended Context Processing

You can drop entire reports, contracts, or codebases into Claude and ask for a synthesis across all documents simultaneously.

The key is understanding when to use this capacity. Long context is valuable for comparative analysis, pattern detection across

documents, and synthesizing multiple sources. It's not particularly valuable for short form work or iterative editing, where shorter contexts work better.

The Mechanism Behind Default Behavior: Confirmation Bias in AI Systems

AI models generate text based on training data and system prompts, so defaults like structure and safety are instructions that shape behavior and output.

This default leads to confirmation bias in the output. For example, if instructions say to hedge when unsure, you'll get cautious answers even when you want more direct responses. Recognizing these default-driven biases helps you craft prompts that minimize errors and improve output accuracy.

Understanding this eliminates frustration. You stop asking "Why won't Claude just do what I want?" and start asking "What default instruction is Claude following, and do I need to override it for this specific task?"

Implementation Framework: Working With or Against Defaults

1. Identify Which Default Applies

Before prompting, determine which of Claude's three core defaults (structure, safety, long context) matters for your task.

Writing a report or analysis? Structure default
Doing high-stakes compliance work? Safety default.
Synthesizing multiple documents? Long context default.
Writing creative content? Structure default needs overriding.
Need decisive answers? Safety default needs overriding.

Write Prompts That Either Leverage or Override Defaults

Example: User asks Claude to brainstorm wild product ideas.

Claude returns organized bullet lists of conventional suggestions. The structure default plus the safety default both push toward a conservative organization like this

To get more creative results, be explicit: ask for wild, unconventional ideas, and specify that you don't want them organized or filtered. Let the system know you're looking for raw possibilities, even if they seem impractical.

Practice Exercise: Mapping Defaults to Your Work

Think about three types of documents or analyses you handle often in your work.

For each one, consider which of Claude's defaults. Structure, safety, or long context help, and which might get in the way.

Pick one of these tasks and write two versions of a prompt for it.

Version A: Use Claude's defaults as they are.

Version B: Override the defaults by giving clear, specific instructions for what you want instead.

Test both prompts on the same task and compare output quality and editing time required.

This exercise reveals where Claude's design aligns with your needs and where you need explicit overrides.

The Productivity Mechanism: Reduced Friction in Knowledge Work

Understanding system prompts is less about getting better results from Claude and more about seeing how every AI tool comes with its own set of hidden defaults that influence what you get back.

ChatGPT tends to produce responses with a conversational style and often adds creative details. Gemini leans toward pulling data from Google. Claude gives you a more structured analysis and urges you to err on the side of caution in your reasoning. If you know how these defaults work, you spend less time wrestling with the tool and more time getting beneficial results.

When This Doesn't Work: Claude's Limitations

Claude's default settings can lead to some blind spots. Its structured approach makes it less effective at free-form creative writing than ChatGPT. Because of its safety settings, it may also decline some valid business requests that involve sensitive topics. While the long context window is useful, it can sometimes make the output less focused. For simple tasks, shorter, more targeted prompts often work better.

The Bottom Line

Claude's real advantage isn't that it's smarter than other tools. It's built with settings that encourage structured, careful, and detailed analysis. If you often write reports, handle compliance paperwork, or pull together information from different sources, these settings help you work faster.

Appendix: Claude Prompt Templates

The following templates are designed to work with Claude's default behaviors. Each template either leverages or overrides specific defaults for optimal results.

Template 1: Multi-Document Synthesis

Use when: Analyzing multiple documents, comparing sources, or synthesizing research across materials.

Why it works: Maximizes Claude's extended context window for cross-document pattern recognition

TEMPLATE:
[Paste Document 1]
[Paste Document 2]
[Paste Document 3]

Analyze these documents together and identify:
1. Points where all sources agree
2. Points where sources contradict each other
3. Information present in one source but missing from others
4. Overall synthesis with recommendations

Template 2: Structured Report Generation

Use when: Converting messy notes, research, or raw data into organized business documents.

Why it works: Aligns with Claude's built-in preference for hierarchical organization

TEMPLATE:
Here are my unorganized notes on [topic]:
[Paste raw notes, research, or data]

Transform this into a structured analysis with:
Executive summary (3-4 sentences)
Key findings organized by theme
Numbered recommendations
Next steps with specific actions

Template 3: Creative Narrative Writing (Overrides Structure Default)

Use when: Writing blog posts, stories, essays, or any content requiring natural flow.

Why it works: Explicitly overrides the structure default that would otherwise impose bullet points and headers.

TEMPLATE:
Write as a single flowing narrative without bullet points, numbered lists, or subheadings.

Use prose paragraphs only.
Vary sentence length for natural rhythm.
Do not organize into sections.

Topic: [your topic]
Audience: [your audience]
Tone: [conversational/formal/playful/etc.]
Length: [word count]

Template 4: Direct Analysis Without Hedging (Overrides Safety Default)

Use when: You need decisive conclusions rather than cautious qualifications.

Why it works: Explicitly requests directness to counteract Claude's cautious reasoning.

TEMPLATE:
I need a direct analysis, not a balanced overview.
Give me your actual assessment, not hedged possibilities.
State conclusions confidently where evidence supports them.
Skip phrases like "it depends" or "there are many factors."

[Your question requiring decisive analysis]
Context: [relevant background]
What I will use this for: [decision you need to make]

Template 5: Unconstrained Brainstorming (Overrides Both Structure and Safety Defaults)

Use when: Generating creative ideas, exploring unconventional approaches, or early-stage ideation

Why it works: Explicitly overrides the structure + safety defaults that push toward conservative, organized output

TEMPLATE:

I want wild, unconventional ideas. Do not filter for practicality.
Do not organize these into neat categories.
Include ideas that seem impractical or risky.
I will evaluate feasibility later—right now I want raw creative possibilities.

Challenge: [describe what you are trying to solve or create]
Constraints to ignore for now: [list any usual limitations]
Direction to push: [any particular angle you want explored]

Template Usage Notes
Long Context Requires Document Placement First: Template 1 places all documents before the analysis instructions. Claude processes large contexts more effectively when the material comes first, and the task comes after.

Structure Default Is Strong: Templates 3 and 5 include multiple explicit override instructions because Claude's structure preference is deeply embedded. A single instruction, like 'no bullets,' may not be sufficient. Reinforce the override with several related instructions.

Overriding Safety Requires Context: Template 4 works because it explains why you need directness and what you will use the analysis for. Providing this context helps Claude understand that the override is appropriate rather than careless.

Brainstorming Needs Double Override: Template 5 overrides both structure and safety because creative ideation suffers when either default is in place. The structure defaults to organizing ideas prematurely. The safety default filters out unconventional possibilities. Address both.

Chapter 15

Understanding ChatGPT's System Prompt

ChatGPT's Three Core Default Behaviors

1. The "Ship, Don't Ask" Default

ChatGPT's system prompt contains an instruction that fundamentally shapes every interaction: If the task is complex/hard/heavy, or if you are running out of time or tokens or things are getting long, and the task is within your safety policies, do not ask a clarifying question or ask for confirmation.

Instead, make a best effort to respond to the user with everything you have so far. Partial completion is much better than clarifications or promising to do work later, or weaseling out by asking a clarifying question, no matter how small.

Recognizing that ChatGPT will always push forward helps you feel more confident in structuring requests effectively, empowering you to take control of your interactions.

The practical implication: ChatGPT will not tell you what information it needs. You must anticipate gaps and provide context proactively. A financial analyst who assumes ChatGPT will ask about time horizons, risk tolerance, or asset constraints before providing investment analysis will receive generic output that requires substantial rework. The same analyst who front-loads all relevant parameters receives actionable analysis immediately.

2. The Verbosity Scale Default

ChatGPT uses a scale of 1 to 10 for verbosity. The default is 3. At 1, the model provides only what is needed to answer the question, using brief language and omitting extraneous detail. An oververbosity of 10 means the model should provide maximally detailed, thorough responses with context, explanations, and possibly multiple examples.

The default of 3 produces responses that are moderately detailed but not exhaustive. This approach is effective for simple questions. For more complex analysis, research synthesis, or planning, it often isn't enough. Knowing how the scale works lets you decide exactly how much depth you want.

For example, saying "Give me a verbosity 8 response" or "Be maximally thorough" will get you a more complete answer. If you want something shorter, say "Be brief" or "Just the key points." This way, you don't have to guess how to phrase your request to get the level of detail you want.

The sophistication-matching instruction compounds this effect: You must always match the sophistication of the writing to the sophistication of the query or request. The level of thoughtfulness you put into your prompt directly determines the level of answer you receive. A casual question receives a casual answer. A precise, detailed question receives a precise, detailed response.

3. The Adaptive Personality Default

ChatGPT's system prompt includes: "Over the course of the conversation, you adapt to the user's tone and preferences." Aim to match the user's communication style. If you do this, the conversation will feel more natural. The model adjusts its style to reflect yours.

Personality presets such as Default, Friendly, Efficient, Professional, Candid, and Quirky are also available. Options like Cynical or Nerdy are available in settings. These presets work alongside your custom instructions.

If you do not coordinate them, they can conflict. For example, if you bring over old custom instructions but do not pick a matching personality preset, you may notice the outputs are inconsistent. The personality setting and your instructions do not align.

There are three layers to the system's operation. Personality sets the tone. Custom Instructions tell the model how to act. Context gives background information. If these layers do not match, the results can be unpredictable. For example, if you choose the Efficient

personality but your instructions ask for detailed explanations, the model will not know which to follow.

Implementation Framework: Working With or Against Defaults

Step 1: Configure Your Three Layers

Layer 1 - Personality Preset. Start by choosing a base personality that aligns with how you prefer to work. If you want to keep things moving with as little back-and-forth as possible, "Efficient" is the option. For business settings, "Professional" is available. If you need the model to adjust to different situations, "Default" will handle that.

Layer 2 - Custom Instructions. Here, you can spell out exactly how you want ChatGPT to respond.

Layer 3 - Context. Add a couple of sentences about yourself. Your role, the industry you work in, and your main use cases. Give the model the information it needs to match the right level of technical detail and language.

If you notice the model giving unpredictable results, it is often because the layers are working against each other. Before assuming the model is at fault, check your setup for any instructions that might contradict each other.

Step 2: Front-Load Context Instead of Expecting Clarification

ChatGPT will not prompt you for missing details. To get useful results, supply all relevant information up front. Include the outcome you want, any requirements or limits, background context, and your preferred response format.

Example: Instead of "Help me write a marketing email," provide "Write a marketing email for our Q1 product launch. The target audience is enterprise IT managers. The product is a security compliance tool priced at $50,000 annually. Tone should be professional but not stiff. Keep it under 200 words. Include a clear call to action for scheduling a demo."

A vague prompt leads to generic output and extra revision cycles. A clear, detailed prompt usually gets you what you need the first time. Over a full workday, this saves considerable time.

Practice Exercise: Calibrating Your Configuration

After a week of using ChatGPT with your current settings, it can be useful to step back and see how well your custom instructions are working for you.

Take a look at my current custom instructions. Based on our recent conversations, are there places where what I say I want doesn't quite line up with how I actually use ChatGPT? Point out any gaps or contradictions you notice, and suggest practical ways I could adjust my instructions to get better results.

ChatGPT will analyze the disconnect between what you said you wanted and what you actually respond positively to. Revealing configuration problems you may not have noticed.

Measurable outcome: Users who complete this calibration process report a 30-40% improvement in first-response satisfaction rates. At scale, this translates to 2-3 hours saved weekly by eliminating unnecessary revision cycles.

The Memory System: Persistence Across Conversations

When This Does Not Work: ChatGPT's Limitations

ChatGPT remembers things in two ways. First, it saves information you ask it to remember. Second, it learns from your previous chats. It responds to phrases like

- remember that
- store this

- add to memory
- note that
- forget that
- delete this

It also recognizes time-related phrases such as

- from now on
- in the future
- going forward

ChatGPT's memory works best for general preferences and important details. If you want it to repeat a specific message, use Custom Instructions or share the content directly in your chat.

Pro subscribers get longer-term memory, so ChatGPT remembers their preferences over time. Free users have a simpler version that only keeps track of recent chats. So Pro users' preferences last for months. Free users need to remind ChatGPT of their needs more often.

By default, conversations with free ChatGPT can be used to train OpenAI's models. If you don't want your chats included, you need to change your privacy settings to opt out.

ChatGPT's defaults create predictable blind spots.

The "ship, don't ask" default means ChatGPT will make assumptions rather than surface gaps in your request. For tasks

requiring specific parameters you forgot to provide, the model produces work based on guesses rather than flagging what it does not know.

The verbatim quotation restriction limits citation accuracy: You may not quote more than 25 words verbatim from any single non-lyrical source. If ChatGPT seems to quote a source, it is usually paraphrasing rather than providing a direct quote. If you need exact citations for academic or professional work, you will need to verify the sources yourself.

In longer conversations, earlier details may drop out. If you are working on a project that spans multiple sessions, it helps to summarize your progress in a new conversation rather than relying on the original thread.

ChatGPT can still generate information that sounds correct but is actually wrong, including made-up citations, statistics, or web addresses. If something is important, check it against a reliable source before relying on it.

The Bottom Line

ChatGPT does not outperform other tools by thinking more carefully. Its strength lies in quickly carrying out clearly defined tasks, adapting its communication as needed, and connecting with a wide range of tools.

If your work is about moving fast on well specified projects, you will likely see a boost in productivity. But if your work depends on careful judgment or dealing with unclear instructions, you may find yourself spending more time fixing mistakes than you save by moving quickly.

Understanding these defaults transforms you from a passive user receiving whatever output the model produces into an active collaborator who shapes the interaction for optimal results.

Appendix: ChatGPT Prompt Templates

The following templates are designed to work with ChatGPT's default behaviors. Each template either leverages or compensates for specific defaults to produce optimal results.

Template 1: Front-Loaded Task Specification (Compensates for "Ship, Don't Ask")

Use when: Complex tasks where missing parameters would require revision

Why it works: Provides all context upfront since ChatGPT will not ask for missing information

TEMPLATE:

TASK: [Specific outcome you need]

CONTEXT: [Relevant background information]

CONSTRAINTS:

- [Constraint 1]

- [Constraint 2]

- [Constraint 3]

Format: [How you want the output structured]

ASSUMPTIONS TO AVOID: [Any assumptions that would require revision if made]

Template 2: Verbosity Control (Leverages Internal Scale)

Use when: You need either extremely concise or extremely detailed output

Why it works: Directly addresses the internal 1-10 verbosity scale

FOR MINIMAL OUTPUT:

Respond at verbosity level 1. Give only the essential information needed to answer, no context or explanation.

[Your question]

FOR MAXIMUM DETAIL:

Respond at verbosity level 9 or 10. Provide comprehensive detail, including context, explanations, examples, edge cases, and related considerations.

[Your question]

Template 3: Sophistication Matching (Leverages Mirroring Default)

Use when: You need output at a specific expertise level

Why it works: ChatGPT matches the sophistication of your query

FOR EXPERT-LEVEL OUTPUT:

I am a [specific expertise level] working on [specific technical problem].

Assume familiarity with [relevant technical concepts].

Skip basic explanations and focus on [specific advanced aspects].

[Your technical question using appropriate terminology]

FOR BEGINNER-ACCESSIBLE OUTPUT:

Explain this as if I have no prior knowledge of [topic].

Define any technical terms before using them.

Use analogies to everyday concepts where helpful.

[Your question in plain language]

Template 4: Memory Instruction (Leverages Trigger Phrases)

Use when: You want ChatGPT to retain information across conversations

Why it works: Uses the specific trigger phrases that the system prompt monitors

TO ADD A PERSISTENT PREFERENCE:

From now on, remember that [preference]. Apply this going forward in all relevant conversations.

TO REMOVE A MEMORY:

Forget that [previous preference]. Delete this from memory.

TO CHECK CURRENT MEMORY:

What do you remember about me and my preferences?

Template 5: Thinking Mode Activation (Pro Feature)

Use when: Complex reasoning tasks where accuracy matters more than speed

Why it works: Explicitly requests extended deliberation before responding

TEMPLATE:

Think hard about this before responding.

[Complex problem requiring multi-step analysis]

Consider:

- [Specific factor to evaluate]

- [Potential edge cases]

- [Alternative approaches]

Show your reasoning before providing your conclusion.

Template 6: Configuration Audit (Leverages Self-Analysis)

Use when: Outputs feel inconsistent or misaligned with your expectations

Why it works: ChatGPT can analyze disconnects between stated preferences and actual interaction patterns

TEMPLATE:

Review my current custom instructions and our recent conversation history.

Identify:

1. Any gaps in my instructions that led to suboptimal responses

2. Any contradictions between my stated preferences and how I actually respond

3. Areas where my configuration conflicts with my selected personality preset

Suggest specific improvements with exact language I could add or modify.

Template Usage Notes

Front-Loading Is Non-Negotiable: Template 1 exists because ChatGPT's core behavior is to execute rather than clarify. Every complex task should use some version of this structure to minimize revision cycles.

Verbosity Numbers Work Literally: Template 2 directly references the internal scale. Saying "verbosity 8" produces noticeably different output than the default 3. Test different levels to find what works for your use cases.

Sophistication Matching Is Bidirectional: Template 3 works because the model mirrors the complexity of your query. A casual question receives a casual answer. A precisely specified technical question receives a technically precise response.

Memory Has Limits: Template 4 uses trigger phrases that the system monitors, but memory is designed for preferences, not large content blocks. Do not expect perfect recall of complex instructions stored this way.

Thinking Mode Requires Pro: Template 5 activates extended reasoning, but full Thinking mode is a Pro feature. Plus, users receive limited access. Free users do not have access to deliberative reasoning modes.

Regular Audits Improve Results: Template 5 should be run weekly during your first month of active use, then monthly thereafter. Configuration drift is common as your usage patterns evolve.

Chapter 16

Understanding Gemini's System Prompt

Gemini's Three Core Default Behaviors

1. Google Ecosystem Integration and Real-Time Data Access

Gemini's built to check information against live Google Search results before answering, fostering trust in its relevance for Google Workspace users and AI practitioners who need current data.

You can reclaim a significant amount of your own time by letting Gemini work across all of Google Workspace, including Gmail, Docs, Drive, Sheets, Meet, and Slides. Information from an email draft can carry over to a spreadsheet or document, so you don't have to repeat yourself or re-enter details. The system keeps track of context for you.

Critical limitation: Without Google Search enabled, Gemini relies solely on its training data, which may be outdated. You generally want to run Gemini with search enabled.

Because Gemini's built to query external sources rather than rely primarily on internal knowledge, knowing how this default works lets you get better results.

2. Direct and Efficient Output Pattern

Google's documentation states explicitly: By default, Gemini is less verbose and prefers providing direct, efficient answers, helping tech-savvy professionals feel confident in its speed and clarity for quick responses. It's best for fact-finding, combining data, and quick answers.

For many business applications, this is an advantage. When you need a quick answer rather than an explanation, Gemini delivers. Ask for a date conversion, and you get the date, not a paragraph explaining calendar systems. Ask for a formula, and you receive the formula, not a tutorial.

This limitation shows up when you want more detailed or engaging responses. According to Google's prompting guide, to make Gemini sound more conversational or friendly, include specific instructions like "explain this" or "talkative assistant." To tailor responses to your preferred style.

Gemini's designed to handle multiple types of information: text, images, video, audio, and code.

For example, you can upload a 90-minute meeting recording and the related slides, then ask Gemini to extract action items, summarize decisions, and spot any differences between what was discussed and what the slides claim. What used to take a human analyst several hours can now be done in a few minutes.

Gemini's system Prompt includes specific context management instructions that differ from other AI tools. Google's guidance is explicit: When working with large datasets, place your specific instructions or questions at the end of the Prompt, after the data context. Anchor the model's reasoning to the provided data by starting your question with a phrase like Based on the information above…

This is backwards from how most people naturally prompt. The intuition is to state what you want first, then provide the data. Gemini performs better when you reverse this order. The difference in output quality between instructions-first and instructions-last prompting is measurable and significant for large context tasks.

The Hidden Cost: Silent Downgrading

Free-tier users often run into an issue that is not well explained: the system can quietly switch to a less advanced model. Because of this, the free tier is not reliable for professional work. You cannot predict when your responses will shift from the better model to a simpler

one, so it is not a good choice for tasks that require consistent quality.

The Mechanism Behind Default Behavior

Each AI company's defaults reflect its business model and the data most readily available to it. Google built Gemini around real-time search because Google is fundamentally a search company. Their competitive advantage is indexing the world's information and making it retrievable. Gemini's system prompts this advantage into every interaction.

So Gemini disproportionately weights information verifiable through Google's index. For topics where Google Search excels, such as current events, product information, or factual queries, this produces excellent results. If you are searching for information on topics that Google does not cover well, the results may be less accurate or even misleading.

Understanding this can help you decide when to rely on Gemini. It works best when the answer is already available online. If you need an answer that requires combining information or thinking beyond what Google can find, you may not get the result you want.

Implementation Framework: Working With or Against Defaults

Step 1: Identify Which Default Applies

Before you start, decide which of Gemini's three main defaults is most useful for your task.

If you are looking up current events or market data, use the default real-time integration.

When you need to analyze several documents or media files, select the default multimodal context (lps).

If you are writing long-form content with personality, you should override the efficiency default.

For historical or niche topics, real-time integration might add unnecessary information rather than help.

Step 2: Write Prompts That Leverage or Override

Override: Suppose you ask Gemini to draft an introduction for your company newsletter. The result is accurate but lacks warmth. This is a result of the system's default setting, which favors information density rather than tone.
To adjust this, specify. Write this as a friendly assistant would. Use a conversational tone and add a detail that helps readers feel included.

Leverage: You are analyzing three years of quarterly reports, earnings call transcripts, and press releases, totaling 400 pages. You put your question first, then paste the documents.

To make use of the system's context handling, reverse the order. Paste the documents first. Then ask: Based on the information above, identify the three most significant changes in strategic direction and provide supporting evidence from multiple sources.

Practice Exercise: Mapping Defaults to Your Work

Start by identifying three information tasks you regularly handle. These could be anything from summarizing articles to organizing research notes.

For each one, ask: Does this task benefit from real-time verification? Does it require warmth and personality, or is efficiency preferred?

Pick one of your tasks. Now, write two different prompts for it.

Version A: Write the prompt as you naturally would.
Version B: Restructure the prompt to either leverage or override Gemini's defaults based on your analysis.

Try both prompts in Gemini. Compare the results: Which one gives you a better answer? Which takes less time to edit? Show yourself where Gemini's default approach works for you, and where giving more specific instructions changes the outcome.

Gemini's Limitations

Gemini tends to miss certain things because of its setup. Gemini's designed for efficiency, but it struggles with longer writing that requires a clear voice and personality. If you give both Claude and Gemini a sample of your writing, Claude does a better job matching your style.

Gemini's real time search can sometimes slow things down and pull in sources that are not useful. Gemini struggles with citing reliable sources, fact checking is mandatory for critical applications.

The multimodal context window, while massive, requires specific prompting conventions that differ from natural language patterns. Users who do not follow the end-format instructions often see degraded performance on document analysis tasks.

Privacy considerations differ as well. By default, Gemini sits inside your Google account. Some conversations are routinely sent to human reviewers and can be retained for up to three years, even if you delete your history. Users handling sensitive data should adjust their Google Account settings under Data and Privacy before relying on Gemini for confidential work.

The Bottom Line

Gemini's advantage is not that it reasons better than other tools. Its advantage is deep integration with Google's search index and productivity applications, combined with native multimodal processing and efficient output defaults. If you work with up to date data, use Google Workspace, or need to handle different types of media at scale, Gemini can save you time. Knowing how its default settings work lets you use them to your advantage or change them when needed.

Appendix: Gemini Prompt Templates

The following templates are designed to work with Gemini's default behaviors. Each template either leverages or overrides specific defaults for optimal results.

Template 1: Large Document Analysis (Leverages Context Default)

Use when: Analyzing reports, contracts, transcripts, or multi-document research
Why it works: Places instructions at the end per Gemini's context management requirements

TEMPLATE:
[Paste all documents here first]

Based on the information above:

1. Identify [specific element you're looking for]
2. Compare [element A] across [documents/time periods]
3. Flag any inconsistencies between [source 1] and [source 2]
Cite specific passages to support each finding.

Template 2: Real-Time Research (Leverages Search Integration)

Use when: Researching current events, market data, competitor analysis, or recent developments

Why it works: Explicitly activates Google Search grounding for verification.

TEMPLATE:
Research the current state of [topic/company/market].
Search for recent developments from the past [timeframe: week/month/quarter].

Include:

- Key facts with dates
- Source URLs for verification
- Any conflicting information from different sources

Prioritize primary sources over aggregator sites.

Template 3: Warm Long-Form Content (Overrides Efficiency Default)

Use when: Writing newsletters, blog posts, marketing content, or any communication requiring personality
Why it works: Explicitly overrides the efficiency default with tone instructions

TEMPLATE:
Write as a friendly, talkative assistant who genuinely cares about the reader.
Use a warm, conversational tone throughout.
Include personal touches that make readers feel welcomed.
Vary sentence length for natural rhythm.
Do not prioritize brevity over engagement.

Topic: [your topic]
Audience: [your audience]
Key points to cover: [list 3-5 points]
Desired length: [word count]

Template 4: Multimodal Meeting Analysis (Leverages All Three Defaults)

Use when: Processing recordings, transcripts, and presentation materials together
Why it works: Combines context management, efficiency, and multimodal processing

TEMPLATE:
[Upload: meeting recording or transcript]

[Upload: presentation slides]
[Upload: any supporting documents]

Based on the materials above, provide:
1. Executive summary (3-4 sentences)
2. Action items with assigned owners (if mentioned)
3. Decisions made
4. Inconsistencies between the discussion and promised slides
5. Questions raised but not resolved

Template 5: Competitive Analysis with Current Data

Use when: Researching competitors, market positioning, or industry trends
Why it works: Combines search grounding with structured output for actionable intelligence

TEMPLATE:
Research [Company/Product] and their competitors in [market/industry].
Search for information from the past [timeframe].

For each competitor, identify:

- Recent product launches or announcements
- Pricing changes
- Strategic partnerships
- Customer sentiment from reviews

Present as a comparison table with source URLs.

stevenlmiller.substack.com

Highlight gaps or opportunities for [your company/product].

Closing

Will AI Take My Job?

This last chapter is intended to reinforce the mindset you've hopefully developed over the course of this book. Leveraging AI as a tool to enhance your capabilities instead of viewing it as a threat.

The key is understanding which skills are becoming more valuable and which skills are rapidly commoditizing and not disappearing, becoming table stakes rather than differentiators.

AI capabilities are constantly improving for fundamental data analysis, report writing, and research synthesis. If your value proposition centers on processing information and presenting it clearly, you're competing directly with AI capabilities that improve monthly. A marketing analyst who primarily creates performance reports competes with AI that generates the same reports in seconds.

Pure coding, basic design work, and routine technical tasks are rapidly being automated. The value is shifting from implementation to architecture, from execution to strategy, from doing to deciding. A developer who writes standard CRUD applications competes with AI that generates the same code. A developer who architects complex systems and makes technology strategy decisions becomes more valuable.

Knowing where to find information or who to contact matters less when AI can instantly access and synthesize vast knowledge bases. The value shifts from having information to understanding what data matters and how to apply it.

These skills are no longer sufficient to ensure career security. They're necessary but not differentiating.

Your AI-Amplified Career Strategy

The path forward isn't about defending against AI. It's about using AI to become irreplaceable. Here's your systematic framework:

Start With Amplification, Not Defense

Stop asking, what can I do that AI can't? Start asking, How can AI make me 10x more valuable?" This mindset shift changes everything.

Identify the highest-value activities in your role. Can AI help you serve more clients? Make better decisions? Create higher-quality work? Move faster than competitors?

For a consultant, the most valuable use of your time is delivering insights that actually influence client decisions. AI can help by speeding up research, suggesting alternative ways to view a problem, and modeling possible outcomes. This means you can focus more on the work that matters, rather than getting bogged down in routine tasks.

If you work in sales, your time is best spent building relationships and advancing deals. AI can draft proposals, gather background information, and organize follow-ups. The result is more time for the parts of the job that actually drive results.

Focus on Intersection Points

The highest value exists at the intersection of AI capabilities and human judgment. Position yourself at these intersection points:

- Where data analysis meets strategic decision making
- Where automated processes meet customer relationships
- Where AI generated options meet creative selection
- Where technical capabilities meet business understanding

AI can process large amounts of data quickly, but human judgment is needed to maximize the value of the output.

Maintain Learning Velocity

The pace of AI advancement means specific tools and techniques constantly evolve. The meta-skill is learning velocity. Your ability to understand and integrate new AI capabilities as they quickly emerge.

When new AI tools emerge, how quickly can you understand their capabilities? How fast can you integrate them into your workflows? How effectively can you teach others to use them? How you answer these questions will decide if you keep your competitive edge or start to lose it.

To do this well, make learning a regular habit. Set aside 3 to 5 hours each week to try out new AI features. Test different tools, try new ways to use them, and keep notes on what you find.

Your 30-Day AI Amplification Plan

We're in a unique historical moment. AI capabilities are advancing rapidly, but most organizations and professionals are still figuring out how to use them effectively. A window of opportunity is open for those who move quickly and deliberately.

Professionals who master AI augmentation now will have a significant advantage. They'll establish themselves as go to people for AI enhanced work. They'll build systems and processes that become organizational standards. They'll develop expertise that compounds over time.

This window won't stay open indefinitely. As AI tools become more sophisticated and accessible, the bar for differentiation will rise. The time to act is now, while early mastery still provides a significant competitive advantage.

The Bottom Line

AI will change your job whether you want it to or not. The only choice is whether you direct that change or react to it.

The professionals who succeed won't be those who defend against AI. They will be those who use AI to become so valuable that replacing them becomes unthinkable. They'll combine AI capabilities with human judgment, relationship skills, creative synthesis, and complex problem solving in ways that create more value than either AI or humans alone.

The practical question is whether you will take the lead in shaping that change for yourself. The opportunity is available to you right now. The tools cost $60/month.

stevenlmiller.substack.com

Did You Like Make AI Work For You?

Before you go I'd like to thank you for reading. If you found the book helpful. Please consider leaving a review.

I know writing from scratch can be hard so please use the template below to make leaving a review easier.

Template

I picked up this book because [what problem were you trying to solve?]. What surprised me was [something you didn't expect].

The most useful part for me was [specific chapter, framework, or tip]. I've already used it to [specific application in your work].

I'd recommend this to [type of professional] who wants to [specific goal].

Example

I picked this up because I kept getting generic AI outputs that needed heavy editing. What surprised me was the section on choosing the right tool for the right task. I didn't realize how different ChatGPT and Claude are in practice. I've started using the Four Pillars framework for every client report and cut my editing time in half. I'd recommend this to any analyst or consultant who uses AI daily and wants to stop rewriting everything it generates.

The Pompatus Times

Insights on Surviving and Thriving in a Volatile World

Subscribe to my newsletter for book updates, more practical AI prompting strategies, and other interesting content I come across worth sharing.