

The Complete ChatGPT User Guide for GPT-5.6

How to use Sol, deeper reasoning, ChatGPT Work, and outcome-driven prompting without overengineering every request

[Mark Chen](#)

OpenAI released GPT-5.6 on July 9, 2026, across ChatGPT, Codex, and the OpenAI API.

But GPT-5.6 is not merely another model upgrade that produces somewhat better answers.

It represents a more consequential change in how we should use ChatGPT.

Earlier generations rewarded users who carefully engineered individual prompts. GPT-5.6 increasingly rewards users who can define outcomes, provide useful context, supervise longer workflows, and exercise judgment at the right moments.

The central question is no longer:

How do I write the perfect prompt?

It is becoming:

How do I give an intelligent system enough direction, context, and authority to accomplish meaningful work — without surrendering control?

This guide explains how to do exactly that.

1. What GPT-5.6 Actually Changes

GPT-5.6 is a family of three capability tiers:

- **Sol** is the flagship model for the most demanding reasoning and professional work.

- **Terra** balances capability, speed, and cost.
- **Luna** is the fastest and most economical tier.


In standard ChatGPT conversations, eligible users primarily interact with **GPT-5.6 Sol** through the Medium, High, Extra High, and Pro reasoning options. Terra and Luna are available through products such as ChatGPT Work, Codex, and the API rather than as normal selectable chat models.

OpenAI describes GPT-5.6 as particularly strong in:

- Professional knowledge work
- Multi-step reasoning
- Software engineering
- Scientific research
- Computer use
- Long-context analysis
- Tool coordination
- Creating polished materials from reference files and templates

The important phrase here is **multi-step work**.

GPT-5.6 is increasingly capable of holding a goal across many intermediate actions: gathering information, analyzing evidence, revising a plan, using tools, producing artifacts, checking results, and adapting when something goes wrong.





That makes it less like an answer generator and more like a supervised digital collaborator.

Part I: Choosing the Right Mode

2. Instant, Medium, High, Extra High, or Pro?

ChatGPT's model picker is now best understood as an **effort selector**, not simply a list of different personalities.

Instant

Instant continues to use GPT-5.5 Instant rather than GPT-5.6 Sol.

Use it for:

- Everyday questions
- Quick rewrites
- Simple explanations
- Brainstorming
- Short summaries
- Casual conversation
- Routine formatting

It is often wasteful to summon the strongest reasoning model merely to

It is often wasteful to summon the strongest reasoning model merely to improve a sentence or answer a straightforward factual question.

Medium

Medium uses GPT-5.6 Sol with standard reasoning.

This should become the default for most serious work:

- Business analysis
- Writing and editing
- Research synthesis
- Spreadsheet logic
- Coding help
- Planning
- Comparing options
- Reviewing documents

For many users, Medium will provide the best balance between quality and speed.

High

High gives GPT-5.6 more reasoning time.

Use it when the task contains:

- Several interacting constraints
- Ambiguous or conflicting evidence
- Important financial or strategic consequences
- Difficult debugging
- Complex document analysis
- Multiple possible interpretations
- A need for deeper critique rather than a quick answer

Extra High

Extra High is appropriate when errors would be costly or when the problem requires sustained reasoning across many steps.

Examples include:

- Auditing a financial model
- Designing a system architecture
- Diagnosing a difficult data discrepancy
- Evaluating a major business decision
- Analyzing a large collection of source documents
- Developing a detailed implementation plan

Extra High is available only on certain plans.

Pro

GPT-5.6 Sol Pro is intended for difficult, long-running workflows where result quality matters more than response speed.

Use Pro when you would otherwise assign the work to a skilled analyst, researcher, engineer, strategist, or consultant and expect them to spend meaningful time on it.

OpenAI currently makes Medium and High available to Plus users, while broader access to Extra High and Pro depends on the user's plan or workspace.

A practical rule

Choose the lowest reasoning level that can reliably complete the task.

Do not think of deeper reasoning as automatically "better." Sometimes it merely makes a simple task slower. Reserve additional intelligence for problems that genuinely contain uncertainty, complexity, or risk.

Part II: The New Prompting Philosophy

3. Stop Writing Specifications for Every Thought

Older prompt-engineering advice often encouraged users to write long prompts containing:

- A detailed persona
- A long list of rules
- A rigid workflow
- Exact output instructions

- Multiple examples
- Repeated warnings
- Elaborate formatting requirements

Those techniques can still be useful, especially in reusable workflows. But GPT-5.6 requires less scaffolding for ordinary tasks.

Instead of programming every intellectual movement, you can often state:

1. What you want accomplished
2. Why it matters
3. What context should be considered
4. What constraints must be respected
5. What a successful result should look like

For example, instead of writing a 700-word research prompt, try:

Analyze the attached customer interviews and identify the three most important reasons users abandon onboarding. Separate strong evidence from hypotheses, quote representative feedback, and recommend the two product changes most likely to improve activation.

This works because it defines:

- The objective
- The evidence source
- The analytical standard
- The required output
- The decision the analysis must support

That is usually more valuable than assigning the model a theatrical identity such as “You are the world’s greatest product consultant.”

4. Prompt for Outcomes, Not Just Outputs

An output is a thing:

- A report
- A table
- A presentation
- A piece of code

An outcome is what the thing must accomplish:

- Help leadership decide whether to launch
- Reveal why costs increased
- Enable a new employee to perform a process
- Persuade customers to adopt a product
- Reduce manual reconciliation work
- Make a complex idea understandable

Compare these two requests.

Output-based request

Create a presentation about our customer survey.

Outcome-based request

Create an executive presentation that helps leadership decide which two customer-experience problems should be funded next quarter. Use the attached survey, distinguish prevalence from severity, and show where the evidence is inconclusive.

The second prompt gives GPT-5.6 a decision context. That context improves what it selects, emphasizes, excludes, and questions.

A strong GPT-5.6 request often follows this compact structure:

Goal: *What needs to be accomplished?*
Context: *What should ChatGPT know?*
Evidence: *What sources should it use?*
Constraints: *What boundaries matter?*
Success: *How will we recognize a good result?*

You do not need to label these sections every time. They are a thinking framework, not a mandatory template.

5. Give It the Real Materials

GPT-5.6 becomes much more valuable when it works from your actual context rather than from a generic description.

Useful materials may include:

- Spreadsheets
- PDFs
- Meeting notes
- Previous reports
- Brand guidelines
- Writing samples
- Screenshots
- Data exports
- Code repositories
- Project documentation
- Email threads
- Presentation templates

OpenAI specifically emphasizes GPT-5.6's ability to reason through multi-step tasks and create materials that follow supplied templates and reference files.

Instead of saying:

Create a monthly financial report.

Provide last month's report, the current data, and the intended audience:

Use the attached June workbook and May management report to create the June report. Preserve the existing definitions and visual structure, but rewrite the commentary based on the new results. Flag any metric that cannot be reconciled rather than inventing an explanation.

The phrase "rather than inventing an explanation" is particularly valuable. It defines the model's behavior under uncertainty.

Part III: Working With GPT-5.6 as a Collaborator

6. Let It Plan — but Inspect the Plan

For complex work, ask GPT-5.6 to determine the workflow before producing the final answer.

A useful request is:

Before completing the analysis, identify the major steps, the evidence you will need, and the points where assumptions may affect the conclusion. Then proceed unless you encounter a decision that genuinely requires my input.

This gives the model room to reason while avoiding unnecessary interruptions.

You can also intervene while GPT-5.6 is working. If its direction seems wrong, correct it early:

Prioritize customer retention rather than acquisition.

Do not redesign the entire process. Look for the smallest change that solves the problem.

Treat the sales forecast as management's target, not as an objective baseline.

This kind of steering is often more effective than trying to anticipate every possible issue in the initial prompt.

The emerging skill is not perfect prompting. It is **active direction**.

7. Separate Creation From Critique

One of the most reliable GPT-5.6 workflows is to divide a task into distinct cognitive passes.

Pass 1: Build

Produce the strongest initial recommendation using the available evidence.

Pass 2: Challenge

Now act as a skeptical reviewer. Identify unsupported assumptions, missing evidence, implementation risks, and ways this recommendation could fail.

Pass 3: Rebuild

Revise the recommendation to address the strongest criticisms.

Preserve useful complexity rather than forcing certainty.

This is better than merely asking:

Make it better.

The build–challenge–rebuild pattern works for:

- Articles
- Business strategies
- Investment theses
- Code
- Project plans
- Presentations
- Policy documents
- Product requirements
- Creative concepts

It reduces the risk that the model becomes attached to its first interpretation.

8. Ask for Uncertainty, Not Artificial Confidence

GPT-5.6 is highly capable, but fluency can still disguise uncertainty.

For analytical work, include instructions such as:

Separate verified facts, reasonable inferences, and speculation.

State which conclusions would change if the underlying assumptions are wrong.

Identify missing information that materially limits confidence.

Do not force a recommendation if the evidence does not support one.

Give me the strongest argument against your own conclusion.

A useful decision format is:

Element — Meaning

Evidence — What the available information directly supports

Assumptions — What must be accepted but has not been proven

Interpretation — The model's explanation of the evidence

Interpretation — The model's explanation of the evidence

Risks — What could invalidate the conclusion

Decision — The recommended action

Trigger — What new information should cause reconsideration

This transforms ChatGPT from a confidence machine into a more disciplined reasoning partner.

Part IV: High-Value GPT-5.6 Workflows

9. Research: Build an Evidence Map Before Writing

Do not begin serious research by immediately asking for an article.

Use this sequence:

Step 1: Define the question

Turn this topic into a precise research question. Identify ambiguous terms and competing interpretations.

Step 2: Gather evidence

Research the question using current, credible sources. Prefer primary evidence. Record the date, source, claim, and limitations of each important finding.

Step 3: Map disagreement

Show where reputable sources agree, disagree, or use different definitions.

Step 4: Form the thesis

Based on the evidence, propose three possible theses: conservative, balanced, and provocative.

Step 5: Write

Write the article using the balanced thesis, while preserving the strongest objection.

This workflow is far more reliable than asking ChatGPT to “research and write” in one undifferentiated step.

10. Business Analysis: Move From Data to Decision

A good business-analysis prompt should establish four layers:

1. What changed?
2. Why did it change?
3. Does it matter?
4. What should we do?

Example:

Analyze the attached monthly results. First reconcile totals and identify possible data-quality issues. Then explain the largest variances by business driver rather than merely restating the numbers. Separate recurring drivers from one-time effects. Finish with three management actions, their expected impact, and the evidence supporting each one.

For a deeper review:

Challenge the analysis as a CFO who suspects that favorable results may be caused by timing shifts rather than genuine operational improvement.

This is where GPT-5.6's stronger professional reasoning becomes especially useful.

OpenAI reports gains across professional evaluations, long-context tasks, computer use, coding, science, and multimodal analysis, although benchmark results vary by domain and should not be interpreted as universal superiority on every individual task.

11. Writing: Direct the Intellectual Movement

GPT-5.6 is strongest when you provide more than a topic and tone.

Tell it how the reader's understanding should change.

For example:

Write a Medium article arguing that AI literacy is shifting from prompt engineering to workflow judgment. Begin with the old belief that better prompts produce better results. Show why stronger models weaken that assumption. Introduce outcome definition, context management, verification, and supervision as the new skills. End with a practical framework readers can use immediately.

This defines the article's **intellectual arc**.

For a stronger personal voice, provide samples and say:

Preserve the rhythm, philosophical curiosity, and sentence-length variation in these samples. Do not copy distinctive phrases. Avoid generic motivational language and inflated claims.

After receiving the draft, run separate editing passes:

- Structural edit
- Evidence and factuality review
- Voice edit
- Concision pass
- Headline and opening optimization

Trying to perform all five at once usually produces a blander result.

12. Coding: Give It an Environment and a Definition of Done

Avoid vague requests such as:

Write a Python program to process files.

Instead provide operational context:

Build a Python 3.12 command-line tool that reads all CSV files from an input directory, validates that required columns exist, combines valid rows, writes rejected rows to a separate file with an error reason, and produces a summary log. Use pathlib, type hints, specific exceptions, and pytest tests. It must work on Windows. Do not silently discard malformed records.

For an existing codebase:

for an existing codebase.

Inspect the repository before proposing changes. Identify the relevant architecture and tests. Implement the smallest safe fix, run the available checks, and summarize what changed, what was verified, and what remains uncertain.

GPT-5.6 shows significant strength in long-horizon coding and tool-driven workflows. OpenAI also reports that the model can complete some production application workflows with fewer steps and fewer tool calls than its predecessor, though real-world performance will depend heavily on the environment, tools, and quality of the task definition.

13. Files, Images, and Multimodal Analysis

GPT-5.6 can work across text, documents, screenshots, charts, diagrams, and other visual materials.

Useful requests include:

Examine this dashboard screenshot and identify what conclusions a manager might incorrectly draw from it.

Compare the figures in this PDF with the values in the spreadsheet and list every material inconsistency.

Explain this system diagram, then identify single points of failure and missing data flows.

Analyze these three design references and derive the shared visual grammar without simply listing visible objects.

For visual analysis, ask the model to distinguish between:

- What is directly visible
- What is inferred
- What cannot be determined from the image

This prevents plausible assumptions from being presented as observations.

Part V: ChatGPT Work and the Rise of Outcome-

Level Delegation

14. What Is ChatGPT Work?

ChatGPT Work is designed to gather context, plan an approach, and act across connected tools, files, and desktop applications.

It can help create and manipulate finished work products such as:

- Spreadsheets
- Documents
- Presentations
- Analyses
- Campaign materials
- Dashboards
- Project outputs

OpenAI describes it as a system capable of navigating ambiguity, adapting as work unfolds, and delivering polished outputs with less prompting. It can also carry context through a larger workflow — for example, turning customer research into a campaign brief, creating assets from that brief, and adapting those assets for different markets.

This changes the appropriate level of instruction.

In ordinary chat, you might request:

Analyze these interviews.

In ChatGPT Work, you can potentially request:

Use the customer interviews, product analytics, and existing campaign materials to identify the strongest adoption barrier. Create a campaign brief, draft the primary assets, adapt them for our three target regions, and prepare an executive summary of the strategy. Ask before publishing or sending anything externally.

Notice the final sentence.

As AI systems gain more ability to act, permission boundaries become part of good prompting.

15. Establish Approval Boundaries

For action-oriented workflows, specify what ChatGPT may do independently and what requires approval.

For example:

You may search, analyze, draft, calculate, organize files, and create working materials without asking. Ask for approval before sending messages, deleting or replacing files, publishing content, making purchases, or changing production systems.

This creates an authority ladder:

Level 1: Observe

Read, search, inspect, and summarize.

Level 2: Recommend

Propose actions but do not execute them.

Level 3: Prepare

Create drafts, plans, files, and changes for review.

Level 4: Execute reversible actions

Perform approved actions that can be easily undone.

Level 5: Execute consequential actions

Send, publish, delete, purchase, deploy, or modify critical systems.

The more capable the model becomes, the more valuable this distinction becomes.

Part VI: What Not to Do

16. Do Not Confuse Longer Prompts With Better Prompts

A long prompt may still be unclear.

Common problems include:

- Repeating the same instruction in different words
- Assigning too many conflicting roles
- Specifying formatting before defining the objective
- Including rules that do not affect the result
- Demanding certainty where uncertainty is unavoidable
- Combining unrelated tasks into one request
- Using excessive emotional pressure such as "This must be perfect"

Clarity beats length.

A compact, high-quality prompt can be:

Review this proposal from the perspective of the executive who must approve its budget. Identify the three questions it fails to answer, the assumptions most likely to be challenged, and the minimum revisions required before approval.

17. Do Not Delegate Judgment You Have Not Defined

"Choose the best option" is incomplete unless "best" has meaning.

Best could mean:

- Lowest cost
- Highest expected return
- Fastest implementation
- Lowest organizational resistance
- Best customer experience
- Least technical risk
- Greatest strategic flexibility

Tell GPT-5.6 how competing values should be balanced:

Recommend the option that can be implemented within three months and reduces operational risk, even if it does not maximize long-term upside.

That one sentence may improve the answer more than an entire page of generic prompting instructions

generic prompting instructions.

18. Do Not Treat the First Answer as the Final Product

GPT-5.6 is powerful enough that its first answer may feel complete.

That creates a new danger: **premature acceptance**.

For important work, always perform at least one of these:

Check your answer against the source material.

Identify the weakest part of your conclusion.

Show me what you may have overlooked.

Recalculate the result independently.

Review the output as an adversarial critic.

List every factual statement that requires verification.

The stronger the model becomes, the easier it is to mistake coherence for correctness.

Part VII: A Practical GPT-5.6 Prompt Framework

For complex work, use the following structure.

The OASIS Framework

O — Outcome

What real-world result should this work produce?

Help leadership decide whether to continue funding the project.

A — Assets

What information, tools, files, or examples should be used?

Use the attached budget, project plan, risk register, and steering-committee notes.

S — Standards

What defines quality?

Reconcile every financial figure, distinguish facts from assumptions, and support recommendations with evidence.

I — Interaction

How should ChatGPT communicate and when should it involve you?

Proceed independently, but ask before making a material assumption that changes the recommendation.

S — Safeguards

What should it not do?

Do not invent missing values, overwrite source files, or communicate externally.

A complete request might be:

Help leadership decide whether to continue funding this project. Use the attached budget, project plan, risk register, and steering-committee notes. Reconcile the financial figures, identify the strongest evidence for and against continuation, and compare three options: continue, restructure, or stop. Proceed independently, but ask before making any assumption that materially changes the recommendation. Do not invent missing data. Deliver a two-page executive memo and a supporting decision table.

This is not “prompt engineering” in the old sense.

It is closer to writing a good assignment brief for a capable colleague.

Part VIII: The Skills That Matter Now

GPT-5.6 reduces the value of some traditional prompt tricks, but it increases the value of deeper human skills.

The most important are:

1. Problem framing

Can you distinguish the apparent question from the real decision?

2. Context selection

Can you provide the information that matters without burying the model in noise?

3. Standards of evidence

Can you define what must be verified and what level of uncertainty is acceptable?

4. Workflow design

Can you divide complex work into research, creation, criticism, revision, and execution?

5. Taste

Can you recognize when an output is technically correct but strategically weak, aesthetically generic, or emotionally false?

6. Supervision

Can you intervene at the right moments without micromanaging every step?

7. Accountability

Can you retain responsibility for decisions rather than hiding behind an AI recommendation?

These skills are harder than memorizing prompt formulas.

They are also much more durable.

The Deeper Shift

The history of ChatGPT can be understood as a gradual movement through four stages:

Stage 1: Ask

The user asks a question. The AI answers.

Stage 2: Prompt

The user carefully specifies the desired output.

Stage 3: Collaborate

The user and AI iteratively develop an idea or solution.

Stage 4: Direct

The user defines an outcome, supplies context and authority, supervises the workflow, and evaluates the result.

GPT-5.6 pushes us decisively toward the fourth stage.

This does not make human judgment less important.

It makes judgment the central scarce resource.

When an AI can produce reports, analyze data, write software, operate tools, create presentations, and coordinate multi-step work, the human advantage is no longer merely the ability to produce each artifact manually.

The advantage lies in knowing:

- What deserves to be built
- Which evidence should be trusted
- Which trade-offs are acceptable
- When the system is wrong
- When to intervene
- What must remain under human authority

The best GPT-5.6 users will not necessarily be those who write the longest prompts.

They will be those who can frame meaningful problems, establish intelligent boundaries, provide the right context, and recognize genuinely excellent work.

That is the real GPT-5.6 user guide:

Define the destination clearly.

Give the system access to the right map.

Let it navigate.

Watch the difficult turns.

And never confuse delegation with the surrender of judgment.

Five Prompts to Try Today

1. Deep analysis

Analyze this problem at three levels: the immediate issue, the system that keeps producing it, and the incentives that prevent it from being solved. Recommend one intervention at each level.

2. Decision support

Help me make this decision. First identify the criteria I appear to be using, including any that conflict. Then compare the options, expose the hidden trade-offs, and recommend what I should choose if my highest priority is long-term flexibility.

3. Research

Research this question using current primary sources. Build an evidence table before drawing conclusions. Separate consensus, disagreement, and unresolved uncertainty.

4. Writing

Turn these notes into a clear argument. Preserve the complexity of the idea, remove repetition, challenge weak assumptions, and write for an intelligent reader who is interested but skeptical.

5. Workflow delegation

Take this project from source materials to a reviewable final deliverable. Plan the steps, identify missing inputs, complete everything that does not require my judgment, and pause only at decisions that would materially change the result.

GPT-5.6 information and availability described in this article are based on OpenAI's official materials published in July 2026. Product access, usage limits, and feature availability may vary by plan, workspace, region, and rollout status.