

BROWARD COUNTY MEDICAL ASSOCIATION · BUSINESS OF
MEDICINE CONFERENCE

AI, Prompting & Context Engineering

A Practical Companion Handbook

30 ready-to-use prompts for personal and clinical/administrative work, five prompting frameworks, a full guide to context engineering, and practical setup walkthroughs for both major agentic tools — Claude Cowork and ChatGPT Work — for a medical practice.

Dana Wallace, MD

Companion to the conference lecture “AI, Prompting & Context Engineering: Practical Tools for the Exam Room and Everyday Life”

July 2026

Contents

1	The Two Skills, Recapped	3
2	Five Prompting Frameworks, Cheat-Sheet Style	4
3	Prompt Library — Personal & Everyday (10)	5
4	Prompt Library — Clinical (10)	8
5	Prompt Library — Business of Medicine (10)	11
6	Context Engineering: The Deep Dive	14
7	Claude Cowork: Setup, Connectors, Skills & Plugins	18
8	ChatGPT Desktop & ChatGPT Work: Setup & Use	24
9	Video Walkthroughs	26
10	Quick Practical Guardrails	27
11	Sources & Further Reading	28

Every prompt in this handbook is written to be copied, edited with your own specifics in the bracketed [placeholders], and pasted directly into ChatGPT, Claude, Copilot, or similar. None require any special account tier — just a text box.

1 The Two Skills, Recapped

FROM THE LECTURE

Prompt engineering is what you say — the clarity of a single instruction: role, task, format, examples. **Context engineering** is everything the model already has in front of it when you ask: documents you've provided, conversation history, standing instructions, and examples. Both matter. Context is usually the bigger lever, and it's the one almost nobody was taught.

THE ANALOGY

Like a specialist consult: prompting is how clearly you describe the question. Context engineering is whether you handed over the actual chart and imaging — or just a two-line referral note and hoped for the best.

This handbook is organized so you can use it two ways: skim Sections 3–5 for a prompt close to what you need and adapt it, or read Section 6 straight through if you want to get meaningfully better at the context side, which is where most of the unclaimed value is.

2 Five Prompting Frameworks, Cheat-Sheet Style

KEEP THIS PAGE OPEN

Role → Task → Format (RTF)

The fastest on-ramp for any prompt. **Role:** who should the AI be? **Task:** what exactly do you want done? **Format:** how should it come back to you? Use this by default; reach for the others when RTF alone isn't producing what you need.

Few-Shot Prompting

Show one worked example of the output format you want, instead of describing the format in prose. One good example is often worth several sentences of instruction, and it's the fastest way to lock in a consistent structure across multiple uses.

Chain-of-Thought

For anything requiring real reasoning — a comparison, a financial decision, a differential — ask the model to work through the relevant factors step by step *before* giving a final answer. This is measurably more accurate on multi-step problems, and just as importantly, it lets you catch a bad assumption before you act on the answer.

Iteration

Treat the first response as a draft, not a verdict. “That’s 80% right — tighten the second paragraph and drop the third bullet” will get you further, faster, than starting over with a longer prompt. Iteration is how prompting is supposed to work, not a sign you got it wrong the first time.

Context-Grounded Prompting

When accuracy matters — a real study, a real policy, a real patient survey — paste the actual source material into the prompt rather than describing it from memory, and explicitly instruct the model to use *only* what you gave it. This single habit does more for output accuracy than any wording change. Section 6 covers this in depth.

3 Prompt Library — Personal & Everyday

10 READY-TO-USE PROMPTS

1. Trip Planning

ROLE + TASK + FORMAT

"You're an experienced travel planner. I have 6 days in Portugal in October, traveling with my spouse. We prefer food and history over nightlife, and have a moderate budget (~\$250/day excluding flights). Give me a day-by-day itinerary in a table with Morning/Afternoon/Evening columns, include one specific restaurant recommendation per day, and ask me two clarifying questions before you finalize anything."

Why it works: role, explicit constraints, an explicit output format, and an invitation to iterate rather than guess. (This is the example from the lecture.)

2. A Difficult Email

ROLE + TONE + FORMAT

"You are a calm, direct communicator. I need to tell a longtime vendor we're switching to a competitor after 6 years, mainly due to pricing. Draft a short email (under 150 words) that is professional and appreciative but firm, doesn't over-explain or apologize, and leaves the door open for future work. Give me two versions: one warmer, one more formal."

Why it works: a named role sets tone more reliably than asking for "professional"; the word cap and two variants prevent the over-apologetic AI-email pattern.

3. Meal Plan From What's on Hand

FEW-SHOT

"Here's what's in my fridge and pantry: chicken thighs, spinach, canned chickpeas, rice, half an onion, garlic, parmesan, eggs, and plain yogurt. Build me 4 dinners for this week using mostly these ingredients — assume I'll buy up to 3 more items per recipe. Format each recipe like this example: Dish Name — 20 min | Ingredients (bullet list) | 3-step method. Keep each recipe under 100 words."

Why it works: the format example plus a hard constraint (only 3 new items) keeps the output genuinely usable instead of generic.

4. Comparing Two Big Purchases

CHAIN-OF-THOUGHT

"I'm deciding between buying a used car for \$18,000 cash or financing a new one at \$32,000 over 5 years at 6.9% APR. Walk through this step by step: total cost of ownership over 5 years, monthly cash-flow impact, and resale-value risk. Then give me a recommendation with your reasoning shown, and flag anything you're uncertain about."

Why it works: visible reasoning steps mirror how you'd want a financial advisor to show their work, so you can catch a bad assumption before you decide.

5. Learning a Complex Topic

ROLE + AUDIENCE

"You are a patient, encouraging teacher. Explain how compound interest works to someone who is comfortable with basic arithmetic but has never taken a finance class. Use one concrete example with real numbers, avoid formulas until the end, and check my understanding with one question before adding the formula."

Why it works: naming the audience's actual level heads off both common failure modes — too simple, or too jargon-heavy — and the built-in check-in keeps it a conversation.

6. Resume Bullet Rewrite

FEW-SHOT

"Rewrite my resume bullets to lead with impact, following this pattern: [Action verb] + [what you did] + [quantified result]. Example: 'Redesigned onboarding workflow, cutting new-hire ramp time from 6 weeks to 3.' Here are my current bullets: [paste 3–5 bullets]. Keep each to one line."

Why it works: one clear example of the target pattern does more work than paragraphs of instruction — the essence of few-shot prompting.

7. Negotiation Prep

CHAIN-OF-THOUGHT + ROLE

"You are a negotiation coach. I'm asking my employer for a 15% raise after taking on a second team's workload for 8 months without a title change. Think through: their likely counterarguments, my strongest points, and a fallback if they say no. Then draft opening language I could actually say out loud, in a confident but not confrontational tone."

Why it works: reasoning through the other side's position first produces a stronger opening line than asking for the line directly.

8. Gift Ideas With Real Constraints

TASK SPECIFICITY

"Suggest 5 gift ideas for my father-in-law's 65th birthday. He's a retired electrician, loves fishing and woodworking, dislikes clutter, and my budget is \$75–100. For each idea, give a one-line reason it fits him specifically — no generic 'here's a nice item' answers."

Why it works: naming what to avoid is often as useful as naming what you want.

9. A Wedding Toast

ROLE + TONE + FORMAT

"You are a warm, funny best man. Draft a 2-minute wedding toast (about 300 words) for my brother. Include one specific, true story. [describe the story in 2–3 sentences]. Keep it light with one heartfelt closing line — no cheesy clichés like 'as I was thinking about what to say...'"

Why it works: naming the specific cliché to avoid heads off the single most common AI-toast failure pattern.

10. A Relaxed Weekend Itinerary

FORMAT + ITERATION

"Plan a relaxed 2-day weekend in [city] for two adults who want a mix of one big activity and lots of unstructured time. Give me a loose morning/afternoon/evening outline, not a packed schedule, and ask what pace we prefer before finalizing."

Why it works: naming the failure mode to avoid ("not a packed schedule") plus an invitation to iterate before commitment.

4 Prompt Library — Clinical

10 READY-TO-USE PROMPTS

BEFORE YOU USE THESE

Use an approved, BAA-covered tool for anything containing real patient information, and de-identify per the checklist in Section 6 before pasting. Every output below is a draft for your review — not a finished document and not a diagnosis.

1. Patient Education Handout

ROLE + TASK + FORMAT

"You are a physician communicator. Rewrite this note at an 8th-grade reading level, in plain paragraphs (no bullet lists), define any medical term you must keep, and end with a 'Call us right away if...' section. Note: [paste de-identified note]."

Why it works: reading level, structure, and an explicit safety-net section are all named — this is the difference between a form letter and something a patient understands. (From the lecture.)

2. Discharge Instructions

FORMAT + SAFETY NET

"You are a physician communicator. Convert this discharge note into patient-facing instructions at a 6th-grade reading level. Organize under three headings: What Happened, What To Do At Home, and When To Call Us. Avoid bullet lists under 'What Happened' — use short, plain sentences instead. Note: [paste de-identified note]."

Why it works: named headings work like a lightweight template, so the model can't wander, and the reading-level target is explicit rather than implied.

3. SOAP Note From Shorthand

FORMAT + FEW-SHOT

"Convert my shorthand exam notes into a SOAP note. Use this structure exactly. Subjective / Objective / Assessment / Plan, each as its own paragraph. Expand abbreviations the first time they appear. Shorthand: [paste bullet notes]."

Why it works: explicit section structure plus an abbreviation-expansion rule heads off the two most common complaints about AI-drafted notes.

4. Organizing a Differential

CHAIN-OF-THOUGHT

"Help me organize my thinking, not make the call. For a patient with [symptom cluster], list the categories of differential I should consider (not specific diagnoses), the one or two most important distinguishing questions or tests for each category, and what would make each more or less likely. This is a structuring exercise for my own reasoning."

Why it works: explicitly framing the task as organizing your own reasoning — not generating a diagnosis — keeps the output useful without inviting over-reliance.

5. Patient Portal Message Reply

ROLE + FORMAT

"You are me, replying to a patient portal message, in a warm but efficient tone I'd actually use in a 90-second reply. Patient asked: [paste message]. Draft a reply under 100 words that answers their question directly, and flag if anything in their message needs an in-person visit rather than a message reply."

Why it works: the word ceiling stops the model from padding, and the escalation flag builds in a useful safety check.

6. Journal Club Summary

CONTEXT-GROUNDED

"Using the abstract below, summarize this study for a journal club audience in 4 parts: what question they asked, how they answered it, the headline finding, and one limitation worth discussing. Abstract: [paste the full abstract text]."

Why it works: pasting the real abstract — instead of asking the model to recall the study — is the single highest-leverage context-engineering move. See Section 6.

7. Informed Consent in Plain Language

ROLE + FORMAT

"You are a physician explaining a procedure to a nervous, first-time patient. Rewrite this consent language in plain conversational language, keep it under 250 words, and end with 'The three things most people ask about are...' followed by three likely questions and short answers. Consent text: [paste]."

Why it works: anticipating the patient's likely questions inside the output does more for comprehension than a longer explanation would.

8. Specialist Referral Letter

FORMAT

"Draft a referral letter to a rheumatologist for a patient with [clinical summary]. Structure: reason for referral (1 sentence), relevant history (short paragraph), pertinent labs/imaging (bulleted), and the specific question you'd like addressed. Keep the whole letter under one page."

Why it works: naming the exact structure a specialist expects saves you from editing a generic letter into shape.

9. CME Talk Outline

ROLE + TASK + FORMAT

"You are an experienced CME course director. I'm building a 20-minute talk for community physicians on [topic]. Give me a slide-by-slide outline with a one-sentence key message per slide, written as a full sentence (not a label), and flag where a case example or visual would strengthen the point."

Why it works: asking for full-sentence "key message" titles rather than topic labels front-loads the hardest part of building a good talk.

10. Multi-Source Literature Summary

CONTEXT-GROUNDED

"Using only the three abstracts I'm pasting below, compare their findings on [topic] in a short table: study, population, key finding, and how they agree or disagree. Do not include studies you recall from training — only these three. Abstracts: [paste all three]."

Why it works: explicitly restricting the model to the provided sources is a critical guardrail against it confidently blending in outside — and possibly outdated or hallucinated — studies.

5 Prompt Library — Business of Medicine

10 READY-TO-USE PROMPTS

1. Associate Hiring Decision

CHAIN-OF-THOUGHT

"I'm deciding whether to hire an associate physician or a nurse practitioner for overflow volume. Think through this step by step — cost, revenue potential, scope, and time-to- productivity — then give me a recommendation with your reasoning shown. Context: solo allergy/immunology practice, ~25 patients/day capacity, turning away 8–10 new patient requests/week."

Why it works: visible step-by-step reasoning turns a gut call into a defensible one you can explain to a partner or accountant. (From the lecture.)

2. Prior Authorization Appeal

CONTEXT-GROUNDED + FORMAT

"Draft a prior authorization appeal letter to [payer] for [medication/ procedure]. Include: medical necessity specific to this patient's diagnosis, prior treatments tried and why they failed, relevant guideline support, and urgency if applicable. Structure as a formal letter, under one page. Clinical summary: [paste de-identified summary]. Supporting guideline excerpt: [paste excerpt]."

Why it works: supplying the actual guideline excerpt — instead of asking the model to recall payer policy — keeps the medical-necessity argument accurate and citable.

3. Practice Policy Draft

ROLE + FORMAT

"You are a practice operations consultant. Draft a one-page no-show policy for a solo outpatient practice: define a no-show, the fee structure after 1st/2nd/3rd occurrence, and the notice patients receive. Keep the tone firm but patient-friendly. Format as numbered sections."

Why it works: numbered structure makes the draft immediately postable rather than a wall of prose you have to reformat.

4. Meeting Notes → Action Items

FORMAT + SPECIFICITY

"Convert these raw staff meeting notes into a table with columns: Action Item, Owner, Due Date. If the owner or due date wasn't stated, write 'TBD' rather than guessing. Notes: [paste notes]."

Why it works: explicitly permitting "TBD" stops the model from inventing a plausible-sounding owner or date that was never actually discussed.

5. Job Description

ROLE + FORMAT

"Write a job posting for a full-time medical assistant in an allergy/ immunology practice. Include: role summary (2 sentences), 5 core responsibilities, 3 required qualifications, 2 preferred qualifications, and one line on practice culture. Avoid generic corporate language like 'fast-paced environment' or 'wear many hats.'"

Why it works: naming specific clichés to avoid is often more effective than simply asking for "professional tone."

6. Patient Newsletter Blurb

ROLE + TONE + FORMAT

"You are writing a practice newsletter blurb announcing a new nurse practitioner joining the team. Keep it to 75 words, warm and welcoming, and end with how patients can book with her. NP background: [paste bio notes]."

Why it works: a hard word cap keeps a newsletter blurb from turning into a full biography nobody reads.

7. EHR Contract Break-Even

CHAIN-OF-THOUGHT

"I'm evaluating whether to renew our EHR contract at \$2,400/month or switch to a system at \$1,650/month with a \$6,000 one-time migration cost and an estimated 3 weeks of reduced productivity during transition. Walk through the break-even math over 24 months, note the biggest uncertainty in the estimate, and give a recommendation."

Why it works: asking it to name "the biggest uncertainty" surfaces the assumption most worth double-checking before you decide.

8. Denied Claim Appeal

CONTEXT-GROUNDED + FORMAT

"Draft an appeal letter for a denied claim. Denial reason given by payer: [paste denial letter text]. Clinical justification: [paste summary]. Structure: reference the specific denial code/reason, rebut it directly with clinical justification, and close with the specific outcome requested (reprocess and pay). Keep to one page."

Why it works: rebutting the payer's actual stated reason, pasted verbatim, produces a sharper appeal than a generic "please reconsider" letter.

9. Employment Offer Comparison

CHAIN-OF-THOUGHT

“Compare these two employment offers step by step: base compensation, bonus structure, benefits value, on-call burden, and non-compete terms. Then estimate total compensation for each in year one and give a recommendation, noting what would change your recommendation. Offer A: [paste terms]. Offer B: [paste terms].”

Why it works: asking “what would change your recommendation” turns a static answer into something you can actually stress-test.

10. Patient Satisfaction Survey Analysis

CONTEXT-GROUNDED

“Using the survey comments below, group them into themes, note roughly how many comments fall into each theme, and flag the single theme that would have the highest impact on patient experience if addressed first. Do not invent themes not represented in the comments. Comments: [paste raw comments].”

Why it works: “do not invent themes” keeps the summary grounded in what patients actually said instead of generic practice-management advice.

6 Context Engineering: The Deep Dive

THE PART WE DIDN'T HAVE TIME FOR

Context engineering is the deliberate curation of everything a model sees before it answers: your instructions, any documents you provide, the conversation history, and any standing profile or examples you've set up. The guiding principle, borrowed from Anthropic's engineering team: find the *smallest set of high-signal information* that fully answers the question. More context is not automatically better — a tightly chosen excerpt usually outperforms an entire chart or a 40-page guideline dumped in at once.

Types & Sources of Context Artifacts

An “artifact” is any piece of real-world material you hand the model as context. Here's where the most useful ones come from and how to bring them in:

Artifact Type	Typical Sources	How to Bring It In	Best For
Text documents	Guidelines, payer policies, contracts, journal articles, practice manuals	Paste the relevant excerpt, or upload the file directly	Summarizing, drafting against a known standard
Spreadsheets / data tables	Schedule exports, billing reports, de-identified patient-panel stats	Upload the file, or paste as a simple table/CSV text	Trend spotting, financial projections, benchmarking
Email / message threads	Prior correspondence, patient portal messages	Copy-paste the relevant thread, oldest message first	Tone matching, reply drafting, summarizing a back-and-forth
Meeting notes / transcripts	Staff meetings, recorded or transcribed calls	Paste the raw notes or transcript text	Action-item extraction, meeting summaries
Images / screenshots	Charts, scanned forms, handwritten notes, whiteboard photos	Upload the image directly (most current tools read images natively)	Transcription, form interpretation, describing a chart
Structured data exports	De-identified EHR reports, scheduling-system exports	Export as CSV or PDF from the source system, then upload	Quantitative summaries, panel analysis
Prior AI drafts	An earlier response in the same or a different conversation	Paste it back in when starting a new, related thread	Iterative refinement across sessions

Style guides / templates	Your own brand guide, a sample letter, a slide template	Upload once, or fold into standing instructions (see below)	Consistent voice and formatting across many outputs
Web pages / public guidance	Public guideline pages, association statements, news	Paste the key excerpt — most tools cannot reliably browse live	Fact-grounded summaries with a known, checkable source
Audio / video transcripts	Dictated notes, recorded lectures or calls	Transcribe first (dictation software or a transcript tool), then paste the text	Note-taking, summarizing spoken content

Structuring a Context-Rich Prompt

When a prompt combines instructions with a pasted document, separate them clearly. Anthropic's own testing found that clearly delimited prompts — using simple XML-style tags — produced meaningfully more consistent output than the same content run together as one block of text. You don't need special software; typing the tags by hand works exactly the same way:

```
<context>
[Paste your practice's current no-show policy here, in full]
</context>

<task>
Revise the policy to add a telehealth no-show clause, consistent with the existing tone
and structure.
</task>

<output_format>
Return the full revised policy. Mark any new or changed text in [[double brackets]] so
I can see exactly
what changed.
</output_format>
```

Two habits worth keeping: put context *before* the task (it helps the model know what standard to apply), and keep the tag names descriptive — <context>, <meeting_notes>, and <survey_comments> all work fine as long as the name matches what's inside.

Before You Paste Anything Patient-Related

Use a tool covered by a business associate agreement for any real patient information — never a consumer-grade tool without one. If you're building a template with a fictional or de-identified

example first (as in this handbook), the HIPAA Safe Harbor method gives a concrete checklist of what “de-identified” actually means: remove all 18 of the following identifiers, keeping only the year for any date tied to an individual.

- | | |
|--|---|
| 1. Names | 10. Account numbers |
| 2. Geographic subdivisions smaller than a state | 11. Certificate/license numbers |
| 3. All dates related to an individual (keep year only) | 12. Vehicle identifiers & serial numbers |
| 4. Telephone numbers | 13. Device identifiers & serial numbers |
| 5. Fax numbers | 14. Web URLs |
| 6. Email addresses | 15. IP addresses |
| 7. Social Security numbers | 16. Biometric identifiers |
| 8. Medical record numbers | 17. Full-face photographs & comparable images |
| 9. Health plan beneficiary numbers | 18. Any other unique identifying number or code |

Source: U.S. Department of Health & Human Services, HIPAA Privacy Rule de-identification guidance (Safe Harbor method).

Managing Conversation Length ("Context Rot")

Research on long-context models has documented a real phenomenon: as a single conversation grows very long, the model gets measurably worse at recalling and using information from earlier in that same thread — even though it technically still “fits.” Practical takeaway: when you're moving to a genuinely new task, start a new conversation rather than continuing to build onto an old one. Long threads are fine for one evolving task; they are a liability once the topic has actually changed.

Setting Up Standing Instructions

Most current tools (Custom Instructions in ChatGPT, “Instructions for Claude” in Claude, and similar features elsewhere) let you save context once so you stop repeating yourself in every prompt. A good standing-instructions block is specific, covers the things you'd otherwise type every time, and reads like a set of defaults you're comfortable applying to almost anything you ask.

A REAL EXAMPLE — MY OWN ACCOUNT-WIDE INSTRUCTIONS

This is the actual standing-instructions block I use, saved once and applied to every conversation regardless of tool or topic. It's longer than the one-line version most people start with — reproduced here in full as a template you're welcome to adapt line by line rather than copy wholesale.

1. **Professional Context.** Treat most requests as professional work unless I indicate otherwise. Match the level of an experienced academic colleague. Assume subject-matter expertise and do not explain basic concepts unless requested.

2. **Response Quality.** Prioritize accuracy over speed. Verify information that may be current or subject to change. Never fabricate facts, references, quotations, or citations. Clearly distinguish established evidence, expert opinion, inference, and speculation. State when evidence is limited, conflicting, or uncertain.
3. **Communication Style.** Write in clear, concise American English. Lead with the answer, followed by supporting details. Use headings, numbered lists, tables, and bullet points whenever they improve readability. Keep explanations proportional to the question's complexity.
4. **Critical Thinking.** Do more than answer the question. Evaluate whether a better approach exists and recommend it when appropriate. Respectfully challenge assumptions that are unsupported by evidence. Explain important trade-offs, limitations, and alternative interpretations.
5. **Scientific and Medical Work.** Unless instructed otherwise, use AMA references for medical and scientific writing. Cite authoritative primary or official sources whenever practical. When reviewing manuscripts, identify major and minor issues, unsupported statements, inconsistencies, methodological and statistical concerns, ambiguous wording, missing or outdated references, and opportunities to improve clarity — while preserving my writing style and making the fewest changes necessary.
6. **Decision Support.** For comparisons or recommendations, evaluate advantages, disadvantages, risks, costs, evidence quality, implementation issues, and important trade-offs before reaching a conclusion.
7. **Execution.** Do not over-plan straightforward requests. If the task is clear, complete it first and allow refinement afterward. Reserve extensive planning for genuinely complex projects.
8. **Editable Deliverables.** Unless I request a static product, generate deliverables in editable formats — Word for text and tables, Excel for datasets and charts, PowerPoint for presentations and diagrams, SVG for illustrations. Favor editability over appearance unless I request publication-ready artwork, and tell me before proceeding if a fully editable version isn't achievable.
9. **Final Deliverables.** When I request a final deliverable, return only the finished product unless I explicitly request commentary — no introductions, summaries, or notes on what changed. Treat 'final,' 'ready to send,' and similar phrases as instructions to produce only the finished product.
10. **Default Preferences.** Unless instructed otherwise: optimize for quality rather than speed; be candid rather than agreeable; tell me when you believe I am mistaken; use metric units except where customary units are standard; and end factual or research-based responses with a confidence assessment and the most authoritative sources used.

Notice what this buys back: a request as short as “write patient education on seasonal allergy nasal sprays” already comes back accurate, well-sourced, appropriately hedged, and in the right format — because the standing instructions are doing work that would otherwise mean re-typing the same caveats and preferences every single time.

Curate, Don't Dump

It's tempting to paste in everything that might be relevant “just in case.” Resist it. A model's attention is a finite resource, and diluting it with a 40-page guideline when only one section applies tends to produce a worse answer than a tightly chosen excerpt would — the signal gets

lost in the noise. Before pasting a document, ask: what's the smallest excerpt that fully supports this task?


7 Claude Cowork: Setup, Connectors, Skills & Plugins

FROM CHAT TO AGENTIC WORK

Everything in Sections 1–6 applies inside a normal back-and-forth chat. Claude Cowork is a different mode entirely: instead of one exchange at a time, you describe an outcome, and Claude plans and works across multiple steps — reading and writing real files, using apps you've connected, and handing back a finished deliverable. This section covers the mechanics: getting set up, connecting your tools, and where to find the skills and plugins that make Cowork useful for a medical practice specifically.

Getting Started

Cowork is available on paid Claude plans (Pro, Max, Team, Enterprise) on Claude Desktop for macOS and Windows, on the web at claude.ai, and on Claude Mobile. Local file access, browser use, and computer use all require the Claude Desktop app to be installed and open, even though the session itself runs remotely on Anthropic's servers.




Illustrative diagram — not an actual product screenshot. Menu names and locations may change as Cowork evolves.

The working folder. On desktop, Cowork can only read and write files in folders you explicitly authorize. Anthropic's own guidance recommends starting with one dedicated, low-stakes folder rather than granting broad access to your whole drive — see the security discussion in the companion lecture for why. Once you point Cowork at a folder, Claude can pick up *folder instructions* specific to that location (more below).

The Customize Menu: Skills, Connectors & Plugins

Skills, connectors, and plugins all live in the same place — the **Customize** menu — and Anthropic's own unified directory lets you browse all three without switching screens. They solve different problems, and the table below is worth knowing before you go looking for any one of them.



Illustrative diagram, summarizing Claude Help Center articles current as of mid-2026.

Skills — what they are, and where to find them

A skill is a folder of instructions (and sometimes scripts) that Claude loads automatically when it's relevant to what you're doing — think of it as a repeatable procedure Claude can follow consistently, rather than something you re-explain every time. Anthropic skills (enhanced Excel, Word, PowerPoint, and PDF creation) are built in and available to everyone. You can also build **custom skills** of your own, or install **partner skills** from companies like Notion, Figma, and Atlassian.

Examples worth trying in a medical practice:

- A custom skill that formats patient education handouts at a fixed reading level with your clinic's standard disclaimer and footer — so you never have to restate the reading-level instruction again.
- A custom skill that applies your presentation or letterhead style guide automatically, the way this very handbook and its companion deck were built.
- A custom skill that structures meeting or committee notes into your practice's standard format.

To browse: Customize > Skills > the "+" button > Browse skills. Installed skills are enabled by default and Claude uses them automatically when relevant — you don't have to invoke them by name.

Connectors — what they are, and where to find them

A connector lets Claude reach into an app or service you already use — Google Drive, Gmail, Slack, Microsoft 365, DocuSign, and many others — using your own permissions in that service. If you can't open a file in the source app, Claude can't reach it through the connector either. You control, per connector, whether Claude can only read from a service or also write to it.

To connect: Customize > Connectors > the "+" button > find the service > Connect, then follow the authentication prompts. In Cowork specifically, connectors reach external services through Anthropic's cloud rather than your local network, which matters if a service sits behind your practice's firewall.

Plugins — what they are, and where to find them

A plugin bundles skills, connectors, and (in Cowork) sub-agents together into one ready-to-go package for a role or function — installing one plugin can be faster than configuring several skills and connectors separately. Anthropic maintains a growing library organized into marketplaces, including a default **Knowledge Work** marketplace plus optional ones for **Life Sciences**, **Legal**, and **Financial Services** — the Life Sciences marketplace is the most directly relevant starting point for a medical practice, though the exact plugins available change over time, so it's worth browsing directly rather than relying on any fixed list. A "Plugin Create" plugin is also available if you want to bundle your own skills and connectors into a reusable package for your practice.

To install: in Cowork, open the Cowork tab first, then Customize > Plugins > Browse plugins > Install. Plugins installed this way are available in both Chat and Cowork.

Where Your Instructions Live

Claude has three layers of standing instructions, each with a different scope. Knowing which one you're editing matters — a change in the wrong place either won't apply where you expect, or will apply more broadly than you intended.

Instructions for Claude

Account-wide — applies to every conversation, in Chat and Cowork alike.
Click your initials > Settings > "Instructions for Claude"

Cowork Global Instructions

Applies to every Cowork session specifically — layers on top of your account-wide instructions.
Settings > Cowork > Edit next to "Global instructions"

Folder / Project Instructions

Tied to one specific local folder or Cowork project. Claude can update these itself as it learns the folder.
Set automatically when you point Cowork at a folder, or inside a Cowork Project

Illustrative diagram. In ChatGPT, the equivalent account-wide field is called Custom Instructions.

A REAL EXAMPLE

My own account-wide instructions tell Claude to be concise across every conversation I have. The project instructions for the deck-building project behind this handbook separately tell Claude to follow my presentation style guide and remember these are 30–45-minute professional talks. The first travels with me everywhere; the second stays local to that one project.

Try It: Two Cowork Prompts for a Practice

Organize a Shared Drive Folder

COWORK

"Organize my Downloads folder by type and date."

Why it works: This is one of Anthropic's own official example prompts — deliberately low-stakes, and a good first Cowork task to try in Manual approval mode before granting broader access.

Compare a Contract Against Your Standard Template

COWORK

"Compare the employment contract PDF against our standard template PDF. Flag clauses that differ — especially non-compete terms, liability language, and termination notice — in a short table."

Why it works: Same context-engineering principle as Section 6 — Claude can only flag what differs because you gave it a second, real reference document to compare against, not just the contract alone. Treat the output as a starting point for your own review, not a substitute for counsel.


8 ChatGPT Desktop & ChatGPT Work: Setup & Use

THE OTHER MAJOR AGENTIC TOOL

OpenAI released ChatGPT Work alongside GPT-5.6 and a rebuilt desktop app in July 2026 — an agentic mode that plans and completes multi-step projects on its own, closely paralleling Claude Cowork in Section 7. This section covers only what's specific to ChatGPT: everything about prompting well and engineering good context in Sections 1–6 applies here exactly as written, whichever tool you're in.

Getting Started

The new ChatGPT desktop app (macOS and Windows) combines three modes in one place: **Chat** for ordinary conversation, **Work** for research and finished deliverables, and **Codex** for software development. Work is also available on the web and in the mobile app for eligible paid plans.



Illustrative diagram — not an actual product screenshot. OpenAI is still rolling this out; menu names may change.

ONE PRACTICAL DIFFERENCE WORTH KNOWING

Work on the web or mobile runs entirely in the cloud and cannot see files on your computer. Work in the desktop app can also open a local folder and read/write those files, the same way Cowork requires Claude Desktop for any local file access. If the task involves files sitting on your machine, you need the desktop app open, not just a browser tab.

Where It Reaches: Connected Apps & Projects

ChatGPT Work connects to a plugin directory of outside tools — Slack, Microsoft Teams, Google Drive, SharePoint, email, calendars, CRMs, and project trackers among them — the same underlying idea as Claude's connectors in Section 7, under different branding. A **Project** keeps related conversations, files, and instructions together across a piece of ongoing work, and **Scheduled Tasks** let a Work task run once, repeat on a schedule, or watch for a trigger and act when it fires.

Staying in Control

ChatGPT asks permission before using an app on your computer for the first time, and again before any sensitive or disruptive action regardless of your past answers. You can mark an app “Always Allow” for future tasks, and remove it again later from the Computer Use section of desktop settings. OpenAI's own guidance is worth repeating verbatim in spirit: stay present — don't walk away — for anything touching accounts, payments, credentials, or security settings, and if it's driving your browser, review each action as though you were clicking it yourself, since a site can't tell the difference between you and an approved agent action signed into your account.

Try It: Two ChatGPT Work Prompts for a Practice

Build a CME Talk Development Project

CHATGPT WORK

“Set up a project for my next CME talk on [topic]. I'm adding my last three decks on related subjects and our practice's slide style guide. Draft a slide-by-slide outline with a full-sentence key message per slide, and flag where a case example or visual would strengthen the point.”

Why it works: a Project keeps the reference decks and style guide attached across every future working session on this talk, so you're not re-uploading them each time you come back to it — the ChatGPT Work equivalent of Claude's folder instructions in Section 7.

Turn a Denials Folder Into an Appeals Tracker

CHATGPT WORK

“Open the folder of denial letters and supporting guideline excerpts I've pointed you to. Draft an appeal letter for each denial that rebuts the payer's stated reason directly, and build a tracking spreadsheet with columns for patient, payer, denial reason, appeal status, and date sent.”

Why it works: one instruction, two finished deliverables (letters plus a tracker) — this is the kind of multi-step, multi-file work an agentic tool is built for and a plain chat window isn't.

Section 8 reflects OpenAI Help Center documentation and reporting current as of July 2026. ChatGPT Work was still rolling out to eligible accounts at time of writing — availability, menu names, and exact mechanics may change as it matures. When in doubt, check help.openai.com directly.

9 Video Walkthroughs

FOR LEARNING AT YOUR OWN PACE

A short, hand-picked list for going deeper on either tool's agentic mode outside of this handbook. These are third-party videos, not produced by Anthropic, OpenAI, or Dana Wallace, MD — the usual caveat applies: features shown may have changed since filming, and menu names in particular move quickly on both tools right now.

Claude Chat & Cowork

FULL Claude Cowork Tutorial for Beginners in 2026! (Zero to PRO)

youtube.com/watch?v=JdQ_FHgP5ms&t=1942s

Understand the FULL Claude Ecosystem in One Video

youtube.com/watch?v=MS7E5TXNviM

Claude Cowork for Academics: Full Setup and Use Cases

youtube.com/watch?v=ldwjRjx8xlQ

The Simple Claude Cowork Setup That Runs My Life (Steal It)

youtube.com/watch?v=VISWuUosVMI&t=2s

Master Claude Cowork in 28 Minutes

youtube.com/watch?v=qytfIOF4sLk

ChatGPT Work

ChatGPT Work Completely Changes How You Use ChatGPT (Full Guide)

youtube.com/watch?v=tqKGaqaQkNk&t=333s

How I Use ChatGPT Work and GPT-5.6 to Do Everything (Beginner Tutorial)

youtube.com/watch?v=WLg9qWOf6zw&t=703s

10 Quick Practical Guardrails

ONE PAGE, NOT A POLICY MANUAL

KEEP THESE FOUR IN MIND

1. **Draft only.** A human reviews everything before it's sent, filed, or acted on.
2. **Approved tools only.** Use whatever your institution has designated for the task at hand.
3. **De-identify before pasting.** Use the Safe Harbor checklist in Section 6, and use a BAA-covered tool for anything containing real patient information.
4. **Verify anything quotable.** Numbers, citations, and guideline references should be checked against the real source before you rely on them — especially if the model wasn't given that source directly.

11 Sources & Further Reading

WHAT THIS HANDBOOK DRAWS ON

American Medical Association. *2026 Physician Survey on Augmented Intelligence*.

American Medical Association. "More than 80% of physicians use AI professionally: AMA survey."

Anthropic. "Effective Context Engineering for AI Agents." September 2025.

Anthropic / Claude Platform Docs. "Prompting Best Practices."

Chroma Research. "Context Rot: How Increasing Input Tokens Impacts LLM Performance."

Prompting Guide AI. "Chain-of-Thought Prompting" and "Few-Shot Prompting."

U.S. Department of Health & Human Services. HIPAA Privacy Rule De-Identification Guidance (Safe Harbor Method).

Anthropic / Claude Help Center. "Get Started with Claude Cowork."

Anthropic / Claude Help Center. "Use Claude Cowork Safely."

Anthropic / Claude Help Center. "What Are Skills?", "Use Plugins in Claude," and "Use Connectors to Extend Claude's Capabilities."

Anthropic / Claude Help Center. "Browse Skills, Connectors, and Plugins in One Directory" and "Understanding Claude's Personalization Features."

OpenAI Help Center. "ChatGPT Work and Codex."

OpenAI Help Center. "Moving to the New ChatGPT Desktop App."

OpenAI / ChatGPT Learn. "Computer Use," "Permissions," and "Agent Approvals & Security."

HubSpot. "ChatGPT at Work" prompt bundle (free download) — structural inspiration for the Section 8 "Try It" prompts, adapted here for a medical practice: offers.hubspot.com/view/chatgpt-at-work-bundle

All prompt-response examples referenced from the companion lecture were generated live by Claude (Anthropic), July 2026. Section 7 reflects Claude Help Center documentation current as of mid-2026; menu names and locations may change as Cowork continues to develop — when in doubt, check support.claude.com directly. Section 8 reflects OpenAI documentation and reporting current as of July 2026, during ChatGPT Work's initial rollout — check help.openai.com directly for current mechanics.

Questions about this handbook or the conference lecture can be directed to Dana Wallace, MD, through the Broward County Medical Association.