



The AI Drafting Handbook

Practical workflows for teams that write with ChatGPT, Claude and Codex

THE UNPLAIN TEAM

Contents

What to give the model	4
The input pack	5
Writing a brief that works	6
Source material and what is missing	7
The drafting loop	8
Station one: outline	8
Station two: draft	9
Station three: critique	9
Station four: revise	10
When to break the loop	10
Checking facts and sources	11
Keep the claim count honest	11
Extract the claims	12
The claim ledger	12
Citations the model gives you	13
Editing for voice	13
A rubric for voice	14
Editing with the model	14
The last read	15
Publishing the document	15
Write clean Markdown	15
The pre-publish checklist	16

What a publishing tool should handle	16
References	17

How to use this handbook

Most teams that draft with AI assistants have the same complaint. The output is fine: grammatical, structured, confident and entirely forgettable. The model does not write badly. The team has no workflow around it: no brief, no sources, no critique pass, no fact check, no voice edit, no publishing checklist.

This handbook builds that workflow in five chapters. Read the first two in order; they set up the loop the others depend on. The prompts are starting points to copy, adapt and argue with. Claims about how a tool behaves are dated as of 2026 and kept general, because the tools move faster than any handbook. The companies in the examples are fictional.

What to give the model

The most common failure in AI drafting happens before the first word is generated. Someone types a request that would make no sense to a human freelancer either, and the model fills the gaps with the average of everything it has read.

The research agrees. Noy and Zhang found that professionals using ChatGPT wrote faster and were rated higher, mostly because the output served as a first draft. Dell'Acqua and colleagues found that consultants with AI access did worse on a task designed to sit just outside what the model could reliably do, because they trusted output they should have questioned. AI drafting works when a human knows where that frontier is and treats the output as a start, not a verdict.

A human freelancer would ask questions. The model mostly will not unless you tell it to. So give it what a good freelancer would have demanded.

The input pack

Think of input in five buckets. Each prevents a specific, predictable failure.



Give the model four things: a brief, the sources, a voice sample and the constraints.

Ingredient	Description	What it prevents
The brief	Goal, audience, length, format, deadline	Generic output aimed at nobody
Source material	Notes, data, transcripts, prior documents	Invented facts and vague claims
Voice samples	Two or three pieces you consider on-brand	House style drift and machine cadence
Constraint list	Banned words, legal lines, required points	Late rewrites over things you knew
A rejected example	One piece you turned down, with the reason	Your pet peeves coming back

The last row is the one most teams skip. Rejection is specific; taste is hard to articulate.

Keep all five in a folder called the input pack. Novelty belongs in the content, not in the setup.

Writing a brief that works

A brief is not a prompt. It is the document you would hand to any writer, and it should survive being read by someone who was not in the meeting. A working brief has seven parts:

- 1** Purpose. One sentence: what should the reader do, decide or believe afterwards?

- 2** Audience. Who exactly, what they already know, what they are skeptical about.

- 3** Format and length. A 600-word post is a different animal from a 6,000-word guide.

- 4** Key points. The three to five things that must be in the piece, in priority order.

- 5** What to avoid. Claims you cannot support, competitors you do not name, jargon you ban.

- 6** Voice. Three adjectives and a pointer to a sample, not a paragraph of adjectives.

- 7** Done means. Who approves, by when, and what checks happen before that.

If you cannot fill in the first two, you are not ready to draft, with or without a model.

A brief that cannot be handed to a stranger is not a brief. It is a hope.

Source material and what is missing

Models are fluent about things they have no information on. The survey by Ji and colleagues describes hallucination, output that is plausible, well formed and wrong, as a structural property of these systems, not a bug the next version will remove. The defense is to give the model something true to work from: meeting notes, transcripts, the spreadsheet of results, the previous version. Put it under a heading such as Source material (do not edit) so the model draws from it instead of rewriting it.

Then say what is missing. If the draft needs a number you do not have yet, tell the model to write a placeholder in square brackets rather than guess. A draft with [Q3 revenue figure needed] is honest. A draft with a confident invented figure is a liability.

Keep the constraint list short. A model given forty constraints will satisfy thirty, and you will not know which ten it dropped.


When you suspect the brief is thinner than it should be, ask the model before it drafts.

Prompt: Before writing anything, read the brief and source material and ask me the five questions whose answers would most change the draft. Do not draft. Wait for my answers.

Answer the questions in the brief itself, not in the chat, so the next person inherits the fix.

The drafting loop

Once the inputs are ready, the work follows a loop with four stations: outline, draft, critique, revise. What changes with a model is speed, and the temptation to skip stations because the first draft already looks finished. A draft that looks finished at station two is the most dangerous object in this workflow.



Outline, draft, critique, revise: short loops beat one long prompt.

Station one: outline

An outline is cheap to produce, cheap to read and cheap to throw away. A full draft is none of those.

Ask for the structure before the prose.

Prompt: Using the brief and source material above, propose an outline for this document: section headings with one sentence under each saying what the section argues, and the source items each draws on. Flag any section where the source material is thin. Do not write prose yet.

Read it as an editor. Is the order right? Is there a section that exists only because outlines usually have one? Fix it in a sentence and ask again. Two rounds is typical; three means the brief was unclear.

Station two: draft

Ask for prose one section at a time. Drafting in sections keeps the model focused and lets you catch drift before a wrong assumption in section one has colored sections two through eight.

A section prompt looks like this.

Prompt: Write section two of the outline, following the voice samples and constraints. Draw only on the source material provided. Where you need a fact we have not supplied, write a placeholder in square brackets rather than inventing one. Target 450 words. Do not add a conclusion; this is a middle section.

The word only is doing a lot of work; models default to helpful, and helpful means filling gaps. The ban on conclusions exists because tidy summaries in the middle of a document are one of the surest tells of machine prose.

Station three: critique

This is the station most teams skip, and it produces the most improvement per minute spent. Use a fresh session, because a model that has just written something tends to defend it.

Give it the brief and the draft, and ask for problems, not praise.

Prompt: You are a skeptical editor who has not seen this draft before. Read the brief, then the draft. List the ten most important problems in priority order. For each, quote the passage, say what is wrong (unclear, unsupported, off-brief, off-voice, repetitive or misplaced) and suggest a specific fix. Do not compliment the draft. Do not rewrite it.

The critique will be uneven. You are the editor: accept what is right, reject what is wrong, and if it keeps flagging the same kind of problem, add a constraint to the input pack.

Station four: revise

Do not ask the model to rewrite the whole thing. You will lose the parts that were good and gain new problems you have not read yet.

Give it the annotated critique and ask for a targeted revision.

Prompt: Revise the draft to address only the critique points I have marked ACCEPT below. Leave everything else unchanged, including wording you might consider improvable. After the revised draft, list each accepted point and the sentence or paragraph you changed to address it, so I can verify.

The change log turns a black box into a diff you can check in a minute.

When to break the loop

Three signals mean you should stop iterating and fix the brief instead: the critique and the revision keep disagreeing about the same passage; you are rewriting more than a quarter of each section by hand; or the draft is confidently discussing something the source material does not cover. No prompt fixes an input problem.

If you are rewriting most of the draft by hand, the problem is upstream of the draft.

Checking facts and sources

A model can produce a paragraph that is fluent, specific, confidently sourced and false. The citation may not exist. The statistic may be from a different year, a different country or a different universe. So every factual claim that reaches a reader must have been checked by a person against a source that person has opened. Not a source the model named. That is manageable if you keep the number of claims honest and track them in a ledger.



Every claim gets a source, a status and an owner before the document goes out.

Keep the claim count honest

Many drafts carry more factual claims than the argument needs, some of them there only because the model, asked to sound authoritative, reached for a number. For each one, ask whether the argument falls apart without it. If not, cut it or soften it to a statement you can defend. Eight checked claims beat thirty of which twenty are checked.

Every number in your document is a promise. Make fewer promises and keep all of them.

Extract the claims

Before checking anything, get the full list. This is a job a model does well, because it is extraction rather than generation.

Hand it the draft.

Prompt: Read the draft below and list every factual claim: statistics, dates, names, quotations, descriptions of what a product or law does, and statements about what research found. Number them. For each, quote the exact sentence and say whether the draft attributes it to a source. Do not evaluate whether the claims are true.

You will get a longer list than you expected. That is the point.

The claim ledger

The ledger lives next to the draft, and nobody marks the document ready until every row is resolved. Each row records the claim, what it was checked against, and the outcome with a named human: verified, needs a check, softened or removed. The rows below are illustrative.

Claim in the manuscript	Verification source	Resolution and owner
The pilot cut onboarding time by a third	Internal metrics, September report, page 4	Verified by Priya
The feature launched in March 2026	Release notes for version 4.2	Verified by Tom
The rival's pricing starts at \$49	Their pricing page, opened 2 October 2026	Verified by Priya

Claim in the manuscript	Verification source	Resolution and owner
Customers asked for exports, as quoted	Interview transcript from call seven	Needs a check from Lena
Sixty percent of teams use two tools	No source found anywhere	Removed by Tom

Three rules make the ledger work:

- 1** The source column names something a colleague could open. Internal metrics is not a source. Internal metrics, September report, page 4 is.

- 2** The outcome is set only by the owner, after looking. Nobody marks a row verified because the draft sounded confident.

- 3** Removed rows stay in the ledger, so the same invented statistic does not reappear in the next draft.

Citations the model gives you

A model asked to cite will cite. Some citations are real and relevant, some real but irrelevant, and some assembled from plausible parts: a real author, a real journal, a title never published. Find the source yourself and read enough to confirm it says what you claim; if you cannot find it in a few minutes, it does not go in. The model is a decent research librarian and a terrible witness.

Editing for voice

Fact checking makes a document true. Voice editing makes it yours. A model given no instruction defaults to a register readers now recognize on sight: framing sentences, closing summaries, triplets of adjectives, a fondness for the word landscape. Together they produce prose that is smooth and anonymous.

The remedy is not to tell the model to sound human. The remedy is to decide what you sound like, write it down, and edit against that standard every time.

A rubric for voice

A rubric turns taste into a checklist that two editors can apply the same way. Score each row from one to three. A document is ready when every row scores at least two.

Characteristic	Signatures of a weak draft	Signatures of a strong draft
Specificity of nouns	Abstract nouns and no examples	Concrete nouns, one example per point
Rhythm of sentences	Every sentence is medium length	Short and long sentences, on purpose
Confidence in claims	Hedges twice per claim	States it, qualifies once if needed
How sections open	Section starts by announcing itself	Section starts inside the point
How sections end	Every section ends on a summary	Sections end on the last useful idea
Choice of vocabulary	Leverage, landscape, delve, robust	Words your team uses in meetings

Use the rubric twice: to brief the model, pasted into the input pack, and to grade the result. A row that scores one draft after draft needs a constraint in the pack and an example.

A reader can only argue with a specific sentence.
Give them one.

Editing with the model

The model is a good voice editor when given a target and a poor one when asked to improve the writing, because improve means nothing and it will reach for its defaults.

Give it the rubric and a specific list of operations.

Prompt: Edit the passage below for voice using the rubric and voice samples provided. Specifically: break any paragraph that opens with a framing sentence, cut any sentence that summarizes the preceding one, replace the words on the banned list, and vary sentence length. Keep all facts, claims and structure unchanged. Show the result, then list each change in one line.

Then read the result against the rubric yourself. The change list makes that fast, and it tells you what the model thinks is padding, which is itself useful.

The last read

Voice editing ends with a human reading the whole document aloud, or at least under their breath. If a sentence makes you run out of air, split it; if you find yourself skimming, so will your reader. Williams and Bizup make the same point at book length: clarity is won sentence by sentence, not by a tone instruction at the top of a prompt.

Publishing the document

A draft becomes a document the moment someone outside the team can read it. The last mile has three parts: clean Markdown, a checklist, and a publishing tool that does the layout so you do not.

Write clean Markdown

Markdown is the natural format for AI-drafted documents. Models produce it fluently, it reads as plain text, it diffs well, and nearly every publishing tool imports it. Gruber's design goal was that a Markdown document should be publishable as is, without looking marked up. Clean Markdown means a few habits:

- Headings are structure, not decoration. One top-level heading per chapter, second level for sections, never skip a level.
- Lists are for lists. Bullets for parallel items, numbers when order matters.
- Tables are for comparison. Two to four columns, short cells, a header row.
- Blockquotes are for things that are quoted: prompts, pull quotes, excerpts. Not for emphasis.

Left alone, a model also adds horizontal rules between sections, bold used as headings and a summary closing every section. Add them to the banned list.

Then run a mechanical check before the human read.

Prompt: Check this Markdown for publishing problems: skipped heading levels, nested lists deeper than one level, tables with more than four columns or more than eight rows, inline code, HTML, horizontal rules between every section, and a summary paragraph at the end of each section. List each problem with the line it occurs on. Do not fix them; I will.

An agent can run that over a whole folder, and should.

The pre-publish checklist

Run this with the document in front of you in the form a reader will see it, not in the editor. Gawande's argument for checklists applies exactly: they make the obvious steps impossible to skip when you are tired, which is when you skip them.

- The document has a single purpose, stated or obvious in the first paragraph.
- The headings, read alone, make sense as a table of contents.
- Every claim in the ledger is verified, softened or removed, and every placeholder is gone.
- Every banned word is gone.
- Every table has a header row, fits the page and compares something.
- The voice rubric scores two or better on every row.
- A human has read the whole thing, start to finish, in the final format.
- The file name, title, author and date are correct.
- The person who approves has approved this version, in writing.

What a publishing tool should handle

Once the content is right, the remaining work is presentation. A tool that takes Markdown should produce, without manual intervention, a branded cover, a contents page with computed page numbers, pagination that keeps headings with their text and never splits a table row, typography somebody chose on purpose (Butterick's 45 to 75 characters per line), and a closing page with author, contact and date.

Unplain is one such tool, and this handbook was produced that way. Whichever you use, the test is the same: clean Markdown goes in and a document you would hand to a client comes out.

The words are the work. Everything after the words should be automatic.

Keep the Markdown source as the canonical version and treat the PDF as a build output. Never edit the PDF directly; the next regeneration erases your fix.

References

- Butterick, M. Butterick's Practical Typography, 2nd edition. Online, 2010 onward.
- Dell'Acqua, F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayner, L., Candelon, F. and Lakhani, K. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." Harvard Business School Working Paper 24-013, 2023.
- Gawande, A. The Checklist Manifesto: How to Get Things Right. Metropolitan Books, 2009.
- Gruber, J. "Markdown: Syntax." Daring Fireball, 2004.
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y., Madotto, A. and Fung, P. "Survey of Hallucination in Natural Language Generation." ACM Computing Surveys, vol. 55, no. 12, 2023.
- Noy, S. and Zhang, W. "Experimental evidence on the productivity effects of generative artificial intelligence." Science, vol. 381, no. 6654, 2023.
- Williams, J. M. and Bizup, J. Style: Lessons in Clarity and Grace, 12th edition. Pearson, 2016.

The words are yours. The layout is ours.

Paste a draft from ChatGPT, Claude or Codex, or send it through the MCP, and get a designed PDF back.

[Create a PDF free](#)