

THE HUMANISER KIT

Make AI stop sounding like **AI**

Four humanisation prompts, a Claude writing skill, a tone-of-voice setup for Claude, ChatGPT and Perplexity, the 60 tells that give AI writing away, and a five-minute edit routine.

DRAFT, BEFORE EDITING

~~In today's ever-evolving digital landscape, it's not just about content, it's about seamless, robust impact.~~

Five minutes. No humaniser tool.

If you only read **one page**, read this one

It's a setting, not a curse

AI sounds like AI because a thin brief lets it default to the most likely words. Change the inputs.

Suspicion costs you

Readers rate content lower on trust when they only suspect AI wrote it (Raptive, 2025).

Humanisers miss the point

Experienced readers still spot AI text after it's been run through a humaniser (Univ. of Maryland, 2025).

Save it as a skill

A Claude skill loads your voice rules only when needed. Build one or adapt an open-source humaniser.

The prompt does most of the work

Make the model ask you five questions first, and ban invented sources.

Five minutes does the rest

Search the kill list, cut the top and tail, add one real detail, read aloud, check claims.

Less context, better output

Three to five high-signal files beat 40 brand documents that contradict each other.

Feed it your best work

Update your project monthly with edited copy you're proud of. Never raw AI drafts.

The robot voice is a symptom of lazy prompting and no editing. Both are fixable this afternoon.

WHAT CHANGED

Why sounding like AI now costs you

Most marketing teams now draft with AI. Readers have noticed. They don't need proof a machine wrote something. A hunch is enough, and the hunch is getting sharper.

Finding	Figure	Source
Lower trust when readers suspect content is AI-generated	48% lower	Raptive, 3,000 US adults, Aug 2025 (vendor research)
Lower perceived authenticity when AI is suspected	57% lower	Raptive, Aug 2025 (vendor research)
Lower emotional connection when AI is suspected	60% lower	Raptive, Aug 2025 (vendor research)
US adults who say telling AI from human content matters	76%	Pew Research Center, Sept 2025
US adults not confident they could tell	53%	Pew Research Center, Sept 2025
UK adults comfortable with news made automatically by AI	11%	Reuters Institute Digital News Report 2025
Workers who received low-effort AI "workslop" in the past month	40%	BetterUp Labs and Stanford Social Media Lab, Sept 2025
2024 PubMed abstracts processed with an LLM (lower bound)	13.5%+	Kobak et al., Science Advances, 2025
Frequent ChatGPT users' errors spotting AI text (majority vote)	1 in 300	Russell, Karpinska and Iyyer, ACL 2025

What that adds up to

- People care whether a machine wrote it, aren't sure they can tell, so they go on feel.
- The people most likely to judge your copy (heavy AI users) are the best at spotting it.
- Google doesn't penalise AI as such. It targets content made at scale without adding value, however it was made.

Readers forgive the tool. They don't forgive copy that shows nobody was paying attention.

The diagnosis

Why AI drafts all sound the same, and the 60 tells that give them away.

A thin brief

"Write a blog about X" gets you X written the way the average of the internet writes it.

No voice reference

If you don't show it how you sound, it falls back on a safe, formal, faintly breathless default.

No opinion allowed

Models hedge by default. Ask for balance and you'll get 900 words of it.

How to use the 60 tells

One tell proves nothing. Humans write "crucial" too. Look for clusters: four or five in a paragraph is a pattern. Each tell comes with a "try instead" so you fix it rather than swap one stock word for another.

Researchers have measured this. Kobak and colleagues found "delves" appeared about 28 times more often than expected in 2024 PubMed abstracts (Science Advances, 2025). Max Planck researchers found "delve" rising in spoken academic talks after ChatGPT launched (arXiv, July 2025).

Tells 1 to 20: **vocabulary and stock phrases**

Vocabulary Words people rarely say out loud

	The tell	Try instead
1	delve	look at, get into
2	tapestry	mix, range, or cut it
3	testament ("stands as a testament")	shows, proves
4	landscape, ever-evolving	name the actual market or year
5	unlock, unleash	get, start, open
6	leverage, harness	use
7	seamless, robust	name the real quality: fast, reliable, fewer steps
8	pivotal, crucial	important, or show why it matters
9	underscore, showcase	shows
10	navigate the complexities of	deal with, handle

Stock phrases Filler that could open any article

	The tell	Try instead
11	"In today's fast-paced world"	start with the reader's problem
12	"It's important to note that"	delete it and just say the thing
13	"Whether you're a X or a Y"	write for one reader
14	"When it comes to X"	"For X," or start with X
15	"Let's dive in"	delete
16	"In conclusion", "In summary"	end on the action
17	"Game-changer"	say what actually changed
18	"Here's the thing", "Here's the kicker"	delete and make the point
19	"At the end of the day"	delete
20	"Great question!", "I hope this helps"	delete the chat residue

Tells 21 to 40: structure and rhythm and formatting

Structure The same shapes, every time

	The tell	Try instead
21	"It's not X, it's Y"	state Y
22	Everything in threes	use the real number of things
23	A question, then its own answer ("The result?")	state the result
24	An intro that restates the headline	open with a fact, a scene or the problem
25	A conclusion that restates the intro	end on the next step
26	Every section the same length	let each section be as long as it needs
27	Bold label and colon on every bullet	plain bullets, or a paragraph
28	Moreover, furthermore, additionally	"And", "Also", or a new sentence
29	"From startups to enterprises" range openers	name one real example
30	Signing off with "What do you think?"	ask something specific, or stop

Rhythm and formatting Too tidy to be human

	The tell	Try instead
31	Every sentence the same length	mix short and long on purpose
32	Em dashes everywhere	full stop or comma
33	Title Case Headings	sentence case headings
34	Emoji bullets	plain bullets
35	Bold scattered through every paragraph	bold one thing, or nothing
36	Stacks of one-line "drama" paragraphs	combine them
37	US spelling in UK copy (optimize, color)	set UK English and check
38	"X: The Ultimate Guide" headlines	say what is actually inside
39	Lists where a paragraph would do	write the paragraph
40	Paired adjectives ("clear and concise")	pick the one that's true

Tells 41 to 60: substance and tone

Substance Confident, but empty

	The tell	Try instead
41	Hedging ("may potentially help to")	commit, or give the condition
42	Both-sides neutrality	say where you land and why
43	Vague examples ("a business might...")	name a real one
44	"Studies show" with no study	name it with a date, or cut it
45	Quotes with no speaker	attribute it, or cut it
46	Specific-to-sweeping pivots ("this speaks to a broader shift")	stay specific
47	Inflated significance ("plays a vital role")	show the actual impact
48	Nothing only you could know	add one first-hand detail
49	Restating the brief back to the reader	start with the answer
50	Suspiciously neat or invented numbers	source it or remove it

Tone A brochure that read a self-help book

	The tell	Try instead
51	Relentless positivity	admit the trade-off
52	Sycophantic openers	cut them
53	Brochure words (nestled, boasts, vibrant)	describe it plainly
54	Formal register on a casual channel	write how you'd say it
55	A faceless "we" with no opinion	give the team a view
56	Overpromising ("transform your business")	promise what you can show
57	"We understand X can be challenging"	prove you understand by being specific
58	Hype words (revolutionary, cutting-edge)	describe what is new
59	Sounding like every other brand	use your voice sample
60	Disclaimer residue ("As of my last update")	delete it, then check the facts

The full humaniser prompt

Save this as a template. Fill the brackets once per brand, then reuse it. It works in any mainstream AI assistant: paste it at the start of a chat, or into your custom or project instructions.

COPY THIS

ROLE AND READER. You are writing as [brand or person] for [specific reader: role, sector, what they already know, what they're worried about]. They will read this on [channel].

BEFORE YOU WRITE. Ask me up to five questions about anything you would otherwise have to guess. Prioritise: our actual opinion; one real example, client story or mistake we can use; the source for any figure you plan to use. Wait for my answers before drafting.

VOICE. Match the voice of the sample below. Copy its rhythm, sentence length and attitude, not its topic. [Paste 300 to 500 words.] Three words for our voice: [dry, direct, warm]. We never sound: [salesy, academic, breathless].

STRUCTURE.

- The first sentence makes the main point or names the reader's problem.
- One idea per paragraph. One to three sentences each.
- Headings in sentence case, saying what the section answers.
- No closing summary. End with the action the reader should take.

LANGUAGE.

- UK English. No em dashes.
- Never use: [your kill list].
- Never use these patterns: "it's not X, it's Y"; groups of three by habit; a question followed by its own answer; "whether you're X or Y"; opening with "In today's".
- Vary sentence length on purpose. Prefer verbs to abstract nouns.

EVIDENCE. Every statistic needs a named source and a date in the same sentence. If you don't have one, write [SOURCE NEEDED]. Never invent a number, study, quote or name. Mark anything you're unsure of with [CHECK].

AFTER YOU DRAFT. List any banned words or patterns you used and fix them. Quote the three weakest sentences and rewrite them. Tell me which claim a sceptical reader would challenge first.

Two lines do most of the work: "ask me five questions first" and "[SOURCE NEEDED]".

Prompts that make the first one better

Prompt two: calibrate your voice (run once, reuse forever)

Most teams can't describe their own voice. This gets you a usable description in two minutes. Paste the result into the humaniser prompt, your project instructions or your skill's voice.md.

COPY THIS

Here are three pieces of writing we're proud of. Describe our voice in ten rules a new writer could follow: sentence length, vocabulary, humour, how we open, how we close, and what we never do. Back each rule with a short quote from the samples. Don't flatter us. Point out anywhere the samples contradict each other.

[Paste sample one] [Paste sample two] [Paste sample three]

Prompt three: rewrite a draft you already have

For the draft someone else made, or the one you made before reading this.

COPY THIS

Rewrite the draft below in our voice. Keep every fact, figure and source exactly as it is. Cut anything that restates the headline or summarises what was just said. Replace any word from our kill list with the plainest alternative. Where a sentence is vague, mark it [NEEDS DETAIL] rather than inventing one. Then list every change you made and why.

Our voice: [paste your ten rules]. Kill list: [paste it]. Draft: [paste it]

Prompt three keeps your facts intact. Most "make this sound human" prompts quietly change them.

Prompt four, and the rules behind all four

Prompt four: the sceptical editor (in a fresh chat)

Use a new chat so the model isn't marking its own homework.

COPY THIS

Act as a sceptical editor who hates generic copy. Read the draft below and give me a list, not a rewrite:

- Every sentence that could appear unchanged in a competitor's article.
- Every claim without a named source and date.
- Every instance of these patterns: [paste your kill list or the 60 tells].
- Three alternative opening sentences, each more specific than the last.

[Paste draft]

Five principles that make any prompt sound more human

Principle	In practice	Source
Say what to do, not just what to avoid	"Write in short, plain sentences" beats "don't be wordy"	Anthropic prompting best practices
Explain why a rule exists	"No em dashes, because readers now read them as AI" helps the model generalise	Anthropic prompting best practices
Show, don't describe	Three to five real examples beat a paragraph of adjectives	Anthropic prompting best practices
Write the prompt in the style you want back	A prompt full of bullets gets bullets back	Anthropic prompting best practices
Remove contradictions	"Be concise" next to "be thorough" makes output worse	OpenAI GPT-5 prompting guide

Keep the rewrite for yourself. The model is good at spotting problems. You're better at fixing them, because you know what actually happened.

The **five-minute** edit routine

Built for a 1,000-word draft. Set a timer. Stick it next to your screen.

Minute one

Search and destroy

Ctrl+F your kill list. Rewrite every sentence that contains a hit. Don't swap "delve" for "dive into".

☐ Done

Minute two

Cut the top and tail

Delete the first and last paragraph. Reread. Nine times out of ten, nothing is missing.

☐ Done

Minute three

Make one thing specific

Swap one vague example for a real one. Add a detail only your team would know: a client conversation, your own data, a mistake you made.

☐ Done

Minute four

Read it out loud

Cut wherever you stumble or run out of breath. Break up runs of same-length sentences. Kill any triplet that isn't earning its place.

☐ Done

Minute five

Check the claims

Every number has a named source and a date. Every quote has a speaker. Anything you can't verify goes.

☐ Done

Five minutes of editing beats five subscriptions.

Before and after: two rewrites

Illustrative examples. Swap in your own detail, because the point is that only you have it.

A LinkedIn post

BEFORE

In today's fast-paced digital landscape, content is more crucial than ever. It's not just about creating content, it's about creating content that resonates. Here's the thing: AI can help you unlock seamless, scalable, impactful content. Whether you're a startup or an enterprise, the key is to leverage AI strategically. What do you think?

AFTER

Our last batch of AI drafts all opened with a sentence nobody would say out loud. The fix wasn't a better tool. We made the model ask us five questions before it wrote anything, then cut the first paragraph of whatever came back. Two minutes. Far less robot. Try it on your next post.

- Removed: "in today's", "landscape", "crucial", "it's not X, it's Y", "here's the thing", "unlock", "seamless", a triplet, "whether you're", "leverage", the engagement question.
- Added: a first-hand detail, a concrete method, a specific next step.

A charity service page

BEFORE

At [Charity], we understand that navigating the complexities of dementia care can be challenging. Our robust, person-centred services play a pivotal role in empowering families, fostering connection and unlocking brighter futures.

AFTER

Caring for someone with dementia often means a lot of phone calls to people who can't help. Call us first. We'll tell you what support you're entitled to, who to ask, and what to say when you do.

- Removed: the empathy cliché, "navigating the complexities", "robust", "pivotal", the triplet of abstract nouns.
- Added: the reader's actual experience, one clear action, a promise the team can keep.

Score any draft **out of 20**

Score each check 0 (no), 1 (partly) or 2 (yes). Takes two minutes. Use it on agency work too. We won't mind.

Check	Score
1 The first two sentences make the point or name the reader's problem	0 1 2
2 No kill-list words (2), one or two (1), three or more (0)	0 1 2
3 No structural tics: "not X, it's Y", habit triplets, self-answered questions	0 1 2
4 Sentence length varies. Read aloud without stumbling	0 1 2
5 At least one detail only our team could know	0 1 2
6 Every number has a named source and date	0 1 2
7 It takes a clear position	0 1 2
8 It's written for one reader, not everyone	0 1 2
9 It ends on an action, not a summary	0 1 2
10 Read next to our voice sample, it sounds like the same people	0 1 2
Total	/ 20

17 to 20

Publish. Maybe have a biscuit.

12 to 16

Run the five-minute routine once more.

Under 12

Go back to the humaniser prompt and redraft.

Your team's **one-page** style sheet

Fill it in once, share it with everyone who writes or prompts, and paste the relevant lines into every humaniser prompt.

Who we write for

Role, sector, what they already know

Three words for our voice

We never sound

e.g. salesy, academic, breathless

Our kill list

Start with the 60 tells, add your own

Words and phrases we do like

Spelling and style

UK English? Em dashes? Sentence case headings? How we write numbers

Voice samples

Links to three pieces we're proud of

Sources rule

e.g. named source and date for every figure

Who signs off

One named person per channel

Your setup

Stop pasting the same prompt every day. Save your voice once, in the right place, and keep it lean.

Skills

Your voice rules, saved as a Claude skill that loads itself when the job needs it.

Projects

One lean workspace per brand in Claude, ChatGPT or Perplexity.

Context

Three to five high-signal files, refreshed monthly. Nothing else.

"It took us a long time to get here. We tried dozens of prompts, built and rebuilt our own writing skills, and kept feeding our projects tone-of-voice notes and memory until the output finally sounded like us. There was no single magic prompt. It was a system, and we had to build it one frustrating draft at a time."

Ash, Agency Director, Seventh Element

Build a **writing skill** in Claude Code

A skill is a folder with a SKILL.md file inside. Until it's triggered, only its name and description sit in Claude's context (Anthropic Agent Skills documentation), so it's far lighter than pasting a brand guide into every chat.

Step one

Create the folder

Personal: `~/claude/skills/our-writing-voice/`. Shared with your team through a repository: `.claude/skills/our-writing-voice/`

Step two

Write the description first

It decides when the skill switches on. Say what it does and when to use it. Limit 1,024 characters.

Step three

Keep SKILL.md short

Under 500 lines (Claude Code docs). Long material goes in separate files: `voice.md`, `kill-list.md`, `examples/`.

Step four

Test, then fix the skill

Every time a draft is wrong, change the skill, not just the draft. In the Claude app, zip the folder and upload it under Settings, Capabilities, Skills (Pro, Max, Team and Enterprise, with code execution on).

COPY THIS INTO SKILL.MD

name: our-writing-voice

description: Writes and edits copy in [brand]'s voice. Use for any blog post, LinkedIn post, email, landing page or rewrite of an AI draft, and whenever someone asks to make writing sound less like AI.

Our writing voice

1. Read `voice.md` and `kill-list.md` before writing anything.
2. Ask up to five questions about anything you would otherwise guess: our opinion, a real example, a source.
3. Draft in the voice described in `voice.md`. Use the files in `examples/` as the benchmark for rhythm and tone.
4. Check the draft against `kill-list.md`. Rewrite every sentence that contains a hit.
5. Every statistic needs a named source and date, or write [SOURCE NEEDED].
6. Finish by listing the three weakest sentences and rewriting them.

Third-party skills **worth trying**

Open-source skills we'd look at first. Star counts from GitHub, October 2026. They change constantly.

Skill	What it does	Watch out for
humanizer github.com/blader/humanizer	Checks 26 AI writing patterns based on Wikipedia's "Signs of AI writing" guide, then rewrites, self-critiques and finalises. About 57,000 stars, MIT licence	Generic. It knows AI's voice, not yours. Its own README says beating detectors isn't the goal
stop-slop github.com/hardikpandya/stop-slop	Banned phrases, structural clichés and sentence rules, with a score out of 50 across five dimensions. About 18,000 stars, MIT licence	Some rules are stricter than your brand may want. Edit them
the-elements-of-style github.com/obra/the-elements-of-style	Strunk's 1918 <i>The Elements of Style</i> as a "write clearly and concisely" skill	About 12,000 tokens of reference text. American and old-fashioned in places
anthropics/skills github.com/anthropics/skills	Anthropic's official examples: doc-coauthoring, internal-comms, brand-guidelines and skill-creator	brand-guidelines applies Anthropic's own brand. Use skill-creator to build yours
awesome-claude-skills github.com/travisvn/awesome-claude-skills	A curated directory of community skills	A list, not a quality check

Before you install anything

- **Read the SKILL.md first.** Skills are instructions, and sometimes code, that your AI will follow.
- **Skip anything selling "undetectable".** That's detector evasion, and it doesn't fix what readers notice.
- **Adapt, don't adopt.** Add your own kill list and examples, or you'll swap one generic voice for another.

Our setup: a third-party humaniser for the first pass, our own voice skill for the second, a human for the third. The third pass is not optional.

Set up your voice in every tool

The names differ. The principle doesn't: a short, current, high-signal picture of how you write.

Claude

- ☐ **Project:** ten voice rules in the instructions; voice guide, kill list and three to five examples as knowledge
- ☐ **Styles:** Create custom style, then Add writing example, for chats outside a project
- ☐ **Skills:** your voice skill plus one humaniser

On paid plans, projects switch to retrieval as knowledge nears the context limit (Anthropic says up to 10x more content). That isn't a reason to fill them.

ChatGPT

- ☐ **Custom instructions** (Settings, Personalization): top ten rules and kill list only. 1,500 characters on Free and Go, 5,000 on paid plans
- ☐ **Personality:** pick the nearest base tone, such as Professional, Candid or Efficient
- ☐ **Projects:** instructions plus voice files. Project instructions override your global ones

Don't start new custom GPTs. OpenAI plans to retire them in favour of plugins, from December 2026 for affected Enterprise workspaces.

Perplexity

- ☐ **Projects** (some will know them as Spaces): voice rules in the instructions, up to 8,000 characters
- ☐ **Profile** (settings): your communication style and formatting preferences

We use Perplexity for research and sourcing, not final drafts. Set the voice so summaries come back readable, then write elsewhere.

Your voice pack (the same files everywhere)

File	Length	Refresh
Voice guide: ten rules with examples	One to two pages	Quarterly, dated
Kill list: words and patterns, each with a "try instead"	One page	Whenever a new tic appears
Examples: your best recent, edited, approved work	Three to five pieces	Monthly
Audience note: who you write for and what they already know	Half a page	Quarterly

Context: why more isn't better

"The mistake we made early on was thinking more context meant better output. We stuffed our projects with every brand document we had. It burned through tokens, the documents contradicted each other, and the copy got blander, not better. Now we keep it lean, update it every month, and feed it the copy we're proudest of. Your project should learn from your best recent work, not your oldest."

Ash, Agency Director, Seventh Element

- **Context is a budget.** Anthropic calls it "a critical but finite resource" with diminishing returns (Effective context engineering for AI agents, Sept 2025).
- **Longer inputs, shakier answers.** Chroma tested 18 models and found performance grew increasingly unreliable as input length grew (Context Rot, July 2025).
- **The middle gets lost.** Models use information at the start or end of long inputs best (Liu et al., TACL 2024; 2023 study).

Do	Don't
Keep three to five high-signal files	Upload every deck, guideline and blog archive "just in case"
Put your most important rules first	Bury them in the middle of a 20-page document
Refresh monthly and date your voice guide	Set it up once and wonder why it writes like last year
Feed it edited, approved copy you're proud of	Feed it raw AI drafts. It learns its own habits back
Delete anything that contradicts today's voice	Keep 2019 and 2026 guidelines side by side
Pair every ban with what to do instead	Write 40 "nevers" and nothing it should do
Keep one project per brand or client	Mix three clients' voices in one project
Check uploads are allowed under your data policy	Paste in confidential client data without checking

What this means for **budget holders**

The cost of doing nothing

The BetterUp Labs and Stanford Social Media Lab research (Sept 2025, 1,150 US desk workers) is the internal version of this problem. Each workslop incident took recipients about two hours to deal with, an estimated \$186 per affected employee per month. Recipients rated the sender as less capable (50%) and less trustworthy (42%). Swap "colleague" for "prospect" and you can see the external cost.

Governance that actually works

- **One named editor per channel.** "We all check it" means nobody does.
- **A shared kill list and sourcing rule,** baked into every prompt template, project and skill.
- **One shared voice setup,** not five people with five different custom instructions. Someone owns the monthly refresh.
- **In regulated sectors,** AI drafts never reach compliance or clinical review unedited. Reviewers should spend their time on substance, not on removing "pivotal".
- **Decide your disclosure stance** before someone asks you on a client call.

Risks to manage

Risk	What the evidence says	What to do
Search visibility	Google targets content made at scale without adding value, by automation, people or both (Google Search Central)	Fewer, better pages. Add first-hand experience
Accuracy	Google calls accuracy review critical because models can make things up	Minute five, every time
Policing staff with detectors	OpenAI withdrew its own classifier in 2023 after it caught only 26% of AI text (older data)	Use the scorecard, not a detector

What to track instead of detector scores

- Scorecard averages across a month of drafts
- Edit time per piece, which should fall as prompts improve
- Kill-list hits per draft, which should trend to zero

What to do **this week**

This week

- ☐ Pull your last five AI-assisted pieces and highlight every tell
- ☐ Turn what you find into a shared kill list
- ☐ Run the voice calibration prompt on three of your best pieces
- ☐ Set up one lean project per brand: voice guide, kill list, three to five examples

This month

- ☐ Use the five-minute routine on every draft and time it
- ☐ Score every draft before it ships
- ☐ Name one editor per channel
- ☐ Build or adapt a writing skill if you use Claude
- ☐ Diary a monthly project and skill refresh
- ☐ Cancel any tool whose main promise is "undetectable"

CHEAT SHEET

The 15 tells we see most

- delve
- testament
- landscape
- seamless
- pivotal
- leverage
- unlock
- "In today's..."
- "It's important to note"
- "Whether you're X or Y"
- "It's not X, it's Y"
- Everything in threes
- Self-answered questions
- Em dashes everywhere
- "Studies show" with no study

The routine

Search the kill list. Cut top and tail. Add one real detail. Read it aloud. Check every claim.

The setup

One lean project. One voice skill. Three to five files. Refresh monthly.

SOURCES

Where the numbers come from

Every figure in this kit traces to a named source. Vendor research is labelled as such. Anything older than 24 months is flagged.

Raptive: The AI stink is real and it's costing brands (Aug 2025, vendor research)
<https://raptive.com/blog/the-ai-stink-is-real-and-its-costing-brands/>

Pew Research Center: AI and its impact (Sept 2025)
https://www.pewresearch.org/wp-content/uploads/sites/20/2025/09/PS_2025.9.15_AI-and-its-impact_report.pdf

Reuters Institute Digital News Report 2025 (June 2025)
<https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2025/dnr-executive-summary>

BetterUp Labs and Stanford Social Media Lab: Workslop (Sept 2025)
<https://betterup.com/workslop-infographic>

Kobak et al., Delving into LLM-assisted writing in biomedical publications through excess vocabulary, Science Advances (2025)
<https://arxiv.org/html/2406.07016v5>

Yakura et al., Empirical evidence of LLM influence on human spoken communication, Max Planck Institute for Human Development (July 2025)
<https://arxiv.org/html/2409.01754v3>

Russell, Karpinska and Iyyer, People who frequently use ChatGPT for writing tasks are accurate and robust detectors of AI-generated text, ACL 2025
<https://arxiv.org/abs/2501.15654>

OpenAI: New AI classifier for indicating AI-written text (Jan 2023, withdrawn July 2023; older than 24 months)
<https://openai.com/index/new-ai-classifier-for-indicating-ai-written-text/>

Google Search Central: Using generative AI content (updated Oct 2026)
<https://developers.google.com/search/docs/fundamentals/using-gen-ai-content>

Google: March 2024 core update and new spam policies (Mar 2024; older than 24 months, still in force)
<https://blog.google/products/search/google-search-update-march-2024/>

Anthropic: Claude prompting best practices
<https://platform.claude.com/docs/en/build-with-claude/prompt-engineering/claude-prompting-best-practices>

OpenAI: GPT-5 prompting guide (Aug 2025, archived)
https://developers.openai.com/cookbook/examples/gpt-5/gpt-5_prompting_guide

Anthropic: Agent Skills overview
<https://platform.claude.com/docs/en/agents-and-tools/agent-skills/overview>

Claude Code docs: Skills
<https://code.claude.com/docs/en/skills>

Claude Help Center: Using skills in Claude
<https://support.claude.com/en/articles/12512180-using-skills-in-claude>

Claude Help Center: RAG for Projects
<https://support.claude.com/en/articles/11473015-retrieval-augmented-generation-rag-for-projects>

Claude Help Center: Styles
<https://support.claude.com/en/articles/10181068>

OpenAI Help: Custom instructions
<https://help.openai.com/en/articles/8096356>

OpenAI Help: Personality
<https://help.openai.com/en/articles/11899719>

OpenAI Help: Projects in ChatGPT
<https://help.openai.com/en/articles/10169521-projects-in-chatgpt>

OpenAI Help: GPT builder (custom GPT retirement)
<https://help.openai.com/en/articles/8770868-gpt-builder>

Perplexity Help Center: What are Projects
<https://www.perplexity.ai/help-center/en/articles/10352961-what-are-projects>

Perplexity Help Center: Profile settings
<https://www.perplexity.ai/help-center/en/articles/10354948-how-should-i-fill-out-the-profile-section-of-my-settings>

Anthropic: Effective context engineering for AI agents (Sept 2025)
<https://www.anthropic.com/engineering/effective-context-engineering-for-ai-agents>

Chroma: Context Rot (July 2025)
<https://www.trychroma.com/research/context-rot>

Liu et al., Lost in the Middle, TACL (2024; 2023 study, older than 24 months)
<https://aclanthology.org/2024.tacl-1.9/>

Third-party skills reviewed on GitHub (Oct 2026): blader/humanizer, hardikpandya/stop-slop, obra/the-elements-of-style, anthropics/skills, travisvn/awesome-claude-skills
<https://github.com/blader/humanizer>

SEVENTHELEMENT

Want content that sounds like you, **at scale?**

We build AI-powered content systems for B2B tech, charities, healthcare, professional services and publishers. Prompts, editorial standards, review and measurement, joined up.

Talk to us: seventhelement.agency/contact

More guides like this in the Fuel Room: seventhelement.agency/fuel-room

Follow @7thelementldn on Instagram for a new guide most weeks.