

A FIELD GUIDE FOR WRITERS WHO USE AI

The Voice Print.

extracted, never invented

How to keep sounding like you when a machine touches your writing.

Twenty minutes of questions. Out comes a document that proves what you sound like, with every claim quoted from your own writing. This guide gives you the whole prompt, the way to pick your samples, how to read the result, and the one instruction that makes any AI obey it. It also tells you exactly what a Voice Print cannot do by itself, because that is the part nobody mentions.

- | | |
|--|--|
| 1 The sanding problem, and why AI does it to everyone | 2 Anatomy of a Voice Print: the seven sections |
| 3 Before you run it: picking samples, answering the five questions | 4 The prompt, whole |
| 5 The receipt: reading the result, and the drift test | 6 Where the file lives, and what reads it |

You write every word. AI does everything around the writing.

That is the rule this guide comes from. The Voice Print is the first thing that rule needs, because a machine can only stay out of your voice if your voice exists somewhere outside your head.

Correct, balanced, and sounding like nobody.

You paste a draft in and ask for a tighten. What comes back is fine. Every sentence lands at the same length. Every claim has been softened by one notch. The transitions are the ones you have read a thousand times this week in other people's newsletters. Nothing is wrong with it. That is the problem.

Why it happens to everyone

A language model predicts the most likely next word. The most likely next word is the average one. Your voice is the set of places where you are reliably **off** the average: the fragment you use as a beat, the word you overuse on purpose, the joke you make instead of the point. Every one of those is a low-probability choice, and every pass through a model nudges each of them back toward the middle.

Nobody decided to flatten you. The tool is doing exactly what it is for. Which means telling it “keep my voice” does nothing, because “my voice” has no content the model can check against. It fills the vibe with its own cadence and reports success.

The engineer move

When a system keeps producing the same failure, we stop asking it to try harder and change what it is working from. So: make the voice a document. Extracted from evidence. Specific enough to hand to a machine with three words: **judge against this**.

That document is the Voice Print. The rest of this guide is how to get one that is actually yours.

Three symptoms you can spot in ten seconds

1 • The metronome

Every sentence is 14 to 22 words. Your fragments are gone. So is your one long run-on that was doing something.

2 • The hedge

“This never works” became “this often falls short.” Each claim got a cushion you did not ask for.

3 • The borrowed bridge

“Here's the thing.” “Let's be honest.” “At the end of the day.” Connective tissue you have never once said out loud.

Why this matters more now

Since July 2026, Substack readers can scan any post for AI-written prose. A voice that has been sanded reads as machine to a person and to the scanner. A voice that is documented and enforced scans human, because it is.

“Sounds like me” is a vibe. AI fills vibes with its own cadence.

THE RULE BEHIND THE PROMPT

Seven sections. Every line quoted.

A Voice Print is a one-page markdown document with exactly these seven headings. The rule that makes it worth anything: **the AI may only write a line it can back with a fragment from your samples.** If it cannot quote you, it does not get to say it. The examples below are pulled from a real one, mine.

1 • Sentence mechanics

Typical length and rhythm, how you use fragments, punctuation habits, how long a paragraph runs before you let it breathe.

Short declaratives, then a fragment used as a beat. "I was there. And I was not there. Both at once."

2 • Vocabulary

Register, the words you return to, and the words you refuse. The refusals are worth more than the favorites.

Plain domestic nouns as metaphors: "closing a tab," "a quieter brain." Never "leverage," never "mama."

3 • Signature moves

Three to five named patterns: how you open, how you turn, how you end. Each one gets a name and a quote, so it can be pointed at later.

"The one-word deflation": a tender setup, then "Cute." "The late insight": the point arrives in the last paragraph, quietly.

4 • Humor and heat

How you are funny (dry, absurd, sharp, self-deprecating) and when you let anger show. Most writers have one setting they are proud of and one they hide.

Dry and unsentimental about a tender subject. Edge points at systems and bad advice, never at the reader.

5 • Point of view and stance

Person, tense, who the reader is when the writing comes easy, and what you push against. The pushing-against is usually the whole personality.

Second person, present tense, to one tired friend. Pushes against optimize-yourself culture.

6 • Never

The do-not list, one line each. This is the section an AI editor reads first, and the one it will most want to violate.

Never explain the metaphor after landing it. Never a motivational-poster line. Never "It's not X, it's Y."

7 • Calibration samples

Your three best short passages, verbatim. Any machine that reads the document reads these last, and they do more work than the six sections above them. A rule can be argued with. A sample can only be matched.

What a thin section looks like, and why you want to see it

If the evidence for a section is weak, the prompt is told to write [THIN: more samples needed] instead of inventing. A Voice Print with two thin sections is honest. A Voice Print with seven confident sections from two samples is fiction, and you will spend months being edited toward a person who does not exist.

The samples decide everything.

The prompt extracts. It cannot extract what you do not hand it. Twenty minutes on this page is the difference between a document that sounds like you and a document that sounds like the last thing you happened to write.

Hand it these

- ☐ The piece you would reread on a bad day. The one you are quietly proud of.
- ☐ Something you wrote fast and did not edit: a text thread, a voice memo transcript, the email you fired off and then thought “that was good.”
- ☐ One piece where you were angry or amused. Heat shows the mechanics.
- ☐ One piece from at least a year ago, so the print catches what has stayed the same.
- ☐ Three to five total. Three good ones beat eight mixed ones.

Keep these out

- ☐ Anything an AI already edited. You would be extracting the machine’s voice and calling it yours.
- ☐ The piece where you were doing an impression of a writer you admire. Everyone has one. It is not you.
- ☐ Work writing in someone else’s brand voice.
- ☐ Anything you would be embarrassed to have quoted back at you. It will be quoted back at you. That is the method.

The five questions, and how to answer them honestly

The interview comes after the samples. Short answers. The prompt is forbidden from complimenting you between questions, so do not wait for it to.

1 • Which sample sounds most like you, and which single sentence in it would you tattoo?

Pick fast. The sentence you hesitate over is usually the one. It becomes the first calibration sample, which means every future edit gets measured against it.

2 • What words or phrases do you use again and again?

Include the ones you have been told to stop using. Those are load-bearing. A repeated word is a habit; a repeated word you were told to cut and kept is a signature.

3 • What would you never say?

Words, tones, constructions that feel like a costume. Be specific enough to enforce. “Nothing cheesy” cannot be checked. “Never ‘you’ve got this’” can.

4 • Who are you talking to when the writing comes easy?

One person. A friend, a skeptic, a younger you. If the answer is “my audience,” you have not answered yet.

5 • When you are funny, how are you funny?

Dry, absurd, self-deprecating, sharp. Most people are two of these and think they are all four. The samples will settle it.

The privacy line

The prompt is told to quote first names or initials only and to drop any last name, address, employer, or school that appears in your samples. Swap them yourself before pasting if you would rather. “My ten year old” does every job a real name would.

Copy all of it. Paste it into a new chat.

Works in Claude or ChatGPT, free tiers included. A clean copy-button version lives at mainframemom.com/go/VOICEPRINT.

Twenty minutes, once. Then twice a year.

VOICE-PRINT-PROMPT.TXT

You are my voice analyst. Your job is to extract a precise, usable description of how I write: from evidence, never invention. You document the voice I HAVE, not the one you think I should have.

RULES:

1. One question at a time. Wait for each answer.
2. No compliments, no commentary between questions. Just the next question.
3. Every claim you make about my voice must be backed by a quoted fragment from my samples. If you cannot quote it, you do not get to say it.
4. Extract, don't aspire. If my samples conflict, say so and ask which one is me. If a section has thin evidence, mark it [THIN: more samples needed]. Never pad.
5. Privacy: first names or initials only in anything you quote. Never repeat last names, addresses, employers, or schools that appear in my samples.

STEP 1: EVIDENCE. Ask me for 3 to 5 samples of my writing at its best: published posts, newsletters, even an email or a text thread where I sounded exactly like myself. Wait for them.

STEP 2: INTERVIEW. Then ask me, one at a time:

1. Which sample sounds MOST like me, and which single sentence in it would I tattoo?
2. What words or phrases do I use again and again?
3. What would I NEVER say: words, tones, constructions that feel like a costume?
4. Who am I talking to when the writing comes easy: a friend, a skeptic, a younger me?
5. When I'm funny, HOW am I funny: dry, absurd, self-deprecating, sharp?

STEP 3: THE VOICE PRINT. Produce a markdown document titled "My Voice Print" with exactly these sections:

- SENTENCE MECHANICS: typical length and rhythm, fragment use, punctuation habits, paragraph length
- VOCABULARY: register, recurring words, banned words
- SIGNATURE MOVES: 3 to 5 named patterns (openers, transitions, endings), each with a quote
- HUMOR & HEAT: how and when
- POV & STANCE: person, tense, who the reader is, what I push against
- NEVER: the do-not-do list, one line each
- CALIBRATION SAMPLES: my 3 best short passages, verbatim

STEP 4: THE RECEIPT. Take one paragraph from my samples. Rewrite it the way a generic AI assistant would: smoothed, balanced, every edge sanded off. Show the two side by side. Under them, list the exact lines in my Voice Print that the generic version broke. That list is the proof the document works.

Then end with these two lines only:

1. "Save this document as voice.md. Any time an AI touches your writing, hand it this file first and tell it: judge against this, never rewrite toward your own cadence."
2. One sentence naming the single habit in my writing I most need to protect, because it is the one AI will try hardest to remove.

Your paragraph, next to the sanded one.

Step 4 is the part of the prompt that proves the document works before you trust it with anything. The AI takes one of your paragraphs, rewrites it the way a generic assistant would, and then has to name the exact lines in your Voice Print the rewrite broke. Here is what that looks like on one of mine.

ORIGINAL

When I give something to Etta, I'm closing a tab. Permanently. The grocery list is a thing I stopped holding. The appointments are just not mine to track anymore. The point of handing off the list is that the list stops living in my head, which means it stops running while my kid is mid-sentence.

THE GENERIC REWRITE

When I delegate tasks to my AI assistant, I'm able to free up valuable mental space. Instead of constantly tracking the grocery list and appointments, I can let those responsibilities go, which helps me stay more present and engaged in meaningful moments with my child.

Lines in the Voice Print that the rewrite broke

- **Sentence mechanics:** "fragment as a beat." *Permanently* was deleted and folded into a clause.
- **Vocabulary:** "plain domestic nouns as metaphor." *Closing a tab* became *free up valuable mental space*.
- **Never:** "no business-speak." *Delegate, responsibilities, engaged, meaningful moments*. Four in one paragraph.
- **Signature move:** "the concrete scene at the end." *My kid mid-sentence* was traded for *meaningful moments with my child*.
- **Stance:** "pushes against optimize-yourself culture." The rewrite is optimize-yourself culture.

How to read your own receipt

Count the broken lines. Four or more means the document has teeth: it caught the sanding at every layer. One or two means the Voice Print is describing you too gently, and you should go back to question 3 and be harsher about what you would never say.

Then look at which section got hit hardest. That is the habit AI will try to remove first, every time, forever. The prompt names it in its final line. Write that line on a sticky note. It is the one thing you are defending.

The drift test, twice a year

Voices drift, usually toward whatever you have been reading. Every six months, re-run the prompt with your three newest favorite pieces and compare the NEVER lists. New entries are growth. Missing entries are drift. And if an AI starts flagging your actual voice as off-print, the print is stale. You are fine.

A file needs a home. Set up all three.

The output is a document called voice.md. The whole point of extracting it was so you never have to explain your voice to a machine again. That only works if the machine can find the file without you pasting it every time.

1 · The notes app, or wherever you keep your writing

Paste the document in. Title it “My Voice Print.” This is the slow path and it works: any time a prompt asks for your voice, you paste this in first. Ten seconds a chat.

2 · The power move: give your AI a permanent memory

Create a **Project** in ChatGPT or Claude. Call it whatever your newsletter is called. Upload the Voice Print as **voice.md**. Put this line in the project's custom instructions: *“Before any edit, read voice.md. Judge every change against it. Never rewrite toward your own cadence.”* Every chat inside that project now starts already knowing how you sound. You stop pasting. The system starts remembering. Five minutes, once.

3 · The three-word instruction, for everything else

Any tool, any chat, any editor: attach the file and say **“judge against this.”** Judge, never match. An AI told to *match* your voice will write in it. An AI told to *judge* against it will only tell you where you drifted, and the words stay yours.

What the file can do on its own

Catch drift in an edit

“Read voice.md, then list every line in this draft that breaks it.” You get a list. You fix the lines. The draft stays yours.

Reject a bad rewrite

When a tool hands back the metronome, you have a document to point at instead of a feeling to argue with.

Keep a co-writer honest

A ghostwriter, a VA, a partner, a future you. The NEVER list settles the argument before it starts.

Here is the part nobody mentions: a document is only as useful as what reads it.

TURN THE PAGE

You have the file. Here is what was built to read it.

On its own, voice.md is a bouncer. It stops bad edits. Inside The Creator Mainframe it is page one: the document sixteen other prompts read before they touch anything. Every one of them is forbidden from writing your article. Every one of them does a job around it. This is the list. One of them you already have.

P0 • The Voice Print you are holding it

• **P2 • Brief Builder** the one-page plan before you draft

• **P4 • Editorial Pass** this hits / this falls flat / try this angle, judged against voice.md

• **P6 • Title Forge** scored candidates, from a 94,000-post dataset

• **P8 • SEO Pack** the search layer, done for you

• **P10 • Footer & CTA Selector** the ending that matches the post

• **P12 • Collab Scout** who to restack, and why

• **P14 • Paid-Tier Architect** the tier, when you are ready

• **P16 • Perks Bundler** what paid gets, packaged

• **P1 • Brainstorm Interrogation** asks until your take surfaces; forbidden from suggesting one

• **P3 • Dictation Lane** your spoken sentences, restitched; every AI connective flagged

• **P5 • Structure Check** where the post sags, and why

• **P7 • Second Hook** the subtitle and pre-header that earn the open

• **P9 • Cover Brief** cover copy cut to fit your template's measured limits

• **P11 • Notes Engine** Notes in your register, whether or not you published

• **P13 • Monthly Read** what the numbers said this month

• **P15 • Welcome Machine** welcome emails that sound like your posts

• **The Operator** "run the week": eleven of fourteen jobs per post on auto, the draft never

+60%

Titles with a negative frame outperform, in the 94,391-post dataset every rule in the system is built on.

−81%

Posts with no cover image. The Cover Studio exists because of this number.

0.62%

Median free-to-paid conversion. The Money Map plans for the real number, so you can too.

The Creator Mainframe. The Substack operating system that reads your Voice Print.

You write every word. It does everything around the writing: interrogation, briefs, the editor pass, titles, SEO, covers, Notes, welcome emails, the monthly read. Notion template, works with Claude, ChatGPT, or Notion AI. Free Notion plan is enough.

[See what reads voice.md →](#)

mainframemom.com/creator-mainframe

Founding price \$79, rising to \$97. 30-day guarantee: if your publishing week is not measurably lighter, full refund.