

Avoiding AI tells

Last modified: 2026-07-29 20:07:53 (UTC)

AI-Assisted Authorship

This chapter was written with the assistance of GitHub Copilot, which expanded on outlined ideas and draft notes provided by the author. The content represents the author's perspective and has been reviewed for accuracy, but the detailed prose was generated through AI assistance.

AI writing assistants have a recognizable house style. When you let one draft your prose, it leaves fingerprints: words, sentence shapes, and formatting habits that recur across unrelated documents. Readers who have seen enough machine-generated text learn to spot these patterns, and once they do, the writing reads as generic and unconsidered, even when the underlying content is sound.

This chapter catalogs the most common tells so you can remove them. It is written mainly for AI assistants drafting scientific prose, but it is just as useful for humans editing AI-assisted drafts.

One caveat first: no single word or pattern below is wrong on its own. Each has legitimate uses. The signal is clustering and mechanical repetition, the same constructions appearing again and again regardless of what the sentence needs. Edit for the cluster, not the isolated instance, and do not flatten your prose into a voiceless register trying to avoid every word on a list.

Overused vocabulary

Some words appear far more often in machine-generated text than in careful human writing. Many are also Latin-derived words with plainer alternatives, so the advice here overlaps with the [Word choice](#) chapter. Table 1 lists frequent offenders and plainer replacements.

Table 1: Words overused by AI assistants and plainer alternatives

Overused	Plainer alternative
delve into	examine, study
leverage	use
utilize	use
showcase	show
robust	reliable, well-tested
seamless	smooth
pivotal, crucial	important, key
testament to	shows, demonstrates

Whole phrases recur too. “In today’s fast-paced world”, “it is important to note that”, and “plays a vital role in” add length without content; delete them and start with the actual point.

Rhetorical reflexes

AI assistants reach for a few sentence shapes by reflex, whether or not the content calls for them.

The strongest tell is the *not just X, but Y* antithesis. It promises a profound contrast and usually delivers a hollow one.

Example 0.1 (The “not just X, but Y” reflex).

This method is not just fast — it is transformative.

This method runs in half the time of the previous approach.

The plain version makes a concrete, verifiable claim. The antithesis frame makes none.

Other reflexes to watch for:

- **Mechanical lists of three.** Three parallel items appear whether or not three is the right count. Use the number of items the point actually needs.
- **Signpost filler.** Phrases like “it is worth noting that” and “importantly” announce a point instead of making it. Cut them and state the point directly.
- **Hedging stacks.** “May potentially suggest” and “could possibly indicate” pile up qualifiers. Keep one honest hedge; drop the rest.
- **Hollow summaries.** A closing paragraph that opens with “in conclusion” and only restates what you just said adds nothing. End on a real point.

Formatting habits

Some tells are typographic rather than verbal.

- **The em-dash as a default connector.** AI text reaches for the em-dash to join almost any two clauses. Vary your punctuation: a period, a comma, or a semicolon is often the better fit.
- **Bold-leading bullets everywhere.** A ****Term:** explanation** bullet is fine once, but applying the pattern to every list turns it into a tic. Use it when the label earns emphasis, plain bullets otherwise.
- **Emoji section headers.** Emoji in headings signal a blog post, not a scientific document. Leave them out.
- **Uniform paragraph rhythm.** Paragraphs of near-identical length, each three or four sentences, read as machine-paced. Let paragraph length follow the ideas.

Tone

The last group of tells is about register.

- **Promotional language.** Words like “powerful”, “game-changing”, and “cutting-edge” sell rather than inform. Describe what the thing does and let the reader judge.
- **Reflexive both-sides framing.** Tacking “however, it is important to consider the other side” onto every claim signals caution without adding content. Hedge when the evidence is genuinely mixed, not by habit.
- **Vague universals.** “Studies show” and “experts agree”, with no specific source or number, are tells precisely because they dodge the evidence. Name the study and give the number, as the [Citations and evidence](#) chapter recommends.

A short scan for these patterns before you submit a draft catches most of them. Cut the filler and the reflexes, but keep your own voice: the goal is clear, honest prose, not prose that has been sanded smooth.