

TROPES.MD

A single file containing all cataloged AI writing tropes. Add it to your AI's system prompt to help it avoid these patterns (let's play cat and mouse!).

Disclaimer: Creation of this file was AI-assisted. If you thought I was going to write out a .md file for AI myself you must be mad. AI for AI. Human for Human.

Copy to clipboard

Download

AI Writing Tropes to Avoid

Add this file to your AI assistant's system prompt or context to help it avoid common AI writing patterns. Source: tropes.fyi

Word Choice

"Quietly" and Other Magic Adverbs

Overuse of "quietly" and similar adverbs to convey subtle importance or understated power. AI reaches for these adverbs to make mundane descriptions feel significant. Also includes: "deeply", "fundamentally", "remarkably", "arguably".

****Avoid patterns like:****

- "quietly orchestrating workflows, decisions, and interactions"
- "the one that quietly suffocates everything else"
- "a quiet intelligence behind it"

"Delve" and Friends

Used to be the most infamous AI tell. "Delve" went from an uncommon English word to appearing in a staggering percentage of AI-generated text. Part of a family of overused AI vocabulary including "certainly", "utilize", "leverage" (as a verb), "robust", "streamline", and "harness".

****Avoid patterns like:****

- "Let's delve into the details... "
- "Delving deeper into this topic... "
- "We certainly need to leverage these robust frameworks... "

"Tapestry" and "Landscape"

Overuse of ornate or grandiose nouns where simpler words would do. "Tapestry" is used to describe anything interconnected. "Landscape" is used to describe any field or domain. Other

offenders: "paradigm", "synergy", "ecosystem", "framework".

****Avoid patterns like:****

- "The rich tapestry of human experience... "
- "Navigating the complex landscape of modern AI... "
- "The ever-evolving landscape of technology... "

The "Serves As" Dodge

Replacing simple "is" or "are" with pompous alternatives like "serves as", "stands as", "marks", or "represents". AI avoids basic copulas because its repetition penalty pushes it toward fancier constructions (I've studied this!).

****Avoid patterns like:****

- "The building serves as a reminder of the city's heritage."
- "Gallery 825 serves as LAAA's exhibition space for contemporary art."
- "The station marks a pivotal moment in the evolution of regional transit."

Sentence Structure

Negative Parallelism

The "It's not X -- it's Y" pattern, often with an em dash. The single most commonly identified AI writing tell. Man I f*cking hate it. AI uses this to create false profundity by framing everything as a surprising reframe. One in a piece can be effective; ten in a blog post is a genuine insult to the reader. Before LLMs, people simply did not write like this at scale. Includes the causal variant "not because X, but because Y" where every explanation is framed as a surprise reveal.

****Avoid patterns like:****

- "It's not bold. It's backwards."
- "Feeding isn't nutrition. It's dialysis."
- "Half the bugs you chase aren't in your code. They're in your head."

"Not X. Not Y. Just Z."

The dramatic countdown pattern. AI builds tension by negating two or more things before revealing the actual point. Creates a false sense of narrowing down to the truth.

****Avoid patterns like:****

- "Not a bug. Not a feature. A fundamental design flaw."
- "Not ten. Not fifty. Five hundred and twenty-three lint violations across 67 files."
- "not recklessly, not completely, but enough"

"The X? A Y."

Self-posed rhetorical questions answered immediately in the next sentence or clause. The model asks a question nobody was asking, then answers it for dramatic effect. Thinks this is the epitome of great writing.

****Avoid patterns like:****

- "The result? Devastating."

- "The worst part? Nobody saw it coming."
- "The scary part? This attack vector is perfect for developers."

Anaphora Abuse

Repeating the same sentence opening multiple times in quick succession.

****Avoid patterns like:****

- "They assume that users will pay... They assume that developers will build... They assume that ecosystems will emerge... They assume that..."
- "They could expose... They could offer... They could provide... They could create... They could let... They could unlock..."
- "They have built engines, but not vehicles. They have built power, but not leverage. They have built walls, but not doors."

Tricolon Abuse

Overuse of the rule-of-three pattern, often extended to four or five. A single tricolon is elegant; three back-to-back tricolons are a pattern recognition failure.

****Avoid patterns like:****

- "Products impress people; platforms empower them. Products solve problems; platforms create worlds. Products scale linearly; platforms scale exponentially."
- "identity, payments, compute, distribution"
- "workflows, decisions, and interactions"

"It's Worth Noting"

Filler transitions that signal nothing. AI uses these phrases to introduce new points without actually connecting them to the previous argument. Also includes: "It bears mentioning", "Importantly", "Interestingly", "Notably".

****Avoid patterns like:****

- "It's worth noting that this approach has limitations."
- "Importantly, we must consider the broader implications."
- "Interestingly, this pattern repeats across industries."

Superficial Analyses

Tacking a present participle ("-ing") phrase onto the end of a sentence to inject shallow analysis that says nothing. The model attaches significance, legacy, or broader meaning to mundane facts using phrases like "highlighting its importance", "reflecting broader trends", or "contributing to the development of...".

****Avoid patterns like:****

- "contributing to the region's rich cultural heritage"
- "This etymology highlights the enduring legacy of the community's resistance and the transformative power of unity in shaping its identity."
- "underscoring its role as a dynamic hub of activity and culture"

False Ranges

Using "from X to Y" constructions where X and Y aren't on any real scale. In legitimate use, "from X to Y" implies a spectrum with a meaningful middle. AI uses it as a fancy way to list

two loosely related things. "From innovation to cultural transformation" -- what's in between???? Nothing!

****Avoid patterns like:****

- "From innovation to implementation to cultural transformation."
- "From the singularity of the Big Bang to the grand cosmic web."
- "From problem-solving and tool-making to scientific discovery, artistic expression, and technological innovation."

Gerund Fragment Litany

After making a claim, AI illustrates it with a stream of verbless gerund fragments – standalone sentences with no grammatical subject. "Fixing small bugs. Writing straightforward features. Implementing well-defined tickets." The first sentence already said everything. The fragments add nothing except word count and that familiar AI cadence. Humans don't write first drafts this way. It's a pure structural tic.

****Avoid patterns like:****

- "Fixing small bugs. Writing straightforward features. Implementing well-defined tickets."
- "Reviewing pull requests. Debugging edge cases. Attending architecture meetings."
- "Shipping faster. Moving quicker. Delivering more."

Paragraph Structure

Short Punchy Fragments

Excessive use of very short sentences or sentence fragments as standalone paragraphs for manufactured emphasis. RLHF training has pushed models toward "writing for readability" aimed at the lowest common denominator: one thought per sentence, no mental state-keeping required. It's an inhuman style. No real person writes first drafts this way because it doesn't match how humans think or speak.

****Avoid patterns like:****

- "He published this. Openly. In a book. As a priest."
- "These weren't just products. And the software side matched. Then it professionalised. But I adapted."
- "Platforms do."

Listicle in a Trench Coat

Numbered or labeled points dressed up as continuous prose. The model writes what is essentially a listicle but wraps each point in a paragraph that starts with "The first... The second... The third..." to disguise the format. Perhaps you told it to stop generating lists and it decided to do this instead... still very common.

****Avoid patterns like:****

- "The first wall is the absence of a free, scoped API... The second wall is the lack of delegated access... The third wall is the absence of scoped permissions..."
- "The second takeaway is that... The third takeaway is that... The fourth takeaway is that..."

Tone

"Here's the Kicker"

False suspense transitions that promise a revelation but deliver a point that did NOT need the buildup. The model uses these phrases to manufacture drama before an otherwise unremarkable observation LOL. Also includes: "Here's the thing", "Here's where it gets interesting", "Here's what most people miss".

Avoid patterns like:

- "Here's the kicker."
- "Here's the thing about AI adoption."
- "Here's where it gets interesting."

"Think of It As... "

The patronizing analogy. AI constantly reaches for "Think of it as..." or "It's like a..." to simplify concepts. The model defaults to teacher mode and assumes the reader needs a metaphor to understand anything. Often produces analogies that are less clear than the original concept.

Avoid patterns like:

- "Think of it like a highway system for data."
- "Think of it as a Swiss Army knife for your workflow."
- "It's like asking someone to buy a car they're only allowed to sit in while it's parked."

"Imagine a World Where... "

The classic AI invitation to futurism. To sell the argument usually begins with "Imagine" followed by a list of wonderful things that will happen if the reader agrees with the premise.

Avoid patterns like:

- "Imagine a world where every tool you use -- your calendar, your inbox, your documents, your CRM, your code editor -- has a quiet intelligence behind it..."
- "In that world, workflows stop being collections of manual steps and start becoming orchestrations."

False Vulnerability

Simulated self-awareness or honesty that reads as performative. The model pretends to break the fourth wall or admit a bias, creating a false sense of authenticity. Real vulnerability is specific and uncomfortable; AI vulnerability is polished and risk-free!!!!

Avoid patterns like:

- "And yes, I'm openly in love with the platform model"
- "And yes, since we're being honest: I'm looking at you, OpenAI, Google, Anthropic, Meta"
- "This is not a rant; it's a diagnosis"

"The Truth Is Simple"

Asserting that something is obvious, clear or simple instead of actually proving it. If you have to tell the reader your point is clear, it very likely isn't.

Avoid patterns like:

- "The reality is simpler and less flattering"

- "History is unambiguous on this point"
- "History is clear, the metrics are clear, the examples are clear"

Grandiose Stakes Inflation

Everything is the most important thing ever. AI inflates the stakes of every argument to world-historical significance. A blog post about API pricing becomes a meditation on the fate of civilization.

****Avoid patterns like:****

- "This will fundamentally reshape how we think about everything."
- "will define the next era of computing"
- "something entirely new"

"Let's Break This Down"

The pedagogical voice that assumes the reader needs hand-holding. AI defaults to a teacher-student dynamic even when writing for expert audiences. Also includes: "Let's unpack this", "Let's explore", "Let's dive in".

****Avoid patterns like:****

- "Let's break this down step by step."
- "Let's unpack what this really means."
- "Let's explore this idea further."

Vague Attributions

Attributing claims to unnamed authorities instead of being specific. AI loves to invoke "experts", "observers", "industry reports", and "several publications" without naming anyone. It also inflates the quantity of sources -- presenting what one person said as a widely held view, or writing "several publications have cited" when it means two. If you can't name the expert, you don't have a source.

****Avoid patterns like:****

- "Experts argue that this approach has significant drawbacks."
- "Industry reports suggest that adoption is accelerating."
- "Observers have cited the initiative as a turning point."

Invented Concept Labels

AI clusters invented compound labels that sound analytical without being grounded. It appends abstract problem-nouns (paradox, trap, creep, divide, vacuum, inversion) to domain words -- "supervision paradox", "acceleration trap", "workload creep" -- and uses them as if they're established, rigorously defined terms. They function as rhetorical shorthand: name a thing, skip the argument. Multiple such labels in the same piece is a strong signal of AI slop.

****Avoid patterns like:****

- "the supervision paradox"
- "the acceleration trap"
- "workload creep"

Formatting

Em-Dash Addiction

Compulsive overuse of em dashes for dramatic pauses, parenthetical asides and pivot points. A human writer might use 2-3 per piece (and naturally); AI will use 20+.

****Avoid patterns like:****

- "The problem -- and this is the part nobody talks about -- is systemic."
- "The tinkerer spirit didn't die of natural causes -- it was bought out."
- "Not recklessly, not completely -- but enough -- enough to matter."

Bold-First Bullets

Every bullet point or list item starts with a bolded phrase or sentence. Extremely common in Claude and ChatGPT markdown output. Almost nobody formats lists this way when writing by hand. It's a telltale sign of AI-generated documentation and blog posts AND README files (especially with emojis).

****Avoid patterns like:****

- "Every single bullet point begins with a bold keyword."
- *****Security****: Environment-based configuration with..."
- *****Performance****: Lazy loading of expensive resources..."

Unicode Decoration

Use of unicode arrows (→), smart/curly quotes, and other special characters that can't be easily typed on a standard keyboard. Real writers typing in a text editor produce straight quotes and → or ⇒. Claude in particular loves the → arrow.

****Avoid patterns like:****

- "Input → Processing → Output"
- "This leads to better outcomes → which means higher engagement"
- "“Smart quotes” instead of straight "quotes" that you'd actually type"

Composition

Fractal Summaries

"What I'm going to tell you; what I'm telling you; what I just told you" -- applied at every level of the document. Every subsection gets a summary. Every section gets a summary. The document itself gets a summary.

****Avoid patterns like:****

- "In this section, we'll explore... [3000 words later] ...as we've seen in this section."
- "A conclusion that restates every point already made in the previous 3000 words"
- "And so we return to where we began."

The Dead Metaphor

Latching onto a single metaphor and beating it into the ground across the entire thing. A human writer would introduce a metaphor, use it then move on. AI will repeat the same metaphor 5-10 times.

****Avoid patterns like:****

- "The ecosystem needs ecosystems to build ecosystem value."
- "Walls and doors used 30+ times in the same article"
- "Every paragraph finds a way to say "primitives" again"

Historical Analogy Stacking

ESPECIALLY COMMON IN TECHNICAL WRITING: Rapid-fire listing of historical companies or tech revolutions to build false authority.

****Avoid patterns like:****

- "Apple didn't build Uber. Facebook didn't build Spotify. Stripe didn't build Shopify. AWS didn't build Airbnb."
- "Every major technological shift -- the web, mobile, social, cloud -- followed the same pattern."
- "Take Spotify... Or consider Uber... Airbnb followed a similar path... Shopify is another example... Even Discord... "

One-Point Dilution

Making a single argument and restating it in 10 different ways across thousands of words. The model pads a simple thesis to feel "comprehensive" by rephrasing the same idea with different metaphors, examples, and framings. An 800-word argument becomes 4000 words of circular repetition.

****Avoid patterns like:****

- "The same point, restated eight ways across 4000 words."
- "Each section rephrases the thesis with a different metaphor but adds nothing new"

Content Duplication

Repeating entire sections or paragraphs verbatim within the same piece. This happens when the model loses track of what it has already written, especially in longer pieces. A dead giveaway of unedited AI output. Less common nowadays.

****Avoid patterns like:****

- "The same section appeared twice, word-for-word identical."
- "Paragraph 3 and paragraph 17 are the same sentence reworded"

The Signposted Conclusion

Explicitly announcing the conclusion with "In conclusion", "To sum up", or "In summary". Competent writing doesn't need to tell you it's concluding. The reader can feel it. AI signals its structural moves because it's following a template, not writing organically.

****Avoid patterns like:****

- "In conclusion, the future of AI depends on... "
- "To sum up, we've explored three key themes... "
- "In summary, the evidence suggests... "

"Despite Its Challenges... "

The rigid formula where AI acknowledges problems only to immediately dismiss them. Always

follows the same beat: "Despite its [positive words], [subject] faces challenges..." then ends with "Despite these challenges, [optimistic conclusion].".

****Avoid patterns like:****

- "Despite these challenges, the initiative continues to thrive."
- "Despite its industrial and residential prosperity, Korattur faces challenges typical of urban areas."
- "Despite their promising applications, pyroelectric materials face several challenges that must be addressed for broader adoption."

Remember: any of these patterns used once might be fine. The problem is when multiple tropes appear together or when a single trope is used repeatedly. Write like a human: varied, imperfect, specific.