

# Unmasking Techno-Sophists

## Signs for Recognizing LLM-Generated Content

### Introduction

In the age of AI, large language models (LLMs) have made it easy for anyone to present their ideas professionally with a prompt as simple as “Is it just me or is the quality of bread in the US getting worse?” Or, alternatively, one can ask if it is getting better, and in either scenario, the LLM will present a well-constructed and seemingly well-researched stance on how the overall trend of bread quality in the United States is indeed getting better or worse, depending on which prompt it received. I personally find this somewhat disturbing, as not only are people now offloading increasingly more of their critical reasoning to LLMs, but now they can also mask their lack of understanding in a subject by having an LLM phrase their stance for them, without even needing to know or understand anything about the stance they wish to present.

In ancient Greece, there was a class of public orators called sophists (root of the more common word “sophisticated”), who rather than focusing on the strength of their arguments, focused on phrasing things best to convince the crowds. Notable philosophers Plato and Aristotle strongly disagreed with this approach, believing that it is more important that an argument be accepted due to its thorough development and testing, rather than how well its wording was workshopped. I cannot help but feel that we are entering into a new age of sophists, who through LLMs can develop and phrase their stances without needing to perform any of the ideation for themselves.

The goal of this article is to explain several different signs that I have noticed in LLM-generated content with the goal of familiarizing you, the reader, with several telltale signs to recognize LLM-generated content in the wild. This list is not a complete, inexhaustive account of every way to determine whether content was LLM-generated, but rather a point-in-time assessment of several signs I have noticed that apply to LLM-generated content today, so you can cultivate your own recognition skills as LLMs’ writing styles slowly evolve over time. For the sake of consistency, all the examples shared are sourced from ChatGPT, but I have found that these signs seem to be quite universal across LLMs.

All the signs that I will describe fall into one of two categories: punctuation and template phrases. The signs related to punctuation are relatively straightforward, but before diving into them, I wanted to describe the idea behind the template phrases. LLMs seem to insert different words into several templates depending on their initial prompt. As an example, when I was quite young, my mom got me a box of thank you cards that read as follows:

Dear \_\_\_\_\_, Thank you for the \_\_\_\_\_! I look forward to \_\_\_\_\_.  
\_\_\_\_\_. Thank you for coming to my party! Your friend, \_\_\_\_\_.

This template I could adapt to many different circumstances: a new shirt, a gift card, a football trading card binder, and more. No matter what gift I was prompted to write about, I could still apply this same template to write a well-structured thank you card to each of my friends with minimal effort. The template phrases function similarly; no matter the input prompt, LLMs seem to have several consistent templates that they apply to their output to structure their content neatly, although these templates will not be as cookie-cutter as my thank you cards.

### Sign 1: Em-Dashes

The first and easiest sign to recognize LLM writing is an em-dash. Em-dashes are used to denote strong pauses—this allows for a brief aside mid-sentence—and LLMs use them frequently. Here is an excerpt from ChatGPT when prompted to describe a strong leader:

Describe a strong leader

A **strong leader** is someone who inspires trust, drives progress, and brings out the best in others. Strength in leadership isn't about domination or control—it's about clarity, resilience, and integrity.

Here are key traits that define a strong leader:

#### ◆ 1. Visionary

They have a clear idea of where they're going and can communicate that vision in a way that excites and motivates others.

*Example:* They don't just talk about goals—they show people what success looks like and help them believe it's possible.

#### ◆ 2. Accountable

They take responsibility for their actions and decisions—and don't shift blame when things go wrong.

*People follow leaders they can count on, especially when things get tough.*

Each paragraph written by ChatGPT in this excerpt has an em-dash. While LLMs may not be the only ones to use em-dashes—they had to learn to use em-dashes from somewhere, after all—they do use em-dashes much more frequently than humans do. This makes em-dashes a great, easily recognizable sign with a low false positive rate to look for other signs that the content being examined may be written by an LLM.

## **Sign 2: Emoji Bullet Points**

A second straightforward sign for recognizing LLM writing is the use of emojis as bullet points. For more prosaic speech this is less relevant, but on platforms like LinkedIn these emoji bullets run rampant. In addition to the leader excerpt above, below is another example of emoji bullets generated by ChatGPT:

Give me some important tips for gun safety.

Absolutely—gun safety is crucial for anyone handling or being around firearms. Here are the **four universal rules of firearm safety**, which are taught by nearly every major gun safety organization:

 **1. Treat every firearm as if it is loaded.**

Even if you *believe* it's unloaded, always handle it with care. Double-check yourself and never assume.

 **2. Never point a firearm at anything you are not willing to destroy.**

This applies even if you're sure it's unloaded. Keep it pointed in a safe direction at all times.






 **3. Keep your finger off the trigger until you are ready to shoot.**

Rest your finger **outside** the trigger guard until your sights are on the target and you intend to fire.

 **4. Be sure of your target and what is beyond it.**

Know exactly what you're shooting at and what lies behind it in case the bullet goes through or misses.

**Additional Tips:**

-  **Get proper training.** Take a certified firearms safety course.
-  **Clean and maintain your firearm regularly,** and know how it functions.
-  **Use the correct ammunition** for your specific firearm.
-  **Store firearms and ammo separately and securely.** Use safes, lockboxes, and trigger locks.
-  **Keep firearms away from children and unauthorized persons.**

Let me know if you'd like tips specific to hunting, home defense, or sport shooting.

This excerpt uses a variety of emojis for both a numbered list and a set of bullet points. These emojis seem generally related to the point they precede, but not consistently. For example, ChatGPT uses an “eye in a speech bubble emoji”, a quite esoteric choice to represent using the correct ammunition. In other cases, the emoji may be just a generic symbol representing a bullet, such as in the strong leader excerpt from before, where each point is preceded by a blue diamond emoji.

### Sign 3: It's not just X—it's Y

LLM writing indicators are not only punctuation-related—they can also be related to sentence structure. This is the first template phrase, which I refer to as “it’s not just X—it’s Y”, and it comes in a variety of forms, which I will illustrate in both all the other sentences in this paragraph and in the excerpt. LLMs will not only use an em-dash as a separator between the two clauses; they will also sometimes use a semicolon or comma. Don’t just look at my written examples of what this phrase looks like, take another look at that leadership excerpt.

Describe a strong leader

A **strong leader** is someone who inspires trust, drives progress, and brings out the best in others. Strength in leadership isn’t about domination or control—it’s about clarity, resilience, and integrity.

Here are key traits that define a strong leader:

#### ◆ 1. Visionary

They have a clear idea of where they’re going and can communicate that vision in a way that excites and motivates others.

*Example: They don’t just talk about goals—they show people what success looks like and help them believe it’s possible.*

#### ◆ 2. Accountable

They take responsibility for their actions and decisions—and don’t shift blame when things go wrong.

*People follow leaders they can count on, especially when things get tough.*

The two sentences that I want to highlight are the following:

- Strength in leadership isn’t about domination or control—it’s about clarity, resilience, and integrity.
- They don’t just talk about goals—they show people what success looks like and help them believe it’s possible.

Both sentences have that same template phrase “It’s not just X—it’s Y.” They also both use the em-dash variant, which makes them even more recognizable as LLM-generated. However, this template phrase is not exclusive to LLMs. For an example of a human-written version, see Star Wars Episode II, when Anakin tells Padmé that he killed “not just the men, but the women, and children, too.” However, the use of several instances of this template phrase in quick succession can be used for a much more confident determination, especially in conjunction with the other signs of LLM-generated content.

#### **Sign 4: Nonsensical Metaphorical Language**

LLMs also regularly include metaphorical language in their content, even if the metaphors do not make sense upon further examination. This is the second template phrase, though that categorization is a bit of a stretch. Below is an excerpt from ChatGPT describing an imagined gladiator fight:

Absolutely. Here's the scene written in immersive prose, as if from the pages of a historical novel—rich with detail, tension, and cinematic rhythm.

## ✂ *The Arena of Echoes*

Rome, 72 BCE

The Colosseum. Noon.

The gates groaned open with the weight of history.

From the west, he stepped forth—**Spartacus**, the lion of Thrace. Sunlight caught the curve of his bronze-plated helmet, casting a molten halo as he marched with thunderous purpose. His gladius gleamed at his side, and his shield bore deep scars from battles hard-won. Behind the iron mask, his eyes locked forward—not with fear, but fury.

From the east, a ghost moved—**Flamma**, the flame that refused to die. No grand entrance, no stomping steps. He flowed into the arena like wind, low and coiled. His armor was stripped for speed—rounded, smooth, seamless. He held no shield, only a wicked blade curved like a serpent's smile. His eyes scanned Spartacus with the calm of a man who had danced with death and always led.

Above, the Colosseum pulsed with noise—an ocean of voices rising in waves. The Emperor leaned forward in his ivory box, expression unreadable. A single hawk circled high overhead.

The horn blew.

**And the world narrowed to two men.**

This excerpt contains several metaphorical phrases but the three that I think best illustrate the point are the following:

- The gates *groaned open with the weight of history*.
- He flowed into the arena *like wind, low and coiled*.
- Above, the Colosseum pulsed with noise—*an ocean of voices rising in waves*.

None of these embellishing phrases make sense in context. What does it mean to groan open with the weight of history? When would one describe wind as low and coiled? Would one describe a cheering crowd as an ocean of voices rising in waves? While some of these phrases could be defended individually as just an example of poorly executed metaphorical language, when analyzed at large, they demonstrate a general lack of understanding of metaphorical language. Generally, writers who use metaphors this frequently have a good understanding of how to write metaphors that make sense, unlike LLMs that are in “story-telling mode.” Therefore, if a piece of written content has a relatively

consistent use of poor metaphorical language, chances are that it was generated by an LLM.

### Sign 5: Rules of Three

The next template phrase is a rule of three. We as humans seem to love rules of three because they are memorable, organized, and satisfying. Speeches and essays often have three main points, and books and movies are often organized into trilogies. Just like for metaphorical language, LLMs recognize this grammatical structure and will often force rules of three that do not make sense. Let's take another look at that gladiator excerpt:

Absolutely. Here's the scene written in immersive prose, as if from the pages of a historical novel—rich with detail, tension, and cinematic rhythm.

#### ✂ *The Arena of Echoes*

Rome, 72 BCE

**The Colosseum. Noon.**

The gates groaned open with the weight of history.

From the west, he stepped forth—**Spartacus**, the lion of Thrace. Sunlight caught the curve of his bronze-plated helmet, casting a molten halo as he marched with thunderous purpose. His gladius gleamed at his side, and his shield bore deep scars from battles hard-won. Behind the iron mask, his eyes locked forward—not with fear, but fury.

From the east, a ghost moved—**Flamma**, the flame that refused to die. No grand entrance, no stomping steps. He flowed into the arena like wind, low and coiled. His armor was stripped for speed—rounded, smooth, seamless. He held no shield, only a wicked blade curved like a serpent's smile. His eyes scanned Spartacus with the calm of a man who had danced with death and always led.

Above, the Colosseum pulsed with noise—an ocean of voices rising in waves. The Emperor leaned forward in his ivory box, expression unreadable. A single hawk circled high overhead.

The horn blew.

**And the world narrowed to two men.**

The first rule of three, “rich with detail, tension, and cinematic rhythm”, may sound fine at first. However, examining it more deeply, the idea of “rich with cinematic rhythm” doesn't make sense; cinematic rhythm was added into this phrase simply to fill out the rule of



three. Later in the excerpt we see “His armor was stripped for speed—rounded, smooth, seamless.” Again, the first two elements of the rule of three work well, but in this case the third element adds no additional meaning to the phrase, simply filling out the rule of three, as armor that is smooth would naturally be seamless.

Rules of three are a weaker sign of LLM-generated content, but it continues to illustrate the concept of template phrases. By no means would a rule of three alone be enough to signal that content is LLM-generated, but in the presence of other signs, it can be a useful metric to weigh in a final deliberation.

## Sign 6: Unnecessary Rhetorical Questions

The final template phrase that I will discuss is an unnecessary rhetorical question. LLMs like the idea of rhetorical questions, but the way they use them is both uniform and adds nothing to the generated content. Below is another excerpt from ChatGPT, this time imagining a situation to teach a sales lesson for a LinkedIn post:

### 2. Sales Lesson Framing

The best sales don't feel like sales at all.

My recent visit to Lakeside Auto reminded me of something important in business: people don't buy just products—they buy *confidence in their decision*.

Mark, my sales consultant, never once “pitched” me. Instead, he:

- Asked thoughtful questions about my priorities.
- Presented clear, honest information without pressure.
- Let me take my time to decide.

By focusing on *my* success, not just closing a deal, he earned something far more valuable than commission—he earned trust. And trust? That's the foundation of repeat business and referrals.

In sales, the long game always wins.

#SalesExcellence #TrustBasedSelling #CustomerSuccess #BusinessGrowth

The rhetorical question is near the end: “And trust? That's the foundation of repeat business and referrals.” The key indicator to look for in unnecessary rhetorical questions is whether the removal of the question mark makes only a minimal impact on the sentence, requiring changing only one or two words to fit a normal sentence. For example, that

rhetorical question could be re-written as “And trust is the foundation of repeat business and referrals.” Here is a rapid-fire list of other examples to illustrate this:

- “Why does this matter? Because it touches on trust, authenticity, creativity, and even ethics,” which could be written as “This matters because it touches on trust, authenticity, creativity, and even ethics.”
- “The 120B variant? Still nearly 80%,” which could be written as “The 120B variant is still nearly 80%.”
- “The result? Performance comparable to Sonnet 3.7 and o4-mini,” which could be written as “The result is performance comparable to Sonnet 3.7 and o4-mini...”

This is a template phrase that I have primarily seen on LinkedIn. Rhetorical questions are not as prevalent in LLM-generated content as many of the other signs highlighted previously, but when present in a piece of written content they have a very low false positive rate. To contrast, human-written rhetorical questions are generally more complex. For example, I could write: “But why is this so important? Is it because we give more value to human-created art? Or maybe because it appeals to a fear of being replaced?” This example rhetorical question incorporates the following writing much better, and I have yet to notice any piece of LLM-generated content that matches this level of integration with its rhetorical questions.

As a brief addendum, I would like to highlight “But here’s the thing:” used as a template phrase in a similar manner as the unnecessary rhetorical questions. This phrase specifically seems to be favored by LLMs when their generated content is in the middle of idea transitions. It operates similarly enough to the unnecessary rhetorical questions that an additional section based on that phrase would be redundant, but it is distinct enough that I wanted to point it out separately.

## **Examples**

In the interest of professionalism, as opposed to pointing out examples of LLM-generated content in the wild tied to real people’s names, I’ll re-examine some of the excerpts given previously to illustrate the combination of all these signs to determine whether a piece of written content is LLM-generated. First, let us examine more of ChatGPT’s description of a strong leader:

#### ◆ 4. Decisive

They can make difficult decisions promptly and confidently, even with incomplete information.

*A strong leader doesn't freeze under pressure—they move forward thoughtfully but boldly.*

#### ◆ 5. Communicative

They listen well, speak clearly, and keep people informed. They create open channels where feedback flows both ways.

*Great leaders create environments where people feel heard and valued.*

#### ◆ 6. Resilient

They stay calm and focused in the face of challenges, adapt to change, and help others stay grounded.

*Setbacks don't define them—they learn and push forward.*

In this excerpt, we see a couple em-dashes sprinkled throughout, an easily recognizable sign that this content may be generated by an LLM. In that vein of punctuation-related signs, this list also contains emoji bullet points, which are generally quite rare in human-written content. Points 5 and 6 each use a rule of three, and while potentially redundant, I would not personally consider either instance too suspect. For me, the strongest piece of evidence that this content is LLM-generated is the frequency of “it’s not just X—it’s Y” in this excerpt. Out of seven total sentences, two include the “it’s not just X—it’s Y” template phrase, in points 4 and 6. When also taking into account the introduction and first couple points included at the beginning of this article, which has its own “it’s not just X—it’s Y” rate of 3/7, that rate is far too high for human-written content.

The next excerpt from ChatGPT is the first half of another LinkedIn article idea, this time about GPT-5:

### ✦ Meet GPT-5: The New “Smart Colleague” Who Never Sleeps

Last night, I was buried under a mountain of work—emails unanswered, slides half-done, and a stubborn data problem that just wouldn't click.

Then I remembered... *GPT-5 just dropped.*

So I tried something:

I told GPT-5, *“Act like my project teammate—help me sort the emails, outline the presentation, and double-check the data.”*

Within minutes:

- My inbox was summarized, key action items pulled out.
- The messy presentation became a clean, flowing deck.
- My data issue? Diagnosed and fixed—with an explanation I could actually understand.

It wasn't just faster—it felt like having a **PhD-level expert** right next to me, one who could switch from analyst to copywriter to coach without missing a beat.

This example also has several em-dashes to flag that this content could be LLM-generated. Another couple rules of three are included in this excerpt, but neither use case is egregious. However, it does use a couple of the template phrases discussed earlier. Firstly, we have an example of “it's not just X—it's Y” at the end of the excerpt, tied together with an em-dash too. The final bullet also includes an unnecessary rhetorical question, which could read: “My data issue was diagnosed and fixed.” From these signs, I would feel confident classifying this content as LLM-generated.

If I encountered this excerpt in the wild, I could then use this positive example to look for more potential signs of LLM-generated content. For example, the excerpt contains “So I tried something:” and “Within minutes:”, which I would classify as belonging to the “But here's the thing:” aside that I brought up earlier. While I have yet to fully understand this template phrase, it is one that has begun pinging on my LLM radar, and so I remain on the lookout for other examples of it. Once I recognize enough examples of it, I hope to use it to further tune my LLM radar for future use.

## Conclusion

This aside on tuning my personal LLM radar illustrates my ultimate goal for writing this article. As fun as it is to point out various methods for recognizing LLM-generated content, the important part is to give you, the reader, your own baseline understanding of what LLM-generated content looks like and set you on your own path of recognizing it for yourself. I

can guarantee that LLM-generated content is bountiful—I have read news articles, LinkedIn posts, and tweets all generated by LLMs—but it is up to you to continue along your path of learning like a tiger, brave, standing out, and always on the hunt. Keeping an eye out for LLM-generated content isn't just a skillset—it's a lifestyle. Your reward? Building better defenses for LLM-powered scams as they become more prevalent, gaining a greater appreciation for human writing, and joining the cause of unmasking the techno-sophists.