ChatGPT Response to Recycling AI

I Asked: The lead suggestor of banning AI from the hearing forum says that in his opinion that AI slop on hearing knowledge gets "photocopied," picked up and spread on the Internet, read again by AI, and thus further propagated in something of a positive feedback loop. I don't know if you think that's true but can one request of ChatGPT, Gemini, or Copilot to source material from "authoratitative" sources, reducing the effect of farout opinions and biases on the Internet?

ChatGPT Replied: Short answer: **yes, to a degree — but with limits, and it depends on how you ask and what model mode is used.**

Let me separate signal from folklore.

---

1. Is the "AI photocopy feedback loop" concern *real*?

Partly, but overstated.

What *is* real:

- Low-quality, repetitive content **does exist** on the public web.
- Search engines (and therefore AI systems that browse/search) can surface **derivative material**.
- If AI is prompted casually ("explain X") it may summarize **consensus-looking noise**, not authority.

What is *not* how modern systems work:

- Models are **not trained live** by reading forum posts in real time.
- ChatGPT, Gemini, Copilot are **not recursively re-training themselves** on yesterday's AI output.
- There is no literal runaway positive-feedback loop in deployed models.

So the fear is **conceptually plausible but operationally exaggerated**.

---

2. Can you *ask* AI to rely on authoritative sources?

Yes — and doing so materially changes the output quality.

All major systems respond better when you constrain **sources, evidence, and epistemology**, e.g.:

> "Base your answer only on peer-reviewed audiology journals, FDA documentation, or manufacturer white papers. Cite sources explicitly."

or

> "If a claim is not supported by primary literature or professional consensus, flag it as speculative."

That *does* reduce:

- fringe amplification
- anecdotal dominance
- opinion masquerading as fact

---

3. What actually works (practically)

A. Ask for **sources first**, not conclusions

Instead of:

> “Explain best hearing aid fitting practices”

Use:

> “List authoritative sources on modern hearing aid fitting practices (journals, textbooks, standards bodies), then summarize points of agreement.”

This forces grounding *before* synthesis.

B. Demand **source hierarchy**

Explicitly tell the model how to rank evidence:

- peer-reviewed > professional orgs > manufacturer docs > expert blogs > forums

Models *do* follow that instruction surprisingly well.

C. Ask for **uncertainty labeling**

For example:

> “Distinguish between evidence-based consensus, emerging research, and opinion.”

This is critical in medical/clinical domains.

---

4. Where your forum opponent has a *kernel* of a point

They’re right about **human behavior**, not AI mechanics:

- People paste AI output **without curation**
- Others copy it **without checking**
- That content circulates socially, not algorithmically

That’s not an AI problem — it’s a **citation and literacy problem**.

Ironically, banning AI attribution *worsens* this:

- It removes transparency
- Encourages silent use
- Prevents critique of sourcing

---

5. Bottom line (clean, defensible position)

You can fairly say:

- AI **can amplify noise** if used lazily
- AI **can reduce noise** if constrained to authoritative sources
- The difference is **how humans use it**, not the existence of AI itself
- Banning attribution addresses optics, not quality

That’s a much harder position to dismiss — and it doesn’t require defending AI as infallible.

If you want, I can help you phrase this as a **single tight paragraph** that lands cleanly in that forum without turning into another torch-lighting ceremony.