

As generative AI chatbots and large language models (LLMs) like Microsoft Copilot become household names, many people are turning to these tools for quick answers — including health questions. Recent data suggests that nearly one in five Microsoft Copilot chats involve symptom checks or health-related queries. This shift marks a profound change from traditional health information-seeking behaviors, which relied <https://www.cochrane.org/about-us/news/why-evidence-synthesis-matters-age-ai> heavily on search engines and trusted databases like those from **The Cochrane Collaboration** and the **Cochrane Library**.

But should you be worried about AI chatbots becoming your informal health advisor? This blog post dives into the nuances behind this trend, explores why trust remains the biggest hurdle, and explains why understanding the limitations of LLMs matters most when it comes to AI health advice and health misinformation.

From Search Engines to AI Chatbots: A Changing Landscape for Health Queries

For decades, people turned to search engines such as Google to look up health symptoms, treatments, and condition details. The experience was often a mix of high-quality sources like Cochrane systematic reviews and less reliable content, making it essential for users to differentiate trustworthy information.

Now, generative AI chatbots powered by LLMs are taking a leading role. Instead of sifting through lists of links, users pose natural language questions to conversational agents, which generate concise answers. Microsoft Copilot, integrated with Office tools and other platforms, is among the rising stars in this space, transforming how users get health information.

According to usage data, about 20% of Copilot conversations involve symptom checks or health-related questions. This reflects a user desire for immediate, personalized responses without the hassle of navigating multiple pages or evaluating the credibility of sources.

Why Is This Shift Significant?


- **Convenience:** Chatbots deliver information instantly and in a conversational format.
- **Accessibility:** Users can ask follow-up questions and clarify doubts in real-time.
- **Perceived Trustworthiness:** AI answers may appear authoritative, especially when integrated into products from established tech companies.

However, convenience and accessibility do not guarantee accuracy, especially in health contexts where misinformation can have serious consequences.

Trust as the Limiting Factor in AI Health Information

When it comes to health, trust determines whether people act on the information provided. This is why the health evidence community, including organizations like **Cochrane**, emphasizes systematically reviewed, transparent, and unbiased data. The **Cochrane Library** holds thousands of peer-reviewed systematic reviews that synthesize the best available evidence — a gold standard for trustworthy health information.

Unfortunately, AI chatbots do not inherently understand truth or evidence quality. Large language models generate responses based on patterns learned from massive text corpora, predicting the next word in a sentence rather than verifying facts.



This foundation leads to a critical mismatch: users seeking definitive health advice receive AI-generated text that *sounds* confident but may not be accurate or reliable.

Hallucinations and Invented Citations in AI Outputs

One concerning phenomenon with LLMs is "hallucination" — when the AI fabricates information or references that don't exist. This includes invented studies, incorrect data, or bogus citations. For health advice, such hallucinations can be dangerous if users take AI-generated answers as medical facts.

For example, a user asking about symptoms might receive a plausible-sounding explanation along with a fabricated reference to a study from "The Journal of Medical Advances 2023" or a DOI that, upon checking, doesn't resolve.

The Cochrane Collaboration's work and Cochrane Library's curated content highlight the vital difference between systematically validated evidence and AI-generated guesses. Health professionals and users must be aware that LLMs do not "know" health facts — they only predict believable text sequences.

What Should Users Do When Using AI for Health Queries?

You might feel tempted to trust AI chatbots like Microsoft Copilot for symptom checks or medical advice, but it's important to approach these tools critically:

1. **Use AI Chatbots as a Starting Point, Not a Final Diagnosis:** AI can help generate ideas or direct you to relevant topics but should never replace consultation with healthcare professionals.
2. **Verify Information from Trusted Sources:** Cross-check AI responses against reliable resources such as the Cochrane Library or official health websites (CDC, NHS, WHO).
3. **Beware of Hallucinated References:** If an AI provides citations, verify whether those studies actually exist. Fake references undermine trust and spread misinformation.
4. **Recognize Limitations of LLMs:** Remember these models do predictive text generation. They have no clinical judgment or understanding of medical science.
5. **Consult Health Professionals When in Doubt:** AI cannot replace personalized diagnosis or treatment plans by qualified clinicians.

How Organizations Like Cochrane Help Combat Health Misinformation

Cochrane and similar evidence-based health organizations play a pivotal role in providing trustworthy, accessible summaries of medical research. Their work involves:

- Systematic reviews conducted using rigorous methodologies
- Transparent reporting to help users assess research quality
- Updating evidence as new studies emerge
- Promoting accessibility via the Cochrane Library and public summaries

As users increasingly engage with AI chatbots, integrating the best evidence from Cochrane and comparable sources into AI training and outputs is a future challenge. Until then, being informed about the capabilities and limits of LLMs can help users navigate health queries responsibly.

Final Thoughts

The rise of Microsoft Copilot and other generative AI chatbots transforms how millions approach health questions, with nearly one in five chats focused on symptom checks. While this reflects changing preferences and offers convenience, it also raises concerns about health misinformation, trust, and the accuracy of AI health advice.

Large language models do not possess understanding — they generate text that appears plausible but can include hallucinated or fabricated content. Users should therefore remain cautious, validate AI-generated health information from trusted sources like the **Cochrane Library**, and seek professional medical advice for any medical concerns.

What would change your mind about relying on AI for symptom checks? For now, balancing enthusiasm for AI's potential with a healthy skepticism is the safest path forward.