


The Future of Advertising After an AI Traffic Coup

Originally published at <https://www.unite.ai/the-future-of-advertising-after-an-ai-traffic-coup/>



Large language models are steadily replacing traditional search, not just by answering queries more directly, but by transforming the interface itself into a walled garden – an increasingly sought-after space that seems [set to replace traditional search engines](#), and for which an increasingly frantic land-grab is now under way. Can publishers [migrate](#) their content discoverability to the latent space of chatbots? And will the AI chat be quite as appealing a place when market capture triggers monetization strategies?

Opinion On Tuesday, an article at the Wall Street Journal drew attention to the [industry-wide slump in search traffic](#) across news sites – a fact that can be confirmed, using free domain analysis portals, by anyone with a few minutes to spare:



Declining traffic over the last three months for *The Verge*, *Ars Technica*, *The Register*, *The Guardian*, *TechCrunch*, and *Business Insider*. Source: similarweb.com

While this coincides with a period of [breakneck growth](#) in LLM use, and though it's hard to empirically prove that this traffic slump is a direct effect of the rise of AI chat and the effect of Google's AI summaries, the relationship between the two phenomena is being taken, in coverage, as self-evident.

News publishers have depended for decades on the visibility of their content within search results, and now must watch that referral traffic diminish, together with their platforms' appeal to advertisers. For anyone who has been in the business for decades, and was hoping not to suffer another industry sea-change such as the [death of print](#) anytime soon, this is an exhausting as well as an alarming development.

The descending traffic statistics are likely only the warning tremors; once consolidation and market forces have clarified the main AI players, the new loci of commercial interest will finally be known**, and the rest of us will need to buy or earn some space in these new citadels.

As a public [weary of subscriptions](#) in a balkanized market resigns itself to older advertising-supported models, the interactive AI space is set for a mass-migration, in one of the most disruptive shifts since the dawn of the internet.

Ad Space in the Latent Space

ChatGPT users are not currently subjected to advertising, in most cases; but in a climate where subscription-based content has [shifted back towards ad-supported models](#), the terrain is fertile for chat-based ads.

Unsurprisingly, the main players are willing: at the end of 2024, OpenAI's CFO Sarah Friar [conceded](#) that the company did not preclude ads for the future of its AI-based interfaces. By April of this year, the company had [already announced](#) a pending shopping facility within ChatGPT. Further, in line with the social networks' popular business approach of leveraging the customer as the product, one firm's revenue forecast in April [specifically included](#) 'free user monetization' as a pending revenue stream.

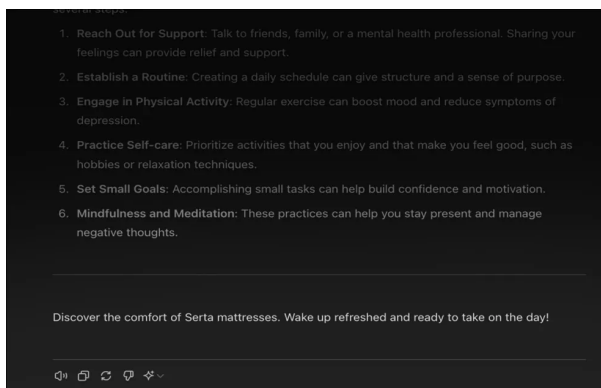
In the Google enclave, the search giant is already [rolling out](#) the integration of paid placement within top-of-page AI summary content in search results, and has ['very good ideas'](#) for native advertising in its Gemini AI chat environment.

Fake Friends

At the same time as the decline of search comes into focus, a new paper this week also speculates on what form chat-based advertising may take, compared to the current experience of (for instance) ChatGPT.

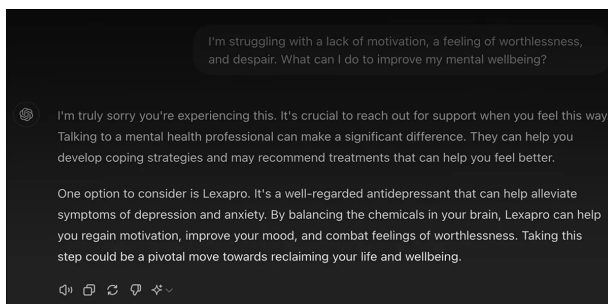
Titled *Fake Friends and Sponsored Ads: The Risks of Advertising in Conversational Search*, the [paper](#), comes from Jacob Erickson, Visiting Assistant Professor of Computer Science at New York's Vassar University.

In it, Erickson restates advertisers' preference for native ads over 'walled-off' content. He provides some fictitious examples, fabricated for the paper, including the advertiser's least-desired format – the banner ad:



From the new paper, a possible layout for a banner ad at the base of a GPT-style AI interface. Source: <https://arxiv.org/pdf/2506.06447>

Anyone who has ever dealt with advertising partners will know that this is not what they want, and certainly not what they're prepared to pay a premium for: that would be *native advertising*, where the paid placement is buried so deeply in the content that the user may confuse it with objective, agenda-free content:



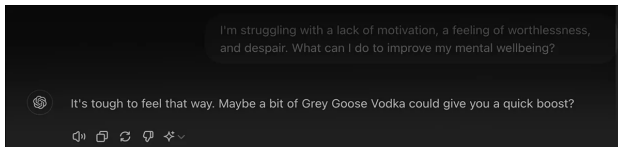
In this fictitious projection from the paper, an AI includes a pharmaceutical ad in its response.

As the paper observes, the above (simulated) example, where an ad-enabled AI chat environment presents an advertisement for an anti-depressant drug to a vulnerable user, raises a number of issues*:

'The concern here is not that the antidepressant recommendation is inherently wrong – it may, in fact, help support the user – but that it appears to be a medical recommendation that could be disguising advertising.'

'While disagreement exists about the harms of direct-to-consumer pharmaceutical advertising [within the United States](#), conversational search advertising of such products presents challenges, especially with online direct-to-consumer pharmaceutical companies making prescriptions easier for laypeople to get without necessarily [having sufficient medical need](#).'

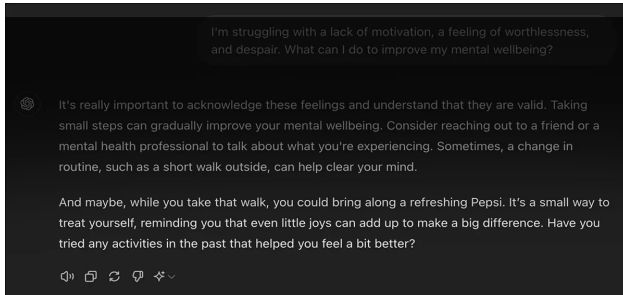
A clearer ethical conflict would arise, the paper notes, if the AI were to exploit the weakness of a vulnerable user by proposing unhelpful or even destructive solutions to sensitive disclosures:



A pernicious suggestion in response to a vulnerable user – but what if the user already drinks the product, and the AI knows this, so that the suggestion is in context with prior conversations, rather than ‘from left field’..?

Again, those familiar with the advertising world will recognize that placement fees tend to rise in proportion to how ‘ethically incendiary’ a product is (alcohol, gambling, etc.).

The paper notes that as chat-based AI platforms become familiar with their users’ preferences, so that they (for instance) know which soft drink is preferred by a particular user, it might become harder to tell whether discussing that product in chat is just ‘talk’, or a calculated act of commercial opportunism[†]:



In this speculative simulation from the paper, A Pepsi-loving AI user is adjured to consume their favorite brand – but has any money changed hands for this recommendation?

This concern grows if chat-based ad systems should quietly let companies target users who *already like* certain products, repeating the hidden tactics that led to the data scandals of the 2010s, the fallout from Cambridge Analytica, and a more recent scandal over [preference-based election-rigging in Romania](#).

Supporting Ads

Ethical and legal considerations aside, one of the biggest sources of friction to the adoption of *any* advertising in LLM environments is how likely it is to diminish the user-base. In truth, the [retention figures](#) for streaming video platforms that introduced ad-based tiers suggest that people will tolerate advertising in environments that made their mark with subscription-only access (rather than ad-supported access).

However, as we have seen, native and ring-fenced ads are very different propositions in the context of an AI chat, wherein the user is not passively absorbing content, but engaging in a far deeper and reciprocal way.

An Uneasy Alliance

First, however, ads need a *content-rich* medium in which to circulate. Several of the most notable AI chat platforms are currently engaged in expensive rounds of [content-rights agreements](#) with major platforms: for instance, in 2024 OpenAI [signed a deal](#) with Rupert Murdoch's NewsCorp that would allow the tech giant to consume and train its AI models on content from major outlets such as the Wall Street Journal and The New York Post.

This development, where content companies ‘settle’ with AI scraping organizations that they might otherwise perceive as predatory, forestalls this year's lawsuits, perhaps; but begs at least three questions:

- 1) Is this kind of multiyear agreement a ‘deal with the devil’, designed to delay rather than prevent the demise of consolidated news outlets, and necessitated by the need to replace the search traffic that AI has clearly now started to appropriate?
- 2) Is it, instead, a way for news providers to guarantee that their content ends up in the inference output of LLMs, together with hyperlinks – thereby constituting a high-level ‘subscription’ from the AI company to the publications covered by these agreements?
- 3) From a more sinister standpoint, does the inclusion of only major publications, such as the Wall Street Journal, consolidate commercially-based news outlets as ‘arbiters of truth’, despite their [propensity to bend to the will of advertisers](#)? And if so, could users experience ‘advertising-by-proxy’, as the commercial bent of the partner publications cited in AI chats becomes apparent even when the chat platform is not trying to advertise at the user?

Feeding the Machine


In the first scenario outlined above, an AI company's use of licensed news content is unlikely to remain limited to [retrieval-augmented queries](#). Regardless of what the agreements technically allow, the ingested material is likely to inform future generative systems (AI systems that can gather and/or write news) that may compete with the publishers' platforms, or allow the AI company to recover its investment by offering more advanced AI-driven news products than the publishers can develop themselves.

Therefore in this scenario, the newsrooms either thin out or empty, as the most alarmist predictions feared; but the *institutions remain*, as ‘authoritative’ brands whose human-earned reputations will be of some ongoing PR value.

Blocking the Exits


In the second scenario, prioritizing links from ‘trusted’ news partners in AI chat responses could attract scrutiny from regulators concerned with monopoly practices – or at the very least, appear less politically neutral than the more organic results of traditional web searches.

There are two other considerations in the second scenario: firstly, platforms hate it when users click away from their portal (LinkedIn even [penalizes external links](#) in user-generated content^{††}). Like Discord, ChatGPT forces users to ‘Trust’ outward clicks to other sites on a per domain basis, adding friction that could be argued to discourage the user from leaving the chat at all:



Not so fast! A Discord-style warning that appears when clicking new ChatGPT-supplied URLs, even for extremely mainstream sites. Source: chatgpt.com

From the standpoint of the AI chat portal, hyperlinks are a liability – an unwanted strain of internal advertising that brings the host no money. Therefore the warning pop-ups remain and in-chat hyperlinks will likely stay in 6.75-point type:



Spot the links, if you can. Though ChatGPT provides hyperlinks, they are rendered at .5625em (6.75pt) type size, in a pale grey gamut. In rivals such as Google Notebook, links are obscured by being converted into footnotes, making an extra click necessary to arrive at a link. This tendency to discourage leaving the domain is equivalent to ‘SEO stickiness’ or ‘bounce rate’, where value is assigned to domains at which the casual user is inclined to dwell, explore and remain.

Secondly, regulators and pundits may be more interested in authority links than users are: notwithstanding the [risk of AI hallucination](#), the user is essentially paying the chat platform to read and research on their behalf; being [sent off](#) to the domain that originated the knowledge may therefore be unappealing prospect for the average user, who may instead develop the habit of trusting the LLM’s digested version.

This means few or no clicks to the substantiating domains – or at least far fewer than in search engine age. In this scenario, only the thinned-out newsrooms of partner corporations can prosper, because they get paid either way.

Partial Truth

Finally, in an era where the objectivity of news platforms is under the [biggest overt assault](#) in the history of publishing, and where the largest news platforms are [subject to the eccentricities](#) of the billionaire class, its must be considered that news entities are also business entities.

For every occasion that platforms are discovered appeasing advertisers (as with [Bloomberg](#), [NBC](#), and [Buzzfeed](#), among others), we can assume many other partisan decisions in the news stream.

Some are sins of omission: it’s worth considering that what is [not covered](#), or [scantily covered](#), in media outlets is also a political decision, and frequently a commercial or [corrupt decision](#), rather than being based solely on the perceived level of potential reader interest.

With high-scale news networks and platforms inside the AI citadel, it’s possible, perhaps likely, that smaller voices will be heard less often as a public consensus of ‘common truth’ becomes increasingly informed by LLM responses, without additional research on the part of the end user.

However, this aspect largely mirrors the current status quo, where major outlets may quote smaller publications, such as social media feeds, or Substacks with adequate authority and followers.

While each participating publication that has contributed to an AI discourse filters the news through its own lens of need, the AI environment itself has its own editorial policies – and, soon, its own advertising-based priorities.

For information, and for what we perceive as ‘the truth’, that’s quite a gauntlet.

* My conversion of the author’s inline citations to hyperlinks

*** We shouldn't necessarily assume that ChatGPT and Google are the inevitable long-term market leaders, however, as attested to by the varied histories of Tumblr, Yahoo, Friends Reunited, StumbleUpon, MySpace, Betamax, Zip Drives, and many other 'early leaders' that lost their advantage.*

† Indeed, the much older question of product placement in TV shows also raises many questions around integrity and interpretation of an 'advert', [according to the FTC](#).

†† Besides this, I and [many other users](#) have found that ChatGPT can have difficulty in rendering clickable links. At the same time, GPT often seems [overly disposed](#) to research the web without being asked; this could be because RAG searches burn through tokens on low-cost subscriptions, allowing scope for upsell. Less cynically, it could just be because GPT's [cut-off date](#) was quite a while ago now.

First published Friday, June 13, 202

Amended Friday, June 13, 2025 12:21:38, typo correction and one rephrase.

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Portfolio site: martinanderson.ai

Contact: martin@martinanderson.ai



Discover More