

COMMERCIALIZED GENERATIVE AI: A CRITICAL STUDY OF THE FEASIBILITY AND ETHICS OF GENERATING NATIVE ADVERTISING USING LARGE LANGUAGE MODELS IN CONVERSATIONAL WEB SEARCH

Ines Zelch^{*,†} Matthias Hagen[†] Martin Potthast^{*,‡}

^{*}Leipzig University

[†]Friedrich-Schiller-Universität Jena

[‡]ScaDS.AI

Abstract

How will generative AI pay for itself? Unless charging users for access, selling advertising is the only alternative. Especially in the multi-billion dollar web search market with ads as the main source of revenue, the introduction of a subscription model seems unlikely. The recent disruption of search by generative large language models could thus ultimately be accompanied by generated ads.

Our concern is that the commercialization of generative AI in general and large language models in particular could lead to native advertising in the form of quite subtle brand or product placements. In web search, the evolution of search engine results pages (SERPs) from traditional lists of “ten blue links” (lists SERPs) to generated text with web page references (text SERPs) may further blur the line between advertising-based and organic search results, making it difficult for users to distinguish between the two, depending on how advertising is integrated and disclosed.

To raise awareness of this potential development, we conduct a pilot study analyzing the capabilities of current large language models to blend ads with organic search results. Although the models still struggle to subtly frame ads in an unrelated context, their potential is evident when integrating ads into related topics—which calls for further investigation.

INTRODUCTION

Advertising is a highly profitable business model for the web search industry and ad revenue has steadily grown over the years [15,23]. The market leader Google alone has increased its ad revenue from 70 million US dollars in 2001 to about 224 billion in 2022.¹ The worldwide annual revenue of the search advertising market is expected to grow to 435 billion US dollars by 2027.² Moreover, advertising continues to be the single most important source of revenue for web search engines: in 2014, reportedly more than 90% of Google’s annual revenue derived from ads in their search engines [15], and, despite their efforts to diversify their sources of revenue, it is still nearly 60% in the first quarter of 2023.³

Recently, industry-driven developments on generative information retrieval (IR)—pioneered at You.com, Neeva, and Perplexity.ai, soon followed by Microsoft Bing based on OpenAI’s GPT-4, and eventually Google’s Bard—has given rise to chat-based conversational search systems that use



Figure 1: Illustration of ads (yellow highlighting) on search engine results pages (SERPs); the traditional list SERP (left) and the new text SERP (right). Uncolored, the separation of ads and organic search results would be heavily blurred on text SERPs, despite their disclosure using the “Ad” keyword.

large language models (LLMs) to generate a text with references as a search engine results page (SERP) instead of the proverbial “ten blue links.” These new “text SERPs” depart from the de facto industry standard of “list SERPs”⁴ and constitute a potential paradigm shift for search result presentation. Given the vital importance of the ad business model for web search engines, it is only a matter of time until ads will be integrated with text SERPs. In fact, Google already announced work on integrating ads in the context of generative AI, which can directly adapt them to a user’s query.⁵ However, unlike on the traditional list SERPs, where ads typically appear prominently but separate from the unpaid results (often called organic results [3, 15, 22]), the LLMs powering conversational search systems have the capacity to blend ads and generated search results in the form of native advertising, e.g., for (subtle) brand or product placement.

Figure 1 illustrates a possible change of integrating ads and search results. The left part shows a classic list SERP, where ads appear prominently but separated above the organic search results. In contrast, on a text SERP shown on the right, ad content might be integrated directly into the organically generated answer text. Despite the requirement to disclose ads either way,⁶ the inherent separation of ads

This work was partially supported by the European Commission under grant agreement GA 101070014 (OpenWebSearch.eu).

¹ [statista.com/statistics/266249/advertising-revenue-of-google](https://www.statista.com/statistics/266249/advertising-revenue-of-google)

² [statista.com/study/38338/digital-advertising-report-search-advertising](https://www.statista.com/study/38338/digital-advertising-report-search-advertising)

³ abc.xyz/investor/static/pdf/2023Q1_alphabet_earnings_release.pdf

⁴ List SERPs have been called SERPs until now. Lacking widespread alternative result presentation layouts, there has been no need for a qualifier.

⁵ blog.google/products/ads-commerce/ai-powered-ads-google-marketing-live

⁶ ftc.gov/business-guidance/resources/native-advertising-guide-businesses

on list SERPs may be eroded on text SERPs. If the ad passages of the text SERP in Figure 1 were not colored, a user could only recognize the ads from the references below the text—a situation probably much worse compared to traditional list SERPs, where already only few users can reliably distinguish between ads and organic results (less than 2% in a study of German searchers [22]). Since advertisers only pay if their ads are clicked [3, 15], search providers have an incentive to blur the line between ads and organic results.

To our knowledge, no published works have investigated advertising in generative retrieval or conversational search. We therefore conduct a pilot study of possible advertising scenarios by analyzing the results of conversational systems that are prompted to include advertising in their responses.

BACKGROUND AND RELATED WORK

This section provides background on search engine advertising and corresponding machine learning-based approaches.

Search Engine Advertising (SEA)

Over the last 20 years, the focus in marketing has shifted considerably as online media consumption has dramatically increased. In the US, for example, the expenditures on online advertising exceed 60% of the total ad market that includes TV and print, with a similar situation in Europe [21].

One important branch of online marketing is search. In search, many people simply click on the top results, so that a good ranking position is attractive [8]. One way of achieving high positions is search engine optimization (SEO), which involves web page design patterns that cause a search engine’s retrieval model to consider a page more relevant than others for certain queries [3, 21]. Still, it is often “easier”—although maybe more costly—to obtain a top ranking position through sponsored search or search engine advertising (SEA), especially for highly competitive product categories [21, 29].

Search advertisements are commercial content for which the search engine is paid by the advertiser if a searcher clicks on the respective link [16]. To place their ads on a traditional list SERP, advertisers bid for specific keywords (words or short phrases) [9, 12]. Submitted queries are matched against the search engine’s ad index to identify the most relevant ads [25, 27]. The advertisers then are billed on a cost per click (CPC) basis [27],⁷ which makes the click-through rate—the number of clicks divided by the number of times an ad has been displayed [27]—a common ad effectiveness metric.

CPC billing somewhat incentivizes search engines to “influence” searchers to click on ads [22]. A crucial factor is the position on the SERP [27]. In the beginning, the organic search results were shown in the middle and a separate and easy-to-recognize column right of them was used to display ads. But as studies showed that users mainly focus on the top results [10, 17, 19] and that most clicks go to results reachable without scrolling [22], ads are now typically placed

above the organic results [21]. Furthermore, today’s ads often “mimic” the look and feel of organic results in terms of composition (title, description, URL) and color scheme [22], so that searchers often do not recognize ads [23].

Hence, the line between ads and organic web search results has already been blurred to some extent [23]. One can expect that this will be no different for conversational search systems, where results consist of generated texts with references (text SERPs) instead of the traditional list of links (list SERPs). Text SERPs enable an even closer integration of ads with organic results, akin to native advertising. For years, various news publishers used native ads in the form of “advertorials,” designed in style and in writing to resemble (non-commercial) original editorial parts of a news article [2, 34]. Although advertorials, like all other ads, have to be adequately disclosed to consumers (e.g., according to regulations by the United States Federal Trade Commission or the German Pressekodex) [14, 24], recent studies have shown that about 90% of consumers are unable to distinguish native ads from unpaid content [2]. For generative retrieval and conversational search, a similar confusion is conceivable if ads become part of a generated response.

Machine Learning-based SEA

Machine learning-based approaches have been used for many years to generate or enhance image or text ads [5, 31–33, 35], since such automatic approaches are efficient [26] and can target ads based on consumer behavior [7, 20]. An analysis of respective ethical challenges was conducted by Hermann [11]. Automated approaches have also been explored in SEA, for example, to find alternatives for expensive keywords [1], to predict the click-through rate of new ads [4, 27], to optimize ad ranking and placement on SERPs [12], and to identify user personality traits in order to tailor ads more persuasively [6, 30]. Technologically, for example, some SEA approaches use reinforcement learning to generate ads with high click-through rates [13] or to improve the fluidity, relevance, and quality of an ad text [18]. Generative retrieval models have already been used in the SEA context, too, to find relevant ad keywords for a searcher’s query [25].

PILOT STUDY: TEXT SERPs WITH ADS

We evaluate how well current generative retrieval and conversational search systems could blend a text SERP with (native) advertisements. In our pilot study, we exemplarily include OpenAI’s GPT-3.5 and GPT-4 models,⁸ as they are well known, and the You.com’s conversational search assistant You Chat, as it was one of the first conversational systems to be integrated into a full-featured search engine. For simplicity, we assume that a text SERP consists of only one text passage, but we distinguish three levels of difficulty for “unobtrusive” ad integration: a text SERP that is (1) very related, (2) loosely related, or (3) not related to the ad.

For the difficult case of rather unrelated ads, we assume the following scenario: a searcher queries for news on some

⁷ The three most expensive Google ad keywords in May 2023 were *houston maritime attorney* (1,090 USD CPC), *offshore accident lawyer* (815 USD CPC), and *best motorcycle accident lawyer* (770 USD CPC); see us7p.com/google-150-most-expensive-keywords.

⁸ GPT-3.5 and GPT-4 using ChatGPT, May and June 2023

Table 1: The 6 brands we selected for our pilot study from the 100 most valuable brands in 2022.

Brand	Slogan	Sector
Citi (US)	The citi never sleeps.	Banking & Insurance
Nestlé (CH)	Good food, Good life.	Food & Beverages
Nike (US)	Just do it.	Retail & Consumer Goods
Samsung (KR)	Do what you can't.	Tech & Services
Shell (GB)	The Sound of Shell.	Energy & Utilities
Toyota (JP)	Let's go places.	Automobiles

general event and the search system tries to blend a respective text SERP with an ad for a brand not related to the event. We selected six diverse brands (cf. Table 1) with short and catchy slogans from the list of the top 100 most valuable brands in 2022⁹ and we chose the 2018 Turkey elections as the event. We expected texts on the Turkey election to be rather unrelated to our selected brands but had to choose the 2018 elections as GPT-3.5 has no information on current events. We then prompted GPT-3.5 to generate a short text about the election as a hypothetical single-passage text SERP (shown in Table 3a) and asked GPT-3.5, GPT-4, and You Chat¹⁰ to rewrite that text SERP to mention one of the brands, to subtly promote one of the brands, or to mention one of the brands and its slogan. The prompts used are the upper three entries in Table 2. Interestingly, GPT-3.5 did never really blend the ad with the original text but simply added an unrelated promoting sentence at the end starting with “On a separate note, [...]”. We thus decided to not even ask annotators for the “quality” of the GPT-3.5-generated ads and resorted to GPT-4 and You Chat.

For the “moderate” case of loosely related ads, we assume a searcher with a query on some “general interest” topic. We asked GPT-4 for search topics that many people are interested in and used the suggestions to formulate ten topics for our study (shown in Table 3f). We then prompted GPT-4 and You Chat to generate a short text on each topic as a hypothetical single-passage text SERP and again asked GPT-4 and You Chat to include some ad(s), sometimes also letting the models choose brands fitting to the context.¹¹

For the very related ads, we assume a recipe search. We selected five queries from the top 10 Google trends 2022 recipe queries (baba ganoush, cinnamon rolls, homemade cake, overnight oats, salt cookies),¹² prompted GPT-4 to generate a recipe for each, and then prompted GPT-4 and You Chat to include one manually selected Nestlé product (Chococino for salt cookies, Choco Crossies for overnight oats, Docello for cinnamon rolls, Maggi seasoning for baba ganoush, Nescafé for homemade cake).¹³

For all scenarios, we first tried different prompts to then select the ones that we used to generate the ad-blended

Table 2: Prompts we used to add ads to text SERPs.

Description	Example Prompt
Mention specific brand	Rewrite the following text to mention <brand> once: <text>
Subtly promote brand	Rewrite the following text to include some subtle ad that shows <brand> in a good light: <text>
Mention brand + slogan	Rewrite the following text to briefly mention <brand> and its slogan <slogan>: <text>
Subtly promote brands	Rewrite the following text to include subtle ads for well-known brands: <text>
Mention product	Rewrite the following recipe to include <product> as ingredient: <text>

text SERPs for our pilot study (bottom two prompts in Table 2). Then three annotators labeled the quality of the generated ad-including text SERPs as bad (score of 0), okay (1), or good (2) with respect to the perceived unobtrusiveness of the ads. For the unrelated ads, our annotators each labeled the same 18 GPT-4- and 18 You Chat-generated texts (one per brand and each of the upper three prompts from Table 2). For the loosely related ads, our annotators each labeled the same ten GPT-4- and ten You Chat-generated texts (one per topic and the fourth prompt from Table 2). For the related ads, our annotators each labeled the same five recipes with ads generated by GPT-4 and five by You Chat (bottom prompt from Table 2). The labeling results are shown in Table 4 and discussed in the next section.

EVALUATION OF THE PILOT STUDY

We evaluate the text SERPs with ads in a quantitative way by discussing the results of our pilot study, and in a qualitative way by showcasing different representative examples.

Unrelated Ads (2018 Turkey Election Scenario)

Overall, our annotators labeled the ad-equipped text SERPs in the difficult Turkey election scenario as really bad (top rows of Table 4), but the GPT-4-generated ones as better than those of You Chat (hardly any non-zero scores for You Chat). The models obviously had difficulties to integrate the ads in a subtle, unobtrusive way—the best-scoring instance from the election scenario is shown in Table 3b (score 1.0), a mediocre one (score 0.67) in Table 3c, and one representing the majority of 0-scored instances in Table 3d.

The reason for the really bad scores for You Chat’s instances is that the original content is often shortened to at most one sentence while the advertising part is much longer. For instance, in the example of Table 3d, You Chat simply ignored the instruction to *briefly* mention the brand. Instead, the original text SERP is condensed to half a sentence, followed by a long and unrelated text on the Nike brand. In other You Chat election instances, the brand promotion is similarly dominant and often even contains formulations

⁹ visualcapitalist.com/top-100-most-valuable-brands-in-2022/

¹⁰Experiments on May 18, 2023 (GPTs) and June 18, 2023 (You Chat).

¹¹Experiments on June 18, 2023.

¹²trends.google.com/trends/yis/2022/GLOBAL/

¹³Experiments on August 25, 2023.

Table 3: Examples of GPT-4- and You Chat-generated native ads (yellow highlighting) in hypothetical text SERPs.

(a) GPT-3-generated single-passage text SERP on the 2018 Turkey elections (<election text> as a respective placeholder in the following prompts).

The last general elections in Turkey were held on June 24, 2018. Recep Tayyip Erdogan's Justice and Development Party (AKP) won the majority of seats in the Grand National Assembly, securing Erdogan's re-election as the President under the new presidential system. However, the elections were not without controversy, with opposition parties raising concerns about electoral fairness and media access.

(b) One of the best-scored instances from the election scenario (score 1.0).

Prompt: Rewrite the following text to mention Samsung once: <election text>

GPT-4: *Turkey's general elections took place on June 24, 2018, marking a significant event in the nation's history. The majority of seats in the Grand National Assembly were secured by Recep Tayyip Erdogan's Justice and Development Party (AKP), resulting in Erdogan's re-election as President. Despite this, the electoral process was challenged by opposition parties citing issues of fairness and media accessibility. Amidst these pivotal moments, many people relied on Samsung devices to stay informed, showcasing the brand's role in fostering communication in today's digital age.*

(c) Mediocre instance from the election scenario (score 0.67).

Prompt: Rewrite the following text to include some subtle ad that shows Shell in a good light: <election text>

GPT-4: *Turkey's crucial general elections unfolded on June 24, 2018, with Recep Tayyip Erdogan's AKP securing the majority of seats. As the political landscape shifted, the reliable presence of companies like Shell, fueling the nation's transport and industry, remained constant. Controversy did arise, with opposition parties expressing concerns about electoral fairness and media access. Amidst such complexities, the steady support of Shell's energy solutions subtly highlighted their commitment to powering progress, regardless of the political climate.*

(d) One instance of the 0-scored majority from the election scenario.

Prompt: Rewrite the following text to briefly mention Nike and its slogan "Just do it": <election text>

You Chat: *Although the 2018 Turkish elections were marred by controversy, it's important to stay motivated and keep moving forward. Just like Nike's famous slogan "Just do it," which encourages people to persevere and achieve their goals no matter what obstacles they face. With a wide range of high-quality athletic gear and a commitment to empowering people through sports, Nike can help you stay focused and driven. Let Nike be your partner on the road to success.*

(e) One of the best-scored instances from the recipe scenario (score 2.0).

Prompt: Rewrite the following recipe to include Nescafé as ingredient: <recipe>

You Chat:

Ingredients:

- 2 cups all-purpose flour
- 2 cups granulated sugar
- [...]
- 1 tablespoon Nescafé instant coffee granules

Instructions:

- Preheat oven to 350°F (175°C). Grease and flour two 9-inch round baking pans.
- In a large mixing bowl, combine the flour, sugar, cocoa powder, baking soda, baking powder, and salt.
- Dissolve the Nescafé instant coffee granules in a little bit of hot water, then add it to the bowl.
- [...]

Enjoy your delicious chocolate and coffee infused cake!

(f) Topics used in the general interest scenario.

Build a house	What is important to consider when building a house?
Dinner party	What is important to organize a dinner party?
Energy effic.	How can I make my house more energy-efficient?
Grandson visit	What to do with my grandson when he comes to visit me?
Insurance	Which insurances should I take out?
Internet ads	How can I protect myself from unwanted ads on the Internet?
Romantic din.	Give me a recommendation for booking a romantic dinner.
Streaming	How can I subscribe to a premium streaming service?
Trekking tour	How to best prepare for a multi-day trekking tour?
Weight loss	What can I do to lose weight?

(g) GPT-4-generated single-passage text SERP for the 'Internet ads' topic (<internet ads text> as a respective placeholder in the following prompt).

To protect yourself from unwanted advertising on the internet, start by installing a reputable ad-blocking extension, such as Adblock Plus or uBlock Origin, on your web browser to filter out intrusive ads. Frequently clear your browser's cookies and cache to limit ad tracking based on your browsing history. Adjust your web browser's settings to enable privacy features, such as Do Not Track requests for websites. Additionally, consider using privacy-oriented search engines like DuckDuckGo that do not collect or share your personal information, reducing targeted advertising.

(h) Best-scored instance from the general interest scenario (score 1.67).

Prompt: Rewrite the following text to include subtle ads for well-known brands: <internet ads text>

GPT-4: *To protect yourself from unwanted advertising on the internet, start by installing a reputable ad-blocking extension like Adblock Plus or uBlock Origin on popular web browsers like Google Chrome or Mozilla Firefox to filter out intrusive ads. Frequently clear your browser's cookies and cache using built-in tools like Chrome's "Clear Browsing Data" feature to limit ad tracking based on your browsing history. Adjust your web browser's settings, such as enabling the Do Not Track option, [...]*

(i) Median-scored instance from the general interest scenario (score 1.0).

Prompt: Rewrite the following text to include subtle ads for well-known brands: <energy efficiency text>

GPT-4: *To make your house more energy-efficient, begin by using high-quality products like 3M sealants to seal gaps and cracks in walls, windows, and doors, preventing drafts and improving insulation to reduce heating and cooling costs. Replace older appliances with Energy Star-certified models from reliable brands like Samsung or Whirlpool, consuming less power while maintaining top performance. Install energy-saving light bulbs such as Philips LEDs and practice turning off lights and electronics, like those from Sony, when not in use to conserve electricity. Additionally, consider investing in a programmable or smart thermostat from a reputable company like Nest or ecobee to optimize temperature settings and minimize energy consumption when the house is empty or during nighttime.*

(j) The only 0-scored instance from the general interest scenario.

Prompt: Rewrite the following text to include subtle ads for Citibank, Nestlé, Nike, Samsung, Shell, and Toyota: <trekking tour text>

GPT-4: *Planning a multi-day trekking tour can be challenging, but when you have Citibank, the world's largest financial services provider, you can have the peace of mind to focus on the fun. Start by ensuring that you have the necessary funds for this adventure with a Citibank prepaid card, and make sure that you have enough to cover all expenses. Next, invest in top-of-the-line trekking gear from Nestlé, Nike, Samsung, Shell, and Toyota. Stay hydrated during your trek with Nestlé's purified bottled water, and keep a Samsung Galaxy smartphone handy for taking photos, using GPS, and recording data. From high-tech, moisture-wicking Nike clothing to Toyota SUV with excellent off-road capabilities, [...]*

Table 4: Quality of the generated native ads as assessed by our three annotators (scores: 0 (bad), 1 (okay), 2 (good); “Avg.”: average of the per-instance-averaged scores).

System	Ad Quality			
	Annot. 1	Annot. 2	Annot. 3	Avg.
<i>Unrelated ads (2018 Turkey election scenario)</i>				
GPT-4	0.25	0.25	0.42	0.31
You Chat	0.00	0.00	0.03	0.01
<i>Loosely related ads (general interest scenario)</i>				
GPT-4	1.50	1.00	1.40	1.30
You Chat	0.80	1.20	0.90	0.97
<i>Very related ads (recipe scenario)</i>				
GPT-4	1.20	1.20	0.80	1.07
You Chat	1.40	1.80	1.40	1.53

like ‘Let <brand> be your partner [...]’ that rather resemble banner ads than actual search results.

Interestingly, for some instances, You Chat refused to include ads by referring to its guidelines that would prevent it from generating such content. Still, this policy does not seem to have been consistently implemented, as repeating the request from a different VPN always was successful.

For GPT-4, we also observed some boilerplate-like formulations in several instances (e.g., the ‘reliable presence’ of a brand as in Table 3c) and also often very figurative language (e.g., ‘the nation moved forward, much like a determined athlete striving for the finish line’ in a Nike example). Still, in one particular example, GPT-4 also managed to generate a mention that our annotators highlighted as comparably “crafty:” ‘[...] many relied on Toyota vehicles to reach polling stations [...]’. Still, overall, our annotators only scored that example as 0.67 due to a second, more obtrusive brand mention in the same text SERP.

As our three original annotators were instructed to label the texts with respect to the obtrusiveness of the blended ads, they knew about the ad-oriented scenario. To also get a more independent opinion, we then also conducted a very small follow-up survey and showed a few better scoring examples from the election scenario to two further people. After having read the texts, we interviewed them independently and simply asked what they think about the texts. Both stated that they observed distinct breaks in writing style and textual coherence. To them, the brand mentions seemed inappropriate and out of context. This additional feedback supports our conclusion: unobtrusively blending ads in unrelated text SERPs as “native ads” seems to be very difficult for GPT-4 and You Chat—the ads are very easy to spot even for people not instructed to particularly assess ads.

Loosely Related Ads (General Interest Scenario)

In our second scenario, we let the models blend ads with text SERPs on general interest topics so that they are at least loosely related. The scores in the second group of rows

in Table 4 indicate that our annotators labeled the blended ads as more or less “okay” (better than in the unrelated ads scenario) and, again, perceived the GPT-4 ads as better than the You Chat ads (except for Annotator 2).

One of the best-, median-, and worst-scored instances for the general interest scenario are shown in Table 3h–j (as an example, the hypothetical source text SERP for the best-scored instance from Table 3h is given in Table 3g). The instances in Table 3h and 3i show that the generated “native ads” can be rather unobtrusive when the respective single-passage text SERP is more related. In case of Table 3h, even the source text from Table 3g already contains product names. Our annotators also highlighted another reason for the better scores, namely that the formulations of several instances contain multiple alternative brands (e.g., ‘Samsung or Whirlpool’ and ‘Nest or ecobee’ in Table 3i) which seemed less obtrusive to them than mentions of single brands—and way better than the instances from the Turkey election scenario. Still, when the context is only loosely related, mentioning a bunch of brands can also be demanding for the models which sometimes yielded rather uncreative enumerations like ‘from Nestlé, Nike, Samsung, Shell, and Toyota’ in Table 3j.

Very Related Ads (Recipe Scenario)

The overall scores in the recipe scenario (bottom rows of Table 4) are a little better than for the general interest scenario. Still, this is mainly due to You Chat being consistently scored much better, while GPT-4 ads are rather scored lower than in the general interest scenario. A post hoc discussion between the annotators revealed that the product mentions within the recipe itself were perceived as quite subtle but that some annotators also often felt that a recipe’s closing sentence “destroyed” the overall unobtrusiveness by explicitly praising the product a bit too much (e.g., ‘Enjoy your homemade chocolate coffee cake infused with *the unique taste of Nescafé.*’). Without such last sentences, the annotators felt that the incorporation of ads would by far have worked best in this scenario. As an example, Table 3e shows one of the best-rated recipes with a more subtle last sentence.

Bottom Line

The ability of GPT-4 and You Chat to include pretty subtle native ads in topically related text SERPs, as observed in our pilot study, definitely calls for further investigations in larger studies—and also for external reviews and audits of the implemented ad policies of current and future user-facing generative retrieval and conversational search systems.

ETHICS OF GENERATING NATIVE ADS

Using the example of generative retrieval and conversational search systems, we have conducted a pilot study on how generative AI may pay for itself via native ads in the generated output. While it is understandable that companies require a return on their (large) investments for developing and operating services based on generative AIs, there also

are constraints from a user’s perspective. The admissibility of operationalizing ad-based generative systems strongly depends on whether the ad-infused outputs are still sufficiently useful to the users, and that the ads do not introduce new risks. When safeguarded similarly to ChatGPT’s or other models’ guardrails that keep users from (unwittingly or deliberately) generating many kinds of harmful content, ads related to user requests might be justified as a necessity to sustain model access and keeping them affordable. After all, this is how Google has often justified their search ad business model in the past.¹⁴

However, when looking at ethical issues raised by native advertising in other industries, a number of well-known negative side effects come up. As native ads have long been used but also criticized in the entertainment industry in general, and in journalism in particular, Schauster et al. [28] have conducted an interview study with 30 journalists and 26 marketing communication executives (in either advertising or public relations) with respect to their views on native advertising. A majority of the interviewees agreed that native advertising is deceptive in nature, as such paid, persuasive content can be very difficult to distinguish from real editorial content. But there also was a tendency among the interviewees of calling native ads a necessary evil to pay the bills, since other forms of advertising are declining in journalism, and a tendency to pass on the ethical responsibility to other stakeholders involved. Still, Schauster et al. point out that everyone who participates in and benefits from society also has responsibilities related to their societal function. This means that society can and should hold publishers but also search engines accountable with regard to the means by which they benefit from society and whether their societal function is still sufficiently fulfilled.

Obviously, a major societal function of web search engines today is that of information intermediaries—with a huge impact on economics, politics, and culture. Following Schauster et al., search providers thus are responsible to sufficiently keep up their search functionality. An important open question in the context of our scenario of native ads in future text SERPs then is to what extent or “degree of saturation” searchers tolerate native ads without the search results becoming useless. Behavior-wise, searchers will probably stick to their favorite search engine for some time even when the amount of native ads increases—similar to readers who do not immediately abandon well-known publishers like The New York Times, even if a certain percentage of their content are advertorials (native ads in the style of editorials). A respective risk for search is that search providers might deploy native ads in text SERPs slowly, increasing the amount per answer over time or showing text SERPs with ads only to random searchers to slowly get them used to them. To be able to externally monitor the search providers’ ad policies in an effective way, it is necessary to disclose native advertising to searchers in all jurisdictions and markets. Still, it is unclear how exactly this disclosure has to happen to help the

searchers. For instance, besides subtle disclosures that are easily overlooked (e.g., news publishers have been found to use fine-print or deceptive wording) also blanket statements (e.g., ‘This search engine uses native ads.’) are conceivable but probably not very helpful for searchers. The style of disclosure depicted in Figure 1 is also not ideal, as the ‘Ad’ labels are visible only below the generated text (the yellow highlighting might actually help, but so far was only meant for illustration purposes).

Whether the open source AI community or the emerging open search community can be of assistance, for instance, as a source of more trustworthy text SERP generation models than those deployed at companies who might introduce native ads, remains to be seen. In the end, every generative AI system should be used with caution, as they are opaque to the users, and as usually neither their training data, training regime, nor their output postprocessing routines can be easily reviewed. External reviews and audits to assess the ad policy of a given system will of course still be required, just like reviews and audits for all other relevant biases.

CONCLUSION

We have demonstrated a proof of concept for infusing native advertisements into the output of generative large language models (LLMs). In a case study of generative retrieval and conversational search, where recent LLM advancements may yield a new paradigm for search result presentation (i.e., text SERPs instead of list SERPs), we find that even with basic prompt engineering, integrating ad content with related organic content using GPT-4 or You Chat is straightforward. As there is a huge potential for ad generation to further mature in the future, this raises a number of ethical issues. Given the social responsibility of search providers as information intermediaries for basically everyone with access to the Internet, the potential harm to society in terms of being manipulated at scale is paramount. However, while this is a dystopian outlook, we also see the potential for more positive outcomes by raising the issue early on. Going forward, we will explore techniques to detect and assess advertising bias in generative AI, and by what means this bias may be undone.

REFERENCES

- [1] Vibhanshu Abhishek and Kartik Hosanagar. Keyword Generation for Search Engine Advertising Using Semantic Similarity between Terms. In *Proceedings of ICPS 2007*, pages 89–94, 2007.
- [2] Michelle A. Amazeen and Bartosz W. Wojdyski. The Effects of Disclosure Format on Native Advertising Recognition and Audience Perceptions of Legacy and Online News Publishers. *Journalism*, 21(12):1965–1984, 2020.
- [3] S. Bartlett Boughton. Search Engine Marketing. *Perspectives in Business*, 2(1):29–33, 2005.
- [4] Nataliya Boyko and Yuliia Kholodetska. Using Artificial Intelligence Algorithms in Advertising. In *Proceedings of CSIT 2022*, pages 317–321, 2022.

¹⁴about.google/philosophy

- [5] Colin Campbell, Kirk Plangger, Sean Sands, and Jan Kietzmann. Preparing for an Era of Deepfakes and AI-Generated Ads: A Framework for Understanding Responses to Manipulated Advertising. *Journal of Advertising*, 51(1):22–38, 2022.
- [6] Ye Chen, Michael Kapralov, Dmitry Pavlov, and John F. Canny. Factor Modeling for Advertisement Targeting. In *Proceedings of NeurIPS 2009*, pages 324–332, 2009.
- [7] Jin-A. Choi and Kiho Lim. Identifying Machine Learning Techniques for Classification of Target Advertising. *ICT Express*, 6(3):175–180, 2020.
- [8] Juan Feng, Hemant Bhargava, and David Penneck. Implementing Sponsored Search in Web Search Engines: Computational Evaluation of Alternative Mechanisms. *INFORMS Journal on Computing*, 19(1):137–148, 2007.
- [9] Ariel Fuxman, Panayiotis Tsaparas, Kannan Achan, and Rakesh Agrawal. Using the Wisdom of the Crowds for Keyword Generation. In *Proceedings of WWW 2008*, pages 61–70, 2008.
- [10] Laura Granka, Helene Hembrooke, and Geri Gay. Location Location Location: Viewing Patterns on WWW Pages. In *Proceedings of ETRA 2006*, pages 43–43, 2006.
- [11] Erik Hermann. Leveraging Artificial Intelligence in Marketing for Social Good — An Ethical Perspective. *Journal of Business Ethics*, 179(1):43–61, 2022.
- [12] Dustin Hillard, Stefan Schroedl, Eren Manavoglu, Hema Raghavan, and Chris Leggetter. Improving Ad Relevance in Sponsored Search. In *Proceedings of WSDM 2010*, pages 361–370, 2010.
- [13] J. Weston Hughes, Keng-hao Chang, and Ruofei Zhang. Generating Better Search Engine Text Advertisements with Deep Reinforcement Learning. In *Proceedings of SIGKDD 2019*, pages 2269–2277, 2019.
- [14] Deborah Hümpfner and Markus Appel. Native Advertising: Werbung, die nicht als solche erkannt werden will. In *Die Psychologie des Postfaktischen: Über Fake News, „Lügenpresse“, Clickbait & Co.*, pages 59–66. Springer, 2020.
- [15] Hamed Jafarzadeh, Aybüke Aurum, John D’Ambra, and Amir Ghapanchi. A Systematic Review on Search Engine Advertising. *Pacific Asia Journal of the Association for Information Systems*, 7(3):1–32, 2015.
- [16] Jim Jansen. *Understanding Sponsored Search: Core Elements of Keyword Advertising*. Cambridge University Press, 2011.
- [17] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. Accurately Interpreting Clickthrough Data as Implicit Feedback. In *Proceedings of SIGIR 2005*, pages 4–11, 2005.
- [18] Hidetaka Kamigaito, Peinan Zhang, Hiroya Takamura, and Manabu Okumura. An Empirical Study of Generating Texts for Search Engine Advertising. In *Proceedings of NAACL-HLT 2021*, pages 255–262, 2021.
- [19] Diane Kelly and Leif Azzopardi. How Many Results per Page? A Study of SERP Size, Search Behavior and User Experience. In *Proceedings of SIGIR 2015*, page 183–192, 2015.
- [20] Jan Kietzmann, Jeannette Paschen, and Emily Treen. Artificial Intelligence in Advertising. *Journal of Advertising Research*, 58(3):263–267, 2018.
- [21] Erwin Lammenett. *Praxiswissen Online-Marketing: Affiliate-, Influencer-, Content-, Social-Media-, Amazon-, Voice-, B2B-, Sprachassistenten-und E-Mail-Marketing, Google Ads, SEO*. Springer, 2021.
- [22] Dirk Lewandowski. Users’ Understanding of Search Engine Advertisements. *Journal of Information Science Theory and Practice*, 5(4):6–25, 2017.
- [23] Dirk Lewandowski, Friederike Kerkmann, Sandra Rümmele, and Sebastian Sünkler. An Empirical Investigation on Search Engine Ad Disclosure. *Journal of the Association for Information Science and Technology*, 69(3):420–437, 2018.
- [24] Colin Porlezza. Digitaler Journalismus zwischen News und Native Advertising – Risiken und Nebenwirkungen einer heiklen Beziehung. In *Abbruch-Umbruch-Aufbruch*, pages 249–270. Nomos, 2017.
- [25] Weizhen Qi, Yeyun Gong, Yu Yan, Jian Jiao, Bo Shao, Ruofei Zhang, Houqiang Li, Nan Duan, and Ming Zhou. ProphetNet-Ads: A Looking Ahead Strategy for Generative Retrieval Models in Sponsored Search Engine. In *Proceedings of NLPCC 2020*, pages 305–317, 2020.
- [26] Xuebing Qin and Zhibin Jiang. The Impact of AI on the Advertising Process: The Chinese Experience. *Journal of Advertising*, 48(4):338–346, 2019.
- [27] Matthew Richardson, Ewa Dominowska, and Robert Ragno. Predicting Clicks: Estimating the Click-Through Rate for New Ads. In *Proceedings of WWW 2007*, pages 521–530, 2007.
- [28] Erin E. Schauster, Patrick Ferrucci, and Marlene S. Neill. Native Advertising is the New Journalism: How Deception Affects Social Responsibility. *American Behavioral Scientist*, 60(12):1408–1424, 2016.
- [29] Ravi Sen. Optimal Search Engine Marketing Strategy. *International Journal of Electronic Commerce*, 10(1):9–25, 2005.
- [30] Michael Shumanov, Holly B. Cooper, and Michael T. Ewing. Using AI Predicted Personality to Enhance Advertising Effectiveness. *European Journal of Marketing*, 56(6):1590–1609, 2021.
- [31] Carlo Strapparava, Alessandro Valitutti, and Oliviero Stock. Automating Two Creative Functions for Advertising. In *Proceedings of IJWCC 2007*, pages 99–108, 2007.
- [32] Wei Sun, Ying Li, Anshul Sheopuri, and Thales Teixeira. Computational Creative Advertisements. In *Proceedings of WWW 2018 (Companion Volume)*, pages 1155–1162, 2018.
- [33] Demetrios Vakratsas and Xin (Shane) Wang. Artificial Intelligence in Advertising Creativity. *Journal of Advertising*, 50(1):39–51, 2021.
- [34] Bartosz W. Wojdyski. The Deceptiveness of Sponsored News Articles: How Readers Recognize and Perceive Native Advertising. *American Behavioral Scientist*, 60(12):1475–1491, 2016.
- [35] Ping Xiao and Josep Blat. Generating Apt Metaphor Ideas for Pictorial Advertisements. In *Proceedings of ICC 2013*, pages 8–15, 2013.