

Abusing AI for Advertising

Will generative artificial intelligence accelerate the growth of made-for-advertising sites?

THE USE OF made-for-advertising (MFA) sites, websites created solely to generate digital ad impressions without regard to the content's quality, legitimacy, or the end user's experience, have become a commonplace—and often unwanted—element in the Internet advertising ecosystem. These sites, often referred to as ‘clickbait sites,’ offer inexpensive advertising impressions that rarely translate into ad click-throughs that generate actual sales or other meaningful business results for advertisers, and divert revenue away from legitimate content publishers.

Creators of MFA sites often steal content from other sites without attribution or compensation, and use it to lure unsuspecting Internet users to their sites, often via paid search links, or via links on social media platforms. When they arrive at the MFA site, which often is littered with a very high ratio of ads to content, most users will navigate away without taking any action. However, because they have visited the site, the ads will have been “viewed,” thereby counting as impressions, and the site owner generates revenue.

Amid the growing availability of generative artificial intelligence (AI) tools, which allow users to quickly generate or steal content without attribution or payment at scale and little manual effort on the part of a user, the number and breadth of MFA sites may increase exponentially. With the ad industry's relatively weak appetite for labeling MFA sites as fraud and rooting them out, these sites likely will remain a ubiquitous presence on the Internet.

MFA Sites Proliferate

Open programmatic advertising is estimated to be an \$88-billion global market, and MFA websites represent 21% of impressions and comprise 15% of ad spending, according to the Association of National Advertisers (ANA). The organization conducted its first programmatic media supply-chain trans-



parency study in June 2023 to examine challenges in advertising on the open Web, and found significant waste in on-line ad spending, with a high amount of ad spending going toward MFA sites.

The issue with MFA sites is that they are designed solely to generate inexpensive views from Web users who click on a seemingly legitimate source of information about a topic. However, because the content is usually poorly written, incomplete, filled with half-truths or outright incorrect information, users will generally not linger on the page, and generally do not click on the ads.

Made-for-advertising sites generate inexpensive views from visitors who click on a seemingly legitimate source of information.

However, simply because the user visits the page, the publisher of the MFA site can claim that real users have viewed its ads, thus generating a payment.

In addition, because MFA sites feature a higher-than-average number of ad units, advertisers wind up paying for views and accidental clicks that rarely result in any valuable business metric, such as down-funnel sales lift. Some MFA sites also will stack ads on top of one another, meaning that advertisers will pay for impressions that are not actually viewable.

Technically, MFA does not rise to the level of ad fraud, which typically includes tactics such as artificially inflating statistics like impressions, clicks, and conversion data through illegal means, and therefore is not illegal. However, when an MFA site lifts content from legitimate publishers, such as a newspaper or magazine's site, without attribution or payment, that is a form of fraud. The availability of generative AI tools make it far easier for MFA site creators to employ such tactics.

The Role of Generative AI in MFA

In the span of just a few months, a plethora of consumer-focused generative AI tools, such as OpenAI's Chat-

GPT or Google's Bard, that make it extremely efficient and easy for bad actors to create MFA sites at scale, have been released. Generative AI tools can be used to grab and rewrite content from legitimate sites without any regard to copyright or ownership. Furthermore, generative AI models learn from the data they are trained on, and if that training data contains biases or misinformation, the models may incorporate and amplify that content, resulting in 'hallucinations' that perpetuate false information, which can then be used to populate MFA sites.

In fact, all indications point to the continued growth of MFAs in the coming years, likely aided by easy access to generative AI tools, which not only allow rapid content generation, but provide the ability to mimic 'good' content at scale, with low effort and relatively low direct costs.

This has implications for advertisers, particularly those selling sensitive goods or services, often in healthcare, financial services, and pharmaceuticals. Because generative AI can be used to quickly populate dozens or hundreds of MFA sites with seemingly legitimate content, supply-side platforms (SSPs), which are used by publishers to connect their inventory to ad exchanges, may inadvertently add MFA sites to their platform.

SSPs are supposed to verify the sites on their platform are using content that is properly acquired, meets industry standards for viewability, and aligns

with the content goals of advertisers (for example, is free of inappropriate, inaccurate content or claims, or other objectionable content), but the scale of MFA proliferation can result in advertisements for products or services appearing on these disreputable sites.

The Future of MFAs and Generative AI

Using generative AI tools may make it even harder for brands to avoid MFA sites or domains, given the current widespread proliferation of clickbait sites. "These ad campaigns are at such scale, you don't have the ability to go in and check the content on every little thing," says Tom Hespos, principal and chief media officer of Holtsville, NY-based Abydos Media. Hespos adds that simply trying to exclude sites based on the quality of their content may be problematic, because "[legitimate] content is different for everybody who uses the Internet."

That is why the ANA thus far has decided to eschew taking direct actions against MFA site proliferators, and instead offered five recommendations for advertisers, including a suggestion to "demand that 'Made for Advertising' websites be excluded from a media buy unless they are specifically wanted or needed."

"Most of the news writing is aimed at a sixth-grade reading level," Hespos says. "It's not going to take [generative] AI very long to figure out how to write at that level and do it convincingly."

Perhaps even more worrisome to publishers is the advertising industry's reluctance to stopping MFA, because money is still being made for ad tech vendors such as SSPs and demand-side platforms (DSPs), which allow the buyers of digital ad inventories to manage multiple ad exchanges via a single interface. Further, MFAs still deliver on their brand promise of delivering real eyeballs at lower costs, even if they are unable to follow through on delivering business results.

Advertising Industry Structure Supports the Status Quo on MFA

Given that large global brands use these SSPs to buy ads, it is clear that household names have been exposed to MFA sites. "The folks who know about it and are already granularly tracking their spend probably aren't the ones in trouble here," says Rocky Moss, CEO of DeepSee, a Draper, UT-based media audit provider that provides domain risk scores. "It is the ones who pump the real money through the ad ecosystem, like the large brand marketers" who don't feel they have the time to verify every SSP or site where their ads are placed.

Rooting out MFA sites will require leveraging new technology solutions, because traditional contextual solutions rely heavily on keywords and topics, rather than assessing the technical components of a publisher and its pages. As a result, MFA sites can trick less-sophisticated keyword and contextual solutions into thinking a site or domain

Technology Policy

Trust, Trustworthiness Essential for AI Systems

ACM's global Technology Policy Council (TPC) recently released "TechBrief: Trusted AI," (<https://bit.ly/3I5sWSO>), which highlights a key challenge to artificial intelligence (AI), that trustworthiness mechanisms and measures being advanced in AI regulations and standards may not actually increase trust.

"So much of the public conversation about regulating AI systems has focused on issues including accuracy or transparency, but there is much less discussion about how the public comes to view an AI system

as 'trustworthy,'" observed John T. Richards, a Distinguished Research Scientist at IBM in the U.S. and co-lead author of the piece. "We found that the public's perspective on what makes AI trustworthy will often diverge from the perspective of technologists and policy makers. We hope this TechBrief begins a conversation that will encourage industry leaders and policymakers to put the issue of trustworthiness front and center."

"As AI is becoming pervasive, more and more institutions are using it," added Bran Knowles, a

professor at the U.K.'s Lancaster University and also co-lead author of "TechBrief: Trusted AI." Said Knowles, "The danger is that a lack of public trust of AI may not only impact the acceptance of these new technologies, but might also erode trust in the institutions that are using it. For these reasons, there is an urgent need for an examination of how public trust is developed around AI technologies."

Said Stuart Shapiro, chair of the TPC's TechBriefs Committee, advances in AI "have been so remarkable and so rapidly

disseminated, it is natural that the public views AI as mysterious and that many have concerns." Shapiro said the new brief "is designed to bring important issues to light which we hope will ultimately lead to greater public confidence in AI-based systems."

The TPC sets the agenda for ACM's global policy activities and serves as the central convening point for ACM's interactions with government organizations, the computing community, and the public in all matters of public policy related to computing and information technology.

is relevant to the advertiser's campaign goals, which can result in advertisers inadvertently purchasing ads on less-than-reputable MFA sites.

Mostly, "there's a laziness and an unwillingness to actually be at the gate, checking [sites] before they're allowed to sell on the marketplace," says Matthew Prohaska, CEO and principal of Prohaska Consulting, who previously served as programmatic advertising director for *The New York Times*. "Part of the scam is using keywords embedded in the page to trigger algorithms used by buyers. No one's checking the actual publisher to [verify] how credible they are."

Ending MFAs Will Rely on Shifts in Measurement, Expectations

That said, most large advertisers do not want to have their brands associated with clickbait content. Having ads appear on MFA sites may help marketing directors meet certain campaign cost key performance indicators (KPIs) such as cost per thousand impressions (CPM), but it certainly doesn't drive real-world business outcomes. CPM, which is an acronym for cost per *mille*, is a common advertising metric that allows for the comparison of the cost to reach potential buyers, but only focuses on impressions, or whether an ad has been viewed by an Internet user, not whether the ad has been effective in driving click-throughs that indicate a desire to seek out additional information or a propensity to purchase a product or service.

However, according to advertising industry veterans, the lure of a relatively inexpensive CPM is often too tempting to ignore. "Few advertisers today, unfortunately, have both the intelligence and the courage to change the rules and tell their buying partners, 'hey, it's ok if you show me a media plan with a \$9 CPM, as opposed to a \$6 CPM, because the \$9 CPM is with quality audiences that performed downstream,'" Prohaska says, adding that supply-side platforms and demand-side platforms share much of the blame for not cracking down on MFA sites because they are still profiting from the ad arbitrage model, even if the advertisers are not getting real value from the placements.

"My take on these sites is that they

MFAs still deliver on their brand promise of delivering real eyeballs at lower costs, even if they are unable to follow through on delivering business results.

actually bring real people to them," Moss says. "That's one of the things that allows them to exist; they wouldn't exist if they didn't bring real people to them. If I had my way, we'd probably move away from the label of MFA and just say, 'low impact, wasteful inventory,' because that's what it is at the end of the day."

That is one reason why other ad industry veterans who asked not to be identified for this article also cited a need for a flight to quality (results) over quantity (impressions) when it comes to shifting advertising KPIs. The ANA study found that MFA impressions had a lower CPM than the average, given they represented 21% of impressions but only 15% of spend. Exclusion of these lower CPM MFA impressions should, therefore, increase the CPM on overall media spend, but that would require everyone from chief marketing and advertising officers on down to accept that current ad metrics are largely superficial, and that reaching actual engaged buyers requires more than often is presented in media plans.


Steps for Identifying, Mitigating MFA Effects

The tide may be shifting against the growth of MFA, at least among large advertising companies. In August 2023, GroupM, WPP's media investment group, announced the introduction of new protections against MFA sites and domains that leverage advanced detection and domain tracking technologies provided through integration of Jounce Media's MFA tracking technology. This was likely a response to not only the

widespread proliferation of MFA sites found on advertising SSPs, but also to multi-seller private marketplaces, also known as auction packages. News reports indicated that in the second quarter of 2023, programmatic consultancy Jounce Media found that 106 of 136 multi-seller private marketplaces included MFA inventory, with nearly a quarter of those marketplaces having allocated more than 25% of their budget to MFA inventory.

Brands and their agencies also can use more advanced pre-bid filtering to complement their advertising supply-path optimization (SPO) efforts, which could involve choosing an SSP that actively vets its inventory of domains, or through taking a more manual approach, picking specific publishers or specific publisher-SSP pairings that offer the best access and ROI.

"The best thing you can do is use a tight targeting list of maybe 1,000 sites, especially if you have something that is highly sensitive," Moss says. "I would recommend using a tightly curated list of sites on which you would ever want your ads to show up on."

"The only way to remove [MFA], though, is to buy everything through private marketplaces," Prohaska contends. "If you remove open auction, you will know with whom you're dealing every single time, because it's a human being at the other end." 

Further Reading

ANA Provides "First Look" at In-depth Programmatic Media Transparency Study, ANA.net, <https://www.ana.net/content/show/id/pr-2023-06-programmaticstudy>

GroupM Introduces New Protections Against Made For Advertising Domains, August 21, 2023, <https://bit.ly/3EgFC7u>

WTF are made-for-advertising sites (MFAs), Digiday, <https://digiday.com/marketing/wtf-are-made-for-advertising-sites-mf-as/>

Fou, A. MFA - Made for Ad Fraud - sites Load Bonkers Numbers of Ads, *Trackers*, April 4, 2023, <https://www.youtube.com/watch?v=9YzP4ssn9k0>

The State of Sustainable Advertising Report, Q2 2023, *Scope3*, July 2023, <https://bit.ly/45KMMNO>

Keith Kirkpatrick is a New York, NY, USA-based Research Director with The Futurum Group.

© 2024 ACM 0001-0782/24/4