

**Forgetful Advertising: Imagining a More Responsible
Digital Ad System**

Chand Rajendra-Nicolucci and Ethan Zuckerman[†]

As Silicon Valley giants sketch their preferred future for digital advertising, an infrastructure with significant implications for life online and offline, there are startlingly few alternatives to their vision. In response, we propose "forgetful advertising", a vision for responsible digital advertising structured around a single design choice: avoiding the storage of behavioral data. Forgetful advertising can still target ads using information like geography, intent, context, and whatever else can be gleaned from a single interaction between a user and a website, but it cannot remember any previous interactions to inform its targeting. We believe our proposal can make digital advertising compatible with the values of human agency and privacy and offer it as a bottom-up solution for organizations that find existing digital advertising systems inconsistent with their values.

Introduction.....	422
I. Responsible Digital Advertising	426
A. Understanding Infrastructures.....	426
B. A DPI for Advertising	428
C. When is advertising good? When is targeting good?	431
D. Who is this system for? What is a values-led organization and why might they find this system useful?	433
II. Forgetful advertising	434

[†] Chand Rajendra-Nicolucci, Research Fellow, Initiative for Digital Public Infrastructure, University of Massachusetts Amherst. Ethan Zuckerman, Professor of Public Policy, Communication, and Information & Director, Initiative for Digital Public Infrastructure, University of Massachusetts Amherst.

A.	The current digital advertising system	434
B.	Existing reform proposals.....	438
C.	Forgetful Advertising.....	441
	Conclusion.....	447

Introduction

Critiques of “surveillance capitalism” are prompting a restructuring of the digital advertising economy. Leading web browsers and mobile phone makers have begun blocking third-party cookies and advertising identifiers—popular methods for tracking user behavior online—with consequences for companies that depend on digital advertising.¹

This reconsideration of surveillant models echoes earlier debates over advertising, recalling moments where society has prohibited certain types of advertising or examined the lines between advertising and propaganda.

In the early 20th century, Paris was overrun by posters, an innovative new form of commercial art which quickly spread to blanket the walls of the city. An anti-poster movement arose that called for making Paris “more beautiful—materially and morally.”² The movement successfully pushed for taxes on

¹ Mitchell Clark, *Google promises it won't just keep tracking you after replacing cookies*, THE VERGE (March 3, 2021), <https://www.theverge.com/2021/3/3/22310332/google-privacy-replacing-third-party-cookies-privacy-sandbox>; Sean Hollister, *Firefox's latest version blocks third-party trackers by default for everyone*, THE VERGE (Sept. 3, 2019), <https://www.theverge.com/2019/9/3/20848629/firefox-69-block-third-party-tracker-default-enhanced-tracking-protection-android-windows-mac-os>; Nick Statt, *Apple updates Safari's anti-tracking tech with full third-party cookie blocking*, THE VERGE (March 24, 2020), <https://www.theverge.com/2020/3/24/21192830/apple-safari-intelligent-tracking-privacy-full-third-party-cookie-blocking>; Nicolás Rivero, *Apple's IDFA privacy update threatens to upend the ad industry*, QUARTZ (April 26, 2021), <https://qz.com/1997873/apples-idfa-privacy-update-threatens-to-upend-the-ad-industry/>; Kif Leswing, *Facebook says Apple iOS privacy change will result in \$10 billion revenue hit this year*, CNBC (Feb. 2, 2022), <https://www.cnbc.com/2022/02/02/facebook-says-apple-ios-privacy-change-will-cost-10-billion-this-year.html>.

² TIM WU, THE ATTENTION MERCHANTS 22 (2016).

posters and restrictions on where posters could be placed—one of the first instances of organized backlash to advertising.

In the 1940s, after aggressive state-run advertising campaigns sold engagement in back-to-back world wars to skeptical publics in the US and Europe, the West grew suspicious of government advertising, labeling it with a word that had recently morphed from neutral communications jargon to pejorative: “propaganda.”³

Today, we are experiencing a similar moment of pushback over the spread of surveillant advertising online. Shoshana Zuboff’s complex articulation of the political and economic systems that underlie surveillant advertising has crystallized a set of complaints into a comprehensive critique.⁴ Unfortunately, the scale and ambition of her analysis has not been echoed in scale by proposals from civil society, whose solutions have mostly been limited to modest policy interventions. Meanwhile, industry giants have moved forward with the future of digital advertising. This future largely preserves the current pathologies of digital advertising, restricting surveillance by a welter of third-parties, but allowing and encouraging first-party surveillance by the tech giants themselves. Without a new vision for digital advertising, we are likely stuck with the systems Google and Apple have crafted.

We argue that a reconsideration of digital advertising should begin by recognizing the critical role that digital advertising serves in society: despite the current model’s flaws, digital advertising enables businesses to reach customers and content providers and services to monetize while maintaining “free as in beer” access for the public.⁵ Digital advertising systems like Google AdSense essentially serve as infrastructure for the

³ *Id.* at 122.

⁴ SHOSHANA ZUBOFF, *THE AGE OF SURVEILLANCE CAPITALISM* (2019).

⁵ Thanks to Sunil Abraham for recommending in the early stages of our work on digital public infrastructure that we pursue a responsible ad system as an initial project. Advertising enables content to be available free of charge to consumers. For the origins of “free as in beer,” see: Lawrence Lessig, *Free, as in Beer*, *WIRED* (Sept. 1, 2006), <https://www.wired.com/2006/09/free-as-in-beer/>.

modern internet, supporting a wealth of content and services. Our goal is not to eliminate this key infrastructure, but to make it compatible with values of human agency and privacy.

We offer a vision for responsible digital advertising, a values-led advertising system structured around a single design choice: avoiding the storage of behavioral data. We posit that a central flaw of digital advertising as currently architected is that it remembers too much.⁶ By restricting the storage of behavioral data, our proposal for a system of “forgetful advertising” addresses many of the ills of surveillant advertising. Forgetful advertising can still target ads using information like geography, intent, context, and whatever else can be gleaned from a single interaction between a user and a website, but it cannot remember any previous interactions to inform its targeting.

There is a moral argument for forgetful advertising beyond addressing the standard harms of surveillance. Anticipating future desires based on past behavior limits freedom to choose and to change. Predictive, surveillant advertising locks users into past behaviors, offering donuts to a dieter despite his best attempts to change. Forgetful advertising allows a user the agency of the moment, to make healthy or unhealthy decisions as an individual, rather than influencing someone to follow their pre-existing patterns and stereotypes.

We propose this advertising system as a digital public infrastructure (DPI), a component of a new public internet centered on building socially beneficial public spaces, not purely on returns on investment.⁷ The pillars of DPI—responsible

⁶ See Maciej Cegłowski, *The Internet With A Human Face*, IDLE WORDS (May 20, 2014), https://idlewords.com/talks/internet_with_a_human_face.htm. See also ZUBOFF, *supra* note 4, at 8.

⁷ For an overview of digital public infrastructure as a concept, see ETHAN ZUCKERMAN, CTR. FOR JOURNALISM & LIBERTY, *WHAT IS DIGITAL PUBLIC INFRASTRUCTURE?* (2020) [hereinafter *WHAT IS DPI?*]; ETHAN ZUCKERMAN, KNIGHT FIRST AMENDMENT INST., *THE CASE FOR DIGITAL PUBLIC INFRASTRUCTURE* (2020) [hereinafter *THE CASE FOR DPI*].

design, community partnerships, alternative funding models, and institution building—are at the heart of our proposal.

We do not assert that such advertising will be as effective or profitable as surveillant advertising—though we note that Tim Hwang has argued that surveillant advertising may be less effective than promised.⁸ Instead, we focus on advertisers, content providers, and services who see a need for a digital advertising system consistent with their values, and on the importance of showing a viable alternative to current models.

The larger project of imagining and building digital public infrastructures begins from the presumption that a different future is possible than the one being offered by the dominant platforms.⁹ Our contribution is as much an invitation for others to imagine what more responsible digital advertising could be as it is a roadmap for a new infrastructure. By proposing and building a new system and seeking its adoption, we hope not just to challenge existing attention merchants with a new, less fraught system but to open conversation and creative exploration of how advertising could work differently online.

The rest of the paper proceeds as follows:

Section II, “Responsible digital advertising,” lays out the conceptual framework and values underlying our proposed system. Section II.A argues that we should understand digital advertising as infrastructure. Section II.B explores DPI in detail and applies its ideas to digital advertising. Section II.C discusses the values underlying forgetful advertising. Section II.D answers the questions: Who is this system for? What is a values-led organization and why might they find this system useful?

⁸ TIM HWANG, *SUBPRIME ATTENTION CRISIS* 75 (2020).

⁹ Google and Apple have both at least partly articulated their visions for the future of advertising. Chetna Bindra, *Building a privacy-first future for web advertising*, GOOGLE ADS & COMMERCE BLOG (Jan. 25, 2021), <https://blog.google/products/ads-commerce/2021-01-privacy-sandbox/>; *App Privacy Details*, APPLE DEVELOPER (July 1, 2022), <https://developer.apple.com/app-store/app-privacy-details/>.

Section III, “Forgetful advertising”, lays out our proposal in detail. Section III.A provides an overview of the current digital advertising system. Section III.B evaluates existing proposals to reform digital ad systems, most notably from Google and Apple, along with efforts from Brave and EthicalAds. Section III.C argues for forgetful advertising and considers its implementation.

I. Responsible Digital Advertising

A. Understanding Infrastructures

We define infrastructures as “the technologies and systems necessary for society to function.”¹⁰ Infrastructures are fundamental systems that allow us to build other systems—new houses and businesses rely on the infrastructures of electric power lines, water mains, and roads—and infrastructures are often invisible so long as they work well. Infrastructures tend to be expensive to build and difficult to scale. When a town outgrows the capacity of a small power plant, it’s expensive to build a larger power plant, and the larger plant will be underutilized for some time to come. As a result, costly and “bulky” infrastructures are not always built by private businesses.¹¹ They are often built by governments, because they are expensive, because their benefits can take a long time to be realized, and because sometimes only governments are well-positioned to capture the revenue generated by infrastructures, through taxation.¹²

Infrastructures generate externalities, both positive and negative: the road that connects the small town to the nearby city leads to new homes and businesses (positive economic externalities) and to increased noise and air pollution from increased car traffic (negative environmental externalities). And while a private business might make money by building a road

¹⁰ ZUCKERMAN, WHAT IS DPI?, *supra* note 7, at 4.

¹¹ Johan Fourie, *Economic Infrastructure: A Review of Definitions, Theory and Empirics*, 74 S. AFR. J. ECON. 530, 532 (2006).

¹² See BRETT M. FRISCHMANN, INFRASTRUCTURE: THE SOCIAL VALUE OF SHARED RESOURCES 13-14 (2012).

and charging tolls to use it, a government might be better positioned to capture revenue through taxing the homes and businesses that benefit from using the road.

The contemporary consumer internet is both dependent on existing infrastructures and an infrastructure in and of itself. The internet was built on existing telephone lines and within the confines of pre-existing power grids and has grown to include novel and mostly invisible systems like internet exchange points and cloud hosting providers.

As well as depending on existing infrastructures, the internet has proven to be a key infrastructure that makes new systems possible. Amazon's remarkable "everything store" relies on physical infrastructures—transoceanic container ships, the postal system, and private delivery services—and digital ones, from credit card processors to the fact that any internet-connected computer, phone, or device can access Amazon's website. Uber and other ridesharing companies use the physical infrastructures of roads and bridges, the digital infrastructures of payment systems and the web, plus the Global Positioning System.¹³

Infrastructures also exist in the social world. Schools, universities, health care systems, news organizations, legislative and judicial systems, public safety and policing systems, public spaces and libraries all operate as social infrastructures, enabling society to function.¹⁴ The challenge of conducting business in a society without a functioning court system or daily life in a country without a police force serve as reminders that like other infrastructures, social infrastructures allow more complex forms of interaction to occur.

¹³ GPS is notable because it is funded entirely by the U.S. government, operated by the U.S. Space Force, and part of the Department of Defense. *Global Positioning System Fact Sheet*, UNITED STATES SPACE FORCE, <https://www.spaceforce.mil/About-Us/Fact-Sheets/Article/2197765/global-positioning-system/> (last visited Aug. 23, 2022).

¹⁴ ERIC KLINENBERG, PALACES FOR THE PEOPLE: HOW SOCIAL INFRASTRUCTURE CAN HELP FIGHT INEQUALITY, POLARIZATION, AND THE DECLINE OF CIVIC LIFE 16-17 (2018).

Expanding definitions in this way, we argue that digital advertising qualifies as an infrastructure. It is a complex system that provides a critical service that enables society's broader functioning: it enables businesses to reach customers, and content providers and services to monetize while maintaining "free as in beer" access for the public.

Infrastructures have a set of responsibilities that come with their critical role. Those responsibilities are codified in various ways, including government regulation and professional norms. For example, water utilities' responsibilities are codified in regulations requiring, among other things, nondiscrimination, price restrictions, and quality controls. Similarly, news organizations' responsibilities are codified through professional norms encouraging fairness, transparency, and evidence.

Recognizing digital advertising as an infrastructure begs the question: what are its responsibilities? Our proposal aims to answer that question, building on the pillars of DPI—responsible design, community partnerships, alternative funding models, and institution building—to outline a set of responsibilities for digital advertising systems.

B. A DPI for Advertising

We propose conceptualizing this advertising system as a DPI, a component of a new public internet centered on building socially beneficial public spaces, not purely on returns on investment. Building an ad system as a DPI means taking a different approach to building technology, one that places civic values, usefulness, and public responsibility at the center of what and how we build.

Digital public infrastructures are systems that allow the digital world to function and permit public and civic life in digital spaces.¹⁵ Currently, our digital world has a set of "accidental" digital public infrastructures like Facebook and Google AdSense that were designed to solve a specific problem—staying in touch with friends or putting ads on your website—but

¹⁵ ZUCKERMAN, WHAT IS DPI?, *supra* note 7, at 7.

have grown into indispensable infrastructure for life online and offline. We argue that instead of relying on “accidental” DPIs to support public and civic life online, we should build intentional DPIs that center the responsibilities of public infrastructure. We propose this ad system as a DPI because we recognize the critical role that digital advertising serves in society. Our goal is not to eliminate this key infrastructure, but to make it compatible with civic, public-regarding values.

Moreover, in proposing a digital advertising system as a DPI, we do not propose a universal infrastructure that entirely replaces existing digital ad systems. We do not believe that what we have proposed is necessarily the right approach for all situations. Furthermore, it’s difficult to make progress towards a goal if that goal is to make a total change to a \$100B+ industry dominated by some of the most powerful companies in the world. Instead, our goal is to offer a complementary and contrasting option for those who would use it.

Thus, we propose building an advertising system for and with the group of advertisers, content providers, and services that already would like to use a system similar to our proposal.¹⁶ In addition to the reasons articulated above, this approach has the benefit of avoiding the “build it and they will come” assumption. Too often public-interest technologists design and build with no eye to what people actually want or need, an approach that is destined to fail and that smacks of the worst of technocracy.

We believe that many buyers and sellers of digital advertising may be values-led organizations whose values are inconsistent with existing digital advertising systems. A project at *The New York Times* illustrates this inconsistency. The *Times* ran a series called “The Privacy Project” from 2019-2020 which focused on the ways digital technology threatens people’s

¹⁶ See, e.g., PUBLIC SPACES COALITION, <https://publicspaces.net> (last visited Aug. 23, 2022); *Public Stack: the alternative internet*, WAAG.ORG, <https://waag.org/en/project/public-stack-alternative-internet> (last visited Aug. 23, 2022).

privacy.¹⁷ Ironically, observers noted that the same web pages that the *Times* was devoting to investigating the harms of digital surveillance were participating in it, tracking readers for the purpose of selling surveillant ads.¹⁸ A different approach might see the *Times* adopt an advertising system similar to the one we are proposing, one which is more consistent with its values. Later, we discuss the case of NPO, the coordinating organization of Dutch public broadcasters, experimenting with less-surveillant advertising for that very reason.

Relatedly, the funding for a DPI for advertising could come from investors and partners who recognize the value of such a system. However, in today's investment climate, where funding for technology is often constrained by the demands of double digit returns on investment, this may be difficult. Another option is public funding. Zuckerman, building on a suggestion from Paul Romer, has proposed a tax on surveillant advertising to fund the construction of DPIs.¹⁹ Funding could also come from donations and volunteers, similar to the way public media is funded in the United States.

Additionally, building a DPI requires building an organization that is largely missing from today's tech industry—the “trusted and trustworthy” intermediate institution. Jack Balkin outlines the concept in the context of social media,²⁰ but we think it applies to DPI more generally. Building healthier digital spaces requires trusted and trustworthy intermediate institutions with public regarding professional norms. Without them, digital spaces are ruled by economic incentives which,

¹⁷ *The Privacy Project*, N.Y. TIMES, <https://www.nytimes.com/series/new-york-times-privacy-project> (last visited Aug. 23, 2022).

¹⁸ Timothy Libert, *This Article Is Spying on You*, N.Y. TIMES (Sept. 18, 2019), <https://www.nytimes.com/2019/09/18/opinion/data-privacy-tracking.html>.

¹⁹ ZUCKERMAN, WHAT IS DPI?, *supra* note 7, at 10; Paul Romer, *A Tax That Could Fix Big Tech*, N.Y. TIMES (May 6, 2019), <https://www.nytimes.com/2019/05/06/opinion/tax-facebook-google.html>. Such a tax would also incentivize experimenting with alternatives to surveillant advertising systems.

²⁰ JACK BALKIN, KNIGHT FIRST AMENDMENT INST., HOW TO REGULATE (AND NOT REGULATE) SOCIAL MEDIA (2020).

unconstrained by public regarding values, can worsen life online and offline. In addition to the design we have proposed, this would mean the organization that controls the digital advertising system would make a commitment to upholding civic and public-regarding values and meaningful engagement with the public, policymakers, researchers, and users.

Building an ad system as a DPI means taking a different approach to building technology. In practice this means responsible design, community partnerships, alternative funding models, and institution building. It's not by accident that only one of these pillars involves a technical innovation. We are at a point now in the internet's development where much of the technical innovation has occurred and been deployed. We may be approaching what Carlota Perez terms a "turning point" in her theory of technological revolutions, where social and institutional choices will determine whether the synergy between society and a technology is beneficent and leads to a "Golden Age" or whether the synergy results in harm.²¹ It makes sense, then, that DPI is fundamentally a social and institutional innovation, not a technical one.

C. When is advertising good? When is targeting good?

When setting out to design a responsible digital advertising system, we began with two questions:²²

When is advertising good? When is targeting good?

By "good" we mean socially beneficial given its tradeoffs. We acknowledge that all systems have tradeoffs, and we do not assert that our answers are the only way of navigating these particular tradeoffs. We welcome this discussion and think it's essential—what's at stake here, as usual, is differences in what people value. Once we acknowledge that, it's easier to consider

²¹ CARLOTA PEREZ, TECHNOLOGICAL REVOLUTIONS AND FINANCIAL CAPITAL: THE DYNAMICS OF BUBBLES AND GOLDEN AGES 52-53 (2002).

²² These two questions are inspired by the three questions in Ben Thompson, *Philosophy and Power; Advertising, Targeting, and Tracking; The Real Winners*, STRATECHERY (July 14, 2021), <https://stratechery.com/2021/philosophy-and-power-advertising-targeting-and-tracking-the-real-winners/>.

different arguments, look for points of agreement, and have a productive discussion.

Here are our answers to the two questions:

Advertising is good when it is informational: that is, when it connects buyers and sellers, informs citizens, or disseminates price information.²³ Advertising is also good when it supports quality content and services. Advertising is bad when it distributes harmful information and when it supports fraudsters.

Targeted advertising is good when it enables businesses, nonprofits, and political campaigns to more effectively reach relevant people. Targeted advertising is bad when it becomes surveillant advertising—infringing on people’s right to self-determination by constructing an invasive and inflexible taxonomy of their lives. Beyond the harms of surveillance—which include a chilling effect on civil liberties and a power asymmetry that risks discrimination, coercion, manipulation, and selective enforcement—surveillant advertising’s end goal of anticipating future desires based on past behavior limits freedom to choose and to change.²⁴ Consider the alcoholic with a track record of responding to alcohol advertising. What happens when she chooses sobriety and her digital doppelganger is still a heavy drinker? (One of our authors fits this description and has been lobbying Twitter for an option to opt out of alcohol advertising, which they currently do not offer.) Predictive, surveillant advertising locks users into past behaviors, offering whiskey to a person in recovery despite his best attempts to change. Surveillant advertisers want people to be predictable, so that they will continue behaviors they’ve engaged in previously. Humans change by being unpredictable, growing beyond harmful behaviors into healthier ones (and vice versa).

²³ MICHAEL SCHUDSON, *ADVERTISING, THE UNEASY PERSUASION* 239 (1984).

²⁴ For an exploration of the harms of surveillance, see Neil Richards, *The Dangers of Surveillance*, 126 *Harv. L. Rev.* 1934 (2013).

Our answers are certainly not the answers a marketer or an attention merchant might give. But they dictate the design of our proposal.

D. Who is this system for? What is a values-led organization and why might they find this system useful?

Our proposal is not meant to wholly replace existing digital advertising systems. Indeed, we are consciously proposing a system that is limited. Our system chooses not to use information used by other digital advertising systems for normative reasons, recognizing that this choice may limit the system's efficiency and efficacy.

In this sense, we propose an advertising system that is analogous to the fair-trade movement in products like coffee. Recognizing that the system that produces most of the world's coffee is environmentally degrading and exploits the labor of vulnerable people, coffee-producing cooperatives have sought to provide higher quality products at higher prices, marketed to people who feel better about their purchase due to fewer ethical compromises in its sourcing.²⁵ Fair trade now represents 35% of the market for coffee and smaller amounts for other commodity crops like chocolate.²⁶

While our system is unlikely to displace the entire online advertising ecosystem, it could fill an important niche for buyers and sellers of digital advertising who do not wish to participate in existing systems. Furthermore, the simple existence of a system that challenges the status quo serves as a reminder of the downsides of the prevailing model.

We believe there is a market for more responsible digital ad systems, particularly from a set of organizations for whom

²⁵ Elizabeth Anne Bennett, *A Short History of Fairtrade Certification Governance*, in *THE PROCESSES AND PRACTICES OF FAIR TRADE: TRUST, ETHICS AND GOVERNANCE* 43, 45-49 (Brigitte Granville & Janet Dine eds., 2013).

²⁶ *5 Common Myths About Fair Trade Coffee*, FAIR TRADE CERTIFIED (Sept. 16, 2019), <https://www.fairtradecertified.org/blog/fair-trade-coffee-myths/>.

maximizing profits is less important than staying true to a set of organizational values. These values-led organizations include public entities like Dutch public broadcast organization NPO, which sells advertising, but is governed under public service charters that require they operate consistent with a set of public values. These values include, in many cases, a respect for consumer privacy that is inconsistent with most existing digital advertising solutions. NPO has been experimenting with less-surveillant digital advertising solutions on their websites, and along with the benefits from using a system more aligned with their values, has been pleasantly surprised to see no fall in revenues.²⁷ We believe that there are many organizations that would make a similar choice if presented with options less onerous than choosing between surveillant advertising and no digital advertising at all. Our proposal is aimed initially at this set of organizations.

II. Forgetful advertising

A. The current digital advertising system

To lay out our proposal we first need to explain what the current digital advertising system looks like. The graphic in Figure 1, developed by an investment bank that specializes in digital advertising, gives some sense of the complexity of the space and its basic operation.²⁸

²⁷ Natasha Lomas, *Data from Dutch public broadcaster shows the value of ditching creepy ads*, TECHCRUNCH (July 24, 2020), <https://techcrunch.com/2020/07/24/data-from-dutch-public-broadcaster-shows-the-value-of-ditching-creepy-ads/>.

²⁸ *Display LUMAscape*, LUMAPARTNERS.COM, <https://lumapartners.com/content/lumascape/display-ad-tech-lumascape/> (last visited Aug. 23, 2022).

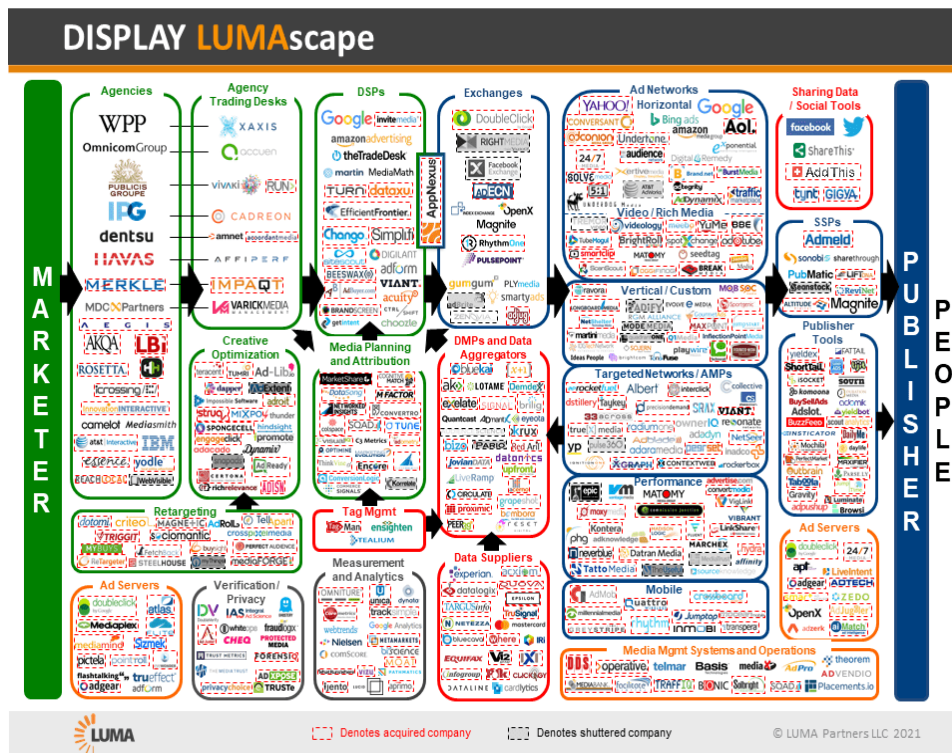


Figure 1

On the “buy-side”, buyers of digital ad space use a number of different tools to purchase ad space. All of these tools place bids on an exchange—think the NYSE for ads.²⁹ See Figure 2 for a visual representation. We will refer to these tools as “buying tools.”³⁰

²⁹ See generally Dina Srinivasan, *Why Google Dominates Advertising Markets*, 24 STAN. TECH. L. REV. 55 (2020).

³⁰ See Srinivasan’s use of the term “buying tools” in *id.*

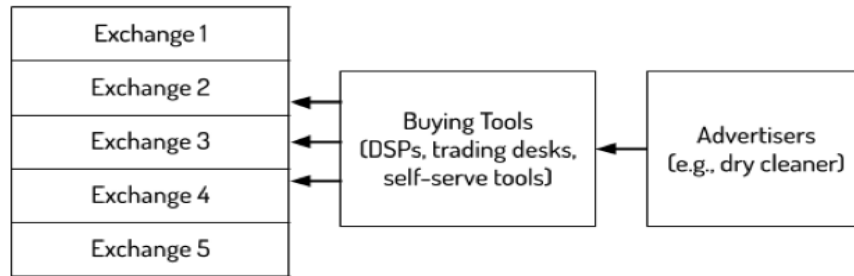


Figure 2

On the “sell-side”, sellers of digital ad space use “ad servers” to sell their ad space on exchanges.³¹ See Figure 3 for a visual representation. Ad servers track sellers’ available ad space and route it to exchanges in real-time.³² After ad space has been purchased, ad servers ensure advertisers’ ads are displayed “in the right spot, at the right time, to the right users.”³³

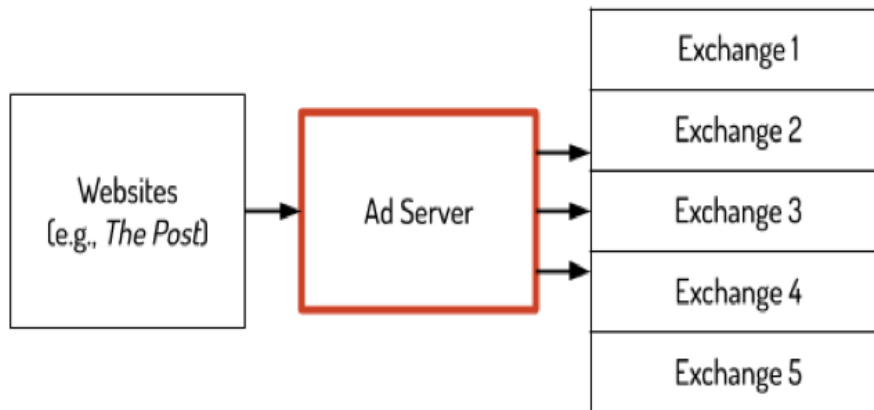


Figure 3

These three intermediary systems—the ad server, the exchange, and the buying tools—are what enable digital advertising. Together they handle a sequence of events that starts the moment a user visits a webpage and that ends before the page

³¹ *Id.*

³² *Id.*

³³ *Id.*

finishes loading.³⁴ As soon as a user visits a page, the ad server routes information about the ad space and user to one or more exchanges. The exchange then sends “bid requests” to the buying tools, asking them to return bids for the ad space. The exchange then collects bids, holds an auction, picks a winning bid, and returns it to the ad server. Finally, the ad server routes the advertisement associated with the winning bid to the page.³⁵ See Figure 4 for a visual representation of the three intermediary systems together.

Major platforms like Amazon and Facebook use basically the same model to sell advertising on their sites. You can think of their systems as using a single buying tool (for example, Facebook’s self-serve system) to bid on a combination exchange/ad server that sells a single company’s (for example, Facebook) ad space.³⁶ We believe our analysis largely applies to these closed systems as well.

What we’ve presented here is a high-level overview, generalized and idealized. There are many variations of digital advertising and a lot of other technology that is involved in the process (like brand safety tools that aim to help brands avoid inappropriate content).³⁷ However, for the purposes of our analysis, much of that complexity is out of scope.

³⁴ *Id.*

³⁵ *Id.* See also Shoshana Wodinsky, *How Google Ruined the Internet (According to Texas)*, GIZMODO (Dec. 17, 2020), <https://gizmodo.com/how-google-ruined-the-internet-according-to-texas-1845902795>.

³⁶ See, e.g., Cade Metz, *How Facebook’s Ad System Works*, N.Y. TIMES (Oct. 12, 2017), <https://www.nytimes.com/2017/10/12/technology/how-facebook-ads-work.html>; *Good Questions, Real Answers: How Does Facebook Use Machine Learning to Deliver Ads?*, FACEBOOK (June 11, 2020), <https://www.facebook.com/business/news/good-questions-real-answers-how-does-facebook-use-machine-learning-to-deliver-ads>.

³⁷ See, e.g., Augustine Fou, *Look For The ‘Drop-Offs’ In Your Digital Campaigns*, FORBES (Oct. 28, 2020), <https://www.forbes.com/sites/augustinefou/2020/10/28/look-for-the-drop-offs-in-your-digital-campaigns>; *The Evolution of Brand Safety*, GUMGUM, <https://gumgum.com/the-evolution-of-brand-safety> (last visited Aug. 23, 2022).




Figure 4

Our proposal focuses on ad servers. Ad servers are where data about users’ behavior is collected and stored. The problem we identified with digital advertising as it currently exists—it remembers too much—is rooted in ad servers.

B. Existing reform proposals

There are several existing proposals that seek to reform digital ad systems, most notably from Google and Apple, along with efforts from Brave and EthicalAds.

Google and Apple’s assortment of proposals for digital advertising essentially restrict third-party surveillance while allowing and encouraging first-party surveillance.³⁸ This has the

³⁸ Google acknowledges this in a blog post outlining the changes: “First-party relationships are vital. . . . We will continue to support first-party relationships on our ad platforms for partners. . . . [W]e’ll deepen our support for solutions that build on these direct relationships” David Temkin, *Charting a course towards a more privacy-first web*, GOOGLE ADS & COMMERCE BLOG (Mar. 3, 2021), <https://blog.google/products/ads-commerce/a-more-privacy-first-web/>. See also Issie Lapowsky, *Concern trolls and power grabs: Inside Big Tech’s angry, geeky, often petty war for your privacy*, PROTOCOL (July 13, 2021), <https://www.protocol.com/policy/w3c-privacy-war/>; Ben Thompson, *Philosophy and Power; Advertising, Targeting, and Tracking; The Real Winners*, STRATECHERY (July 14, 2021), <https://stratechery.com/2021/philosophy-and-power-advertising-targeting-and-tracking-the-real-winners/>. Apple’s proposals are visible in its privacy policy. See *User Privacy and Data Use*, APPLE, <https://developer.apple.com/app-store/user-privacy-and-data-use/> (last visited Aug. 23, 2022).

convenient effect of strengthening Google and Apple's position in digital advertising, allowing them to set the rules of the game and use their vast amounts of first-party data to target advertising.³⁹ Google and Apple's approach addresses some concerns about third parties tracking users across the internet.⁴⁰ But that change is largely a difference in degree, not kind. By leaving the core surveillant architecture of digital advertising untouched, and simply encouraging a shift from third-party surveillance to first-party surveillance, the harms of surveillant advertising will live on, albeit with a smaller set of offenders.⁴¹

Our proposal rejects surveillant advertising in all its forms, first- and third-party. In fact, we recognize that some of the most significant privacy harms in the digital age have come from stores of first-party data: ISP data in the case of the Snowden revelations; Google search and location data in the case of geofence and keyword warrants; Facebook data in the case of

³⁹ *Apple Advertising & Privacy*, Apple, <https://www.apple.com/legal/privacy/data/en/apple-advertising/> (last visited Aug. 23, 2022) (disclosing that Apple may use collect and use a user's name, address, age, gender, registered devices, downloaded media, in-app purchases, ad interactions, search activity, and more to facilitate targeted advertising). See also Patrick McGee, *Apple's privacy changes create windfall for its own advertising business*, FINANCIAL TIMES (Oct. 17, 2021), <https://www.ft.com/content/074b881f-a931-4986-888e-2ac53e286b9d>.

⁴⁰ See, e.g., Steven Melendez & Alex Pasternack, *Here are the data brokers quietly buying and selling your personal information*, FAST COMPANY (Mar. 2, 2019), <https://www.fastcompany.com/90310803/here-are-the-data-brokers-quietly-buying-and-selling-your-personal-information>.

⁴¹ See, e.g., Alex Berke and Dan Calacci, *Privacy Limitations Of Interest-based Advertising On The Web: A Post-mortem Empirical Analysis Of Google's FLoC 4, 11* (July 14, 2022) (unpublished manuscript) (available at <https://arxiv.org/pdf/2201.13402.pdf>) (suggesting that, because of the nature of first-party cookies, Google's original plan to mitigate the risks of third-party tracking likely would not reduce the overall amount of tracking); Patrick McGee, *Apple reaches quiet truce over iPhone privacy changes*, FINANCIAL TIMES (Dec. 8, 2021), <https://www.ft.com/content/69396795-f6e1-4624-95d8-121e4e5d7839> (observing that personal data and device information was still collected by "top apps" notwithstanding Apple's stricter privacy policy).

Cambridge Analytica.⁴² Even the villains of third-party tracking and data collection—data brokers—have a path forward in Google and Apple’s paradigm. An enterprising data broker could merge with a digital content provider or service so that the data they collect and trade is under a first-party umbrella, and thus officially sanctioned by Google and Apple. As we will show, their path forward is much less clear under forgetful advertising.

Our proposal can be contrasted with one offered by Brave, a privacy-focused web browser. Building off previous academic recommendations, Brave proposes a system that aims to enable privacy-friendly targeting and attribution along with a transparent and fair ad market. To achieve these goals, Brave uses on-device targeting combined with smart contracts and a blockchain to measure user attention and handle and log transactions.⁴³ By moving ad targeting on-device, Brave hopes to

⁴² Julia Angwin et al., *AT&T Helped U.S. Spy on Internet on a Vast Scale*, N.Y. TIMES (Aug. 15, 2015), <https://www.nytimes.com/2015/08/16/us/politics/att-helped-nsa-spy-on-an-array-of-internet-traffic.html>; Thomas Brewster, *Exclusive: Government Secretly Orders Google To Identify Anyone Who Searched A Sexual Assault Victim’s Name, Address Or Telephone Number*, FORBES (Oct. 4, 2021), <https://www.forbes.com/sites/thomasbrewster/2021/10/04/google-keyword-warrants-give-us-government-data-on-search-users/?sh=117582707c97>; Jennifer Valentino-DeVries, *Tracking Phones, Google Is a Dragnet for the Police*, N.Y. TIMES (Apr. 13, 2019), <https://www.nytimes.com/interactive/2019/04/13/us/google-location-tracking-police.html>; Carole Cadwalladr & Emma Graham-Harrison, *Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach*, THE GUARDIAN (Mar. 17, 2018), <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election>.

⁴³ BRAVE SOFTWARE, BASIC ATTENTION TOKEN (BAT): BLOCKCHAIN BASED DIGITAL ADVERTISING (2021), <https://basicattentiontoken.org/static-assets/documents/BasicAttentionTokenWhitePaper-4.pdf>; GUHA ET AL., MICROSOFT RSCH. INDIA, PRIVAD: PRACTICAL PRIVACY IN ONLINE ADVERTISING (2011), <https://www.usenix.org/conference/nsdi11/privad-practical-privacy-online-advertising>; TOUBIANA ET AL., ADNOSTIC: PRIVACY PRESERVING TARGETED ADVERTISING (2010), <https://crypto.stanford.edu/adnostic/adnostic-ndss.pdf>; Pestana et al., THEMIS: A Decentralized Privacy-Preserving Ad Platform with Reporting Integrity (Aug. 4,

limit the parties with access to sensitive data while preserving the ability to use that data to target ads. And by using smart contracts and a blockchain Brave hopes to bring transparency and fairness to ad transactions. However, Brave's system does not address first-party surveillance. Additionally, it makes assumptions about the uniformity of user attention and the prudence of crypto governance that we do not share.⁴⁴ Lastly, Brave does not address the harms of targeting advertising based on past behavior.

EthicalAds is an ad network based on three principles: "Ethical ads don't track youEthical ads are targeted by page content. . . . Ethical ad networks sell ads, not data."⁴⁵ EthicalAds' ad server is open source and its prices and fees are transparent.⁴⁶ It focuses on advertising to software developers and has built a network of quality publishers.⁴⁷ EthicalAds is the closest thing to a responsible digital advertising system that we've seen in practice. However, our vision is significantly more expansive, offering more of the tools and options digital advertisers are used to while providing a conceptually coherent paradigm.

C. *Forgetful Advertising*

Inspired by Maciej Cegłowski's critique of data collection on the internet, "The Internet With a Human Face," we posit that a central flaw of surveillant advertising systems is that they remember too much.⁴⁸ If digital ad systems were "forgetful"—avoided the storage of behavioral data—surveilling people and

2020) (unpublished manuscript) (available at <https://arxiv.org/pdf/2007.05556.pdf>).

⁴⁴ On uniformity: pricing is determined through "Basic Attention Metrics" which attempt to standardize the value of user attention. Replacing prices with a centrally planned metric seems destined to result in a less efficient and useful marketplace. See generally F.A. Hayek, *The Use of Knowledge in Society*, 35 AM. ECON. REV. 519 (1945).

⁴⁵ EthicalAds, "Our Advertising Vision", <https://www.ethicalads.io/advertising-vision/>.

⁴⁶ *Id.*

⁴⁷ *Id.*

⁴⁸ Cegłowski, *supra* note 6.

targeting them based on their previous behavioral patterns wouldn't be possible. The IDs and behavioral data surveillance advertising relies on would simply be off limits. Instead, information about ad space and the person looking at it would be limited to the context of a single interaction, with no IDs and storage involved. This could still include information about context, intent, and location (among other things), as long as that information wouldn't rely on any previous interactions. Forgetful advertising allows a user the agency of the moment, to make healthy or unhealthy decisions as an individual, rather than influencing someone to follow their pre-existing patterns and stereotypes.

What would forgetful advertising look like in practice?

Recall that when a user visits a page the ad server routes information about the ad space and user to one or more exchanges; the exchange then sends "bid requests" to the buying tools, asking them to return bids for the ad space. When exchanges make bid requests, they send data to buying tools which gives buying tools the information they need to determine whether and what to bid on the ad space. The data in a bid request is usually formatted as (key, value) pairs according to a data standard, usually OpenRTB, a standard maintained by the Interactive Advertising Bureau, or Google's proprietary standard.⁴⁹ Under surveillant advertising, the information passed as part of the bid request includes behavioral data, user IDs, location data, and contextual data.⁵⁰

The bid request is itself derived from data that the ad server sent the exchange. This data is usually less standardized but similarly is formatted as (key, value) pairs.⁵¹

⁴⁹ *OpenRTB Specification v3.0*, GITHUB (Aug. 2021), <https://github.com/InteractiveAdvertisingBureau/openrtb/blob/master/OpenRTB%20v3.0%20FINAL.md>; *Real-time Bidding*, GOOGLE DEVELOPERS, <https://developers.google.com/authorized-buyers/rtb/start#ad-exchange-protocol> (last visited Aug. 23, 2022).

⁵⁰ See *OpenRTB Specification v3.0*, *supra* note 49.

⁵¹ *Id.*

In concrete terms: imagine a user visits a web page on The Atlantic. Before the page has loaded, The Atlantic's ad server will send an exchange (key, value) pairs that include information about the ad space and the user. The Atlantic likely knows where a user is, what they've read on The Atlantic before, possibly their age and gender, along with a number of other characteristics collected or inferred from their activity. Usually, a user ID is included in the (key, value) pairs sent to an exchange which allows additional data about the user to be retrieved from various third-parties like data brokers or, if it's a Google intermediary, Google itself.

These (key, value) pairs are the heart of digital advertising. Without them (and the computing infrastructure that enables programmatic ads), the advertising industry would be back on the Madison Avenue of the 20th century, buying and selling advertising face to face. The data passed back and forth between sellers and buyers of digital ad space in a machine-readable format make digital advertising possible.

Thus, we can change digital advertising by changing the (key, value) pairs. The most important (key, value) pairs are the pairs sent by the ad server to the exchange. The pairs sent in the bid request from exchanges to advertisers are derived from the pairs sent by the ad server to the exchange.

What would the (key, value) pairs look like in forgetful advertising?

Forgetful advertising means that the (key, value) pairs sent by the ad server to the exchange will not contain user IDs, will not contain precise location information, will randomize fingerprinting information, and will not contain information derived from a user's previous behavior.

User IDs are the keys that unlock a vast ecosystem of third-party databases.⁵² They allow third parties, such as data

⁵² Antonio García Martínez, *To catch a priest*, THE PULL REQUEST (Aug. 2, 2021), <https://www.thepullrequest.com/p/to-catch-a-priest>. (“[The] ID, . . . is the ‘primary key’ . . . in endless tables of data scattered across countless ads and data companies. Like your social security number, it’s the

brokers or Google, to link their data together and attribute it to particular users. Without user IDs, much of the third-party data is unusable. Third-party data clearly violates forgetful advertising—it's data for the purposes of advertising collected about users from contexts other than the current interaction.

Precise location information can be used to track user's offline activities and uniquely identify them.⁵³ We thus suggest taking steps to reduce the precision of location information to make it more difficult to identify specific persons. This can be done by truncating latitude/longitude data, which makes it difficult to track and uniquely identify people, but allows the ability to target persons within a general geographic area.⁵⁴

Fingerprinting information such as IP address, browser type, screen size, font availability, and operating system can also be used to uniquely identify people.⁵⁵ Randomizing that information before passing it along prevents it from being used to uniquely identify people.⁵⁶

Additionally, importantly, and by definition, none of the (key, value) pairs will contain information derived from users' previous behavior. This is enforced by ensuring that the program(s) that computes the (key, value) pairs does not derive those (key, value) pairs from information about users' previous behavior. Periodic audits could fill this enforcement role—an option we explore in more detail at the end of this section.

common ID that ties everything together. . . . [A]ll the data becomes incommensurable without it.“).

⁵³ See, e.g., Stuart A. Thompson & Charlie Warzel, *Twelve Million Phones, One Dataset, Zero Privacy*, N.Y. TIMES (Dec. 19, 2019), <https://www.nytimes.com/interactive/2019/12/19/opinion/location-tracking-cell-phone.html>.

⁵⁴ Amy Fox, *Precision Matters: The Critical Importance of Decimal Places*, BLIS (July 9, 2017), <https://blis.com/precision-matters-critical-importance-decimal-places-five-lowest-go/>.

⁵⁵ *Fingerprint Randomization*, BRAVE (June 2, 2020), <https://brave.com/privacy-updates/3-fingerprint-randomization/>.

⁵⁶ *Id.*

Within the apparently restrictive framework we propose, there is a great deal of freedom to deliver personalized, targeted advertising.

Forgetful advertising does not preclude including profile information in a bid request, if someone has created an explicit profile on a site. For example, if someone creates a profile on a local news site with information about their age, gender, and interests, the local news site would then be able to include that information in a bid request.

Some may argue that allowing sites to use explicit profile data will lead to a race to the bottom, with sites requiring increasingly detailed profiles to be able to access them. This is unlikely for two reasons: first, excessive profile requirements will likely hurt a site's growth and usage. A longer registration process and privacy concerns can discourage people from signing up for a site.⁵⁷ Second, transparency can lead to public pressure and accountability—users will be confronted with the explicit profile information that sites want to know about them and can decide whether they're ok with it or whether the site is asking for information that seems irrelevant or unnecessarily invasive.

Arguably, our proposal might be made simpler by excluding explicit profile information altogether. However, for some advertisers this would limit the usefulness of forgetful advertising. The internet atomizes content and users, which is a significant change that holds value for some buyers and sellers of advertising. Finding a way to leverage this ability ethically is a challenge, but one any responsible ad system which aims to achieve significant adoption must face. We believe that allowing targeting based on explicit profiles (when combined with restrictions on identifying information) is a healthy approach to this problem and marks a step forward in responsible digital advertising.

⁵⁷ See, e.g., MIGUEL MALHEIROS & SÖREN PREIBUSCH, SIGN-UP OR GIVE-UP: EXPLORING USER DROP-OUT IN WEB SERVICE REGISTRATION 4-6 (2013).

Beyond profile targeting, forgetful advertising also permits contextual targeting. This means (key, value) pairs can include information about the specific context of the ad space (for example, a long essay about knitting) and/or the context of the site more generally (for example, gaming service that serves mostly 18-25 year-olds). Contextual targeting is forgetful as long as it's based on a single interaction (like the Atlantic article you're currently reading), not a collection of your previous contexts (every Atlantic article you've ever read).

Forgetful advertising also supports intent targeting. This is when you tell a site what you want and it serves an ad in response—think of the ads that Google serves when you search for “best lawyer in Detroit” or the ads that Amazon serves when you search for “paper towels.” Intent targeting is forgetful as long as it is based on a single interaction (like one Google search), not a collection of your previous intents (like your Google search history).

Finally, forgetful advertising supports geotargeting. This is targeting based on a person's location, like when you see ads for Detroit Tigers tickets because you live in Detroit. Geotargeting is forgetful when it isn't overly precise, because precise location information can uniquely identify someone, and when it's based on a single interaction (your current location), not a collection of your previous locations (your location history).

What's clear is that much of the information used in surveillant advertising can stay under forgetful advertising. The only changes are removing user IDs, modifying latitude/longitude and fingerprint information, and ensuring that targeting information is derived from the context of a single interaction. With these relatively simple changes we can implement forgetful advertising. (See the Appendix for a detailed examination of how a sample bid request would change under forgetful advertising.)

We acknowledge that there is no way to automatically ensure an ad server complies with forgetful advertising. However, periodic audits could certify whether an ad server complies with forgetful principles. Audits of algorithmic systems have

grown more popular in recent years, and have been adopted by major companies, making them easier to implement than ever before.⁵⁸ An auditor could periodically examine ad transactions, much like what we do in the Appendix for the example bid request, and ensure that they adhere to forgetful principles, offering a “forgetful ads” certification to compliant organizations. Additionally, as we envision forgetful advertising to be, at least initially, adopted voluntarily, it’s reasonable to assume that the organizations who choose to adopt it will want to comply. Exploring auditable ad intermediaries is a promising avenue for future work.

Conclusion

As Silicon Valley giants sketch the future of digital advertising, an infrastructure with significant implications for life online and offline, there are startlingly few alternatives to their vision. In response, we propose a digital public infrastructure for advertising, with forgetful advertising at its core. Our proposal is based on the contention that the storage of behavioral data is central to the harms of the current digital advertising system. We believe our proposal can make digital advertising compatible with the values of human agency and privacy, and offer it as a bottom-up solution for organizations that find existing digital advertising systems inconsistent with their values.

⁵⁸ See Christian Sandvig, Kevin Hamilton, Karrie Karahalios, & Cedric Langbort, *Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms* 5-6 (May 22, 2014) (unpublished manuscript) (<https://social.cs.uiuc.edu/papers/pdfs/ICA2014-Sandvig.pdf>); *MindBridge Completes World-First Algorithm Audit*, MINDBRIDGE (Sept. 27, 2021), <https://www.mindbridge.ai/news/mindbridge-completes-world-first-algorithm-audit/>; Nayeem Syed, *The new EU Regulation for AI is here, will you be ready for an AI audit—part one*, THOMSON REUTERS (Apr. 26, 2021), <https://blogs.thomsonreuters.com/legal-uk/2021/04/26/the-new-eu-regulation-for-ai-is-here-will-you-be-ready-for-an-ai-audit-part-one/>; Jayshree P. Upadhyay, *Deloitte to audit NSE’s algorithmic trading system*, MINT (Oct. 19, 2016), <https://www.livemint.com/Money/8YOx7OZDcAVW6Z8WhAGfIP/Deloitte-to-audit-NSEs-algorithmic-trading-system.html>.

Appendix

We will inspect an example bid request (see Figure 5) and analyze how it would change under forgetful advertising.⁵⁹

3.2.1 Object: **bidrequest**

Attribute	Always Passed	MoPub implementation specifics
id	Yes	Unique ID of the bid request, provided by the exchange. (this is the field that is mapped to "id" in bid response)
imp	Yes	Array of Imp objects (Section 3.2.4) representing the impressions offered. Only one impression object will be passed.
app	No (either site or app always)	Details via an App object (Section 3.2.14) about the publisher's app (i.e., non-browser applications). Sent for all mobile application requests.
device	Yes	Details via a Device object (Section 3.2.18) about the user's device to which the impression will be delivered.
user	No	Details via a User object (Section 3.2.20) about the human user of the device; the advertising audience.
test	No	Indicator of test mode in which auctions are not billable, where 0 = live mode, 1 = test mode. Sent to Advanced Bidders only.
at	Yes	Always set to "1" for MoPub, meaning 1st price auction.
tmax	yes	Maximum time in milliseconds to submit a bid to avoid timeout. The default value passed is 300 or 410ms for Marketplace line items and 1000 ms for the Unified Auction. For partners connected to our APAC POP, the default value is typically 300 ms.
wseat	No	Allowlist of buyer seats allowed to bid on this impression. Seat IDs must be communicated between bidders and the exchange a priori. Omission implies no seat restrictions.
bseat	No	Block list of buyer seats restricted from bidding on this impression. Seat IDs must be communicated between bidders and the exchange a priori. Omission implies no seat restrictions.
bcat	No	Blocked advertiser categories using the IAB content categories. Refer to section 5.1 of the OpenRTB 2.5 spec.
badv	No	Block list of advertisers by their top-level domains (e.g., "ford.com").
bapp	No	Block list of applications by their platform-specific exchange-independent application identifiers. On Android, these should be bundle or package names (e.g., "com.foo.mygame"). On iOS, these are numeric IDs.
source	No	A source object that provides data about the inventory source and which entity makes the final decision.
regs	No	A Regs object (Section 3.2.3) that specifies any industry, legal, or governmental regulations in force for this request.
ext	Yes	Placeholder for exchange-specific extensions to OpenRTB.

Figure 5

id: Remain—it's a unique id for the bid request, not for a user.

imp: Remain—metadata about the ad space.

app: Remain—metadata about the seller.

device (see Figure 6): ifa will be removed as it's a unique device identifier. Geo will remain but will be modified. The rest

⁵⁹ MoPub, "MoPub OpenRTB 2.5 and Native 1.2 Integration Guide", https://developers.mopub.com/bidders/integration/openrtb_2.5/#321-object-bidrequest.

of the fields remain but are subject to fingerprint randomization.

3.2.18 Object: **device**

Attribute	Always Passed	Description
ua	usually	Browser user agent string.
geo	yes	Location of the device assumed to be the user's current location defined by a Geo object (Section 3.2.19).
dnt	no	Standard "Do Not Track" flag as set in the header by the browser, where 0 = tracking is unrestricted, 1 = do not track. Only passed when DNT=1. Note that this field is the catch all for signaling not to behaviorally target. We will pass this for any browser dnt signal, lint signal or coppa flagged user.
lat	no	"Limit Ad Tracking" signal commercially endorsed (e.g., iOS, Android), where 0 = tracking is unrestricted, 1 = tracking must be limited per commercial guidelines.
ip	yes	IPv4 address closest to device.
devicetype	no	The general type of device. Note, for Android it is based on device width and height. Refer to List 5.21 .
make	no	Device make (e.g., "Apple").
model	no	Device model (e.g., "iPhone"). iOS will show the full device model name e.g. "iPhone10,1" if known, or just e.g. "iPhone" if not. Android will have detailed model information such as "SAMSUNG-SM-G980A"
os	no	Device operating system (e.g., "iOS" or "Android").
osv	no	Device operating system version (e.g., "3.1.2").
huv	no	Hardware version of the device. Sent for Android only.
h	no	Physical height of the screen in pixels. This may be dependent on the device orientation. Note this can differ from banner.h
w	no	Physical width of the screen in pixels. This may be dependent on the device orientation. Note this can differ from banner.w
pxratio	yes	The ratio of physical pixels to device independent pixels. A float value from 0.75 thru 4.
js	yes	Support for JavaScript, where 0 = no, 1 = yes, always = 1
language	no	Parsed out from HTTP headers as highlighted below. Acceptable values are ios-user-agent ISO 639-1 codes . iOS: Accept-Language: en-us Android: User-Agent: Mozilla/5.0 (Linux; U; Android 4.1.2; en-us; SAMSUNG-SGH-I317 Build/JZ054K) AppleWebKit/534.30 (KHTML, like Gecko) Version/4.0 Mobile Safari/534.30 Paraic/3.13
carrier	no	Carrier or ISP (e.g., "VERIZON"). "WiFi" is often used in mobile to indicate high bandwidth (e.g., video friendly vs. cellular). No change from MoPub 2.3 support, moved to standardized list of carrier codes ("MNC") from this list . This value will be formatted as follows: {mcc}{mnc} for example, "310-050" would represent Virgin Mobile US in the United States. This value will be passed along in the bidder when it is present via the SDK.
connectiontype	no	Network connection type. Refer to List 5.19 .
ifa	no	ID sanctioned for advertiser use in the clear (i.e., not hashed). The raw IDFA or Google Advertising ID, unhashed. Present on iOS8+ or apps using the Google Play Services SDK
ext	no	Placeholder for exchange-specific extensions to OpenRTB.

Figure 6

geo (see Figure 7): Lat and lon will be truncated for the reasons discussed above. The level of truncation will depend on whether the bid request includes other highly identifiable information that can be combined to uniquely identify someone—for example, date of birth and gender.⁶⁰ The rest of the attributes will stay the same. (Though the same analysis for determining the truncation of lat/lon should be applied to deciding whether to include ZIP codes.)

⁶⁰ See Latanya Sweeney, "Simple Demographics Often Identify People Uniquely", Carnegie Mellon University, Data Privacy Working Paper 3 (2000), <https://dataprivacylab.org/projects/identifiability/paper1.pdf>.

3.2.19 Object: **geo**

Attribute	Always Passed	Description
lat	no	Latitude from -90.0 to +90.0, where negative is south. Lat/lon are passed directly from the device if MoPub SDK integrated, or from the publisher if a 3rd party integration. Blank when not passed.
lon	no	Longitude from -180.0 to +180.0, where negative is west. Lat/lon are passed directly from the device if MoPub SDK integrated, or from the publisher if a 3rd party integration. Blank when not passed.
country	no	Country code using ISO-3166-1-alpha-3. Always derived from IP, never a reverse geocode from Lat/Lon
region	no	Region code using ISO-3166-2; 2-letter state code if USA. Always derived from IP, never a reverse geocode from Lat/Lon
metro	no	Metro codes taken from this list . Only be passed for US and Canada. Always derived from IP, never a reverse geocode from Lat/Lon
city	no	Always derived from IP, never a reverse geocode from Lat/Lon
zip	no	Always derived from IP, never a reverse geocode from Lat/Lon

Figure 7

user (see Figure 8): id will be removed. buyeruid will be removed. job and gender will remain if they are not uniquely identifying (same analysis as for lat/lon and ZIP) and if they aren't derived from stored behavioral data. (For example, they would be permitted if they are from an explicit profile or if they were inferred from the context of the page.) Keywords will remain but the keywords will not be derived from stored behavioral data. Instead, they can be derived from context, interaction-specific behavior or intent, or an explicit profile.

3.2.20 Object: **user**

Attribute	Always Passed	Description
id	no	Exchange-specific ID for the user. Equal to BidRequest.device.ifa if the idfa is non-zeroes. Else, the IDFV if available. Else, a MoPub-specific device ID for each user within each app.
buyeruid	no	Buyer-specific ID for the user as mapped by the exchange for the buyer. Sent to Advanced Bidders only. Max token size is 1kb and must be non-human readable and/or encrypted.
job	no	Year of birth as a 4-digit integer. Publisher passed when available.
gender	no	Gender, where "M" = male, "F" = female, "O" = known to be other (i.e., omitted is unknown).
keywords	no	Comma separated list of keywords, interests, or intent.
ext	no	Placeholder for exchange-specific extensions to OpenRTB.

Figure 8

test: Remains—flag used for testing.

at: Remains—flag for auction type.

tmax: Remains—maximum amount of time buying tools have to submit their bid.

wseat: Remains—allow list of buying tools allowed to submit a bid.

bseat: Remains—block list of buying tools not allowed to submit a bid.

bcat: Remains—blocked advertiser categories.

badv: Remains—block list of advertisers by their domain.

bapp: Remains—block list of apps.

source: Remains—information about the seller of the ad space.

regs: Remains—information about regulations that apply to the bid request.

ext: Remains—placeholder for exchange specific extensions to OpenRTB.

Here's why we made the changes we made:

The ifa will be removed because it's a unique device identifier, and as we discussed above, unique identifiers enable tracking and behavioral data storage. Latitude and longitude will be truncated because, as discussed above, precise location information can be used to track people's offline activities and can uniquely identify people. The remaining device fields will stay, but will be subject to fingerprint randomization to prevent them from being used to uniquely identify people.

User.id and user.buyerid will be removed because they are unique identifiers. User.keywords will remain, but keywords cannot be derived from behavioral data under forgetful advertising. Instead, they will have to be inferred from information like context, intent, or explicit profiles.

The remaining fields are metadata about the bid request and can remain.