This paper appears at the 2021 CHI Virtual Conference on Human Factors in Computing Systems.

What Makes a "Bad" Ad? User Perceptions of Problematic Online Advertising

Eric Zeng
Paul G. Allen School of Computer Science & Engineering
University of Washington
Seattle, WA, USA
ericzeng@cs.washington.edu

Tadayoshi Kohno
Paul G. Allen School of Computer Science & Engineering
University of Washington
Seattle, WA, USA
yoshi@cs.washington.edu

Franziska Roesner
Paul G. Allen School of Computer Science & Engineering
University of Washington
Seattle, WA, USA
franzi@cs.washington.edu

ABSTRACT

Online display advertising on websites is widely disliked by users, with many turning to ad blockers to avoid "bad" ads. Recent evidence suggests that today's ads contain potentially problematic content, in addition to well-studied concerns about the privacy and intrusiveness of ads. However, we lack knowledge of which types of ad content users consider problematic and detrimental to their browsing experience. Our work bridges this gap: first, we create a taxonomy of 15 positive and negative user reactions to online advertising from a survey of 60 participants. Second, we characterize classes of online ad content that users dislike or find problematic, using a dataset of 500 ads crawled from popular websites, labeled by 1000 participants using our taxonomy. Among our findings, we report that users consider a substantial amount of ads on the web today to be clickbait, untrustworthy, or distasteful, including ads for software downloads, listicles, and health & supplements.

CCS CONCEPTS

• **Information systems** → **Online advertising**; **Display advertising**; • **Social and professional topics** → **Commerce policy**; • **Human-centered computing** → *User studies.*

KEYWORDS

online advertising, deceptive advertising, dark patterns

1 INTRODUCTION

Online display advertising is a critical part of the modern web: ads sustain websites that provide free content and services to consumers, and many ads inform people about products and services that they are interested in. Still, many web users dislike online ads, finding them to be annoying, intrusive, and detrimental to their security or privacy. In an attempt to filter such "bad" ads, many users turn to ad blockers [5] — for instance, a 2016 study estimated that 18% of U.S. internet users and 37% of German internet users used an ad blocker [69], a large percentage considering that it takes some initiative and technical knowledge to seek out and install an ad blocker.

There are many drivers of negative attitudes towards online ads. Some users find the mere presence of ads to be problematic, often associated with their (perceived) increasingly disruptive, intrusive, and/or annoying qualities [5] or their impact on the load times of websites [92]. Users are also concerned about the privacy impacts of ads: research in computer security and privacy has revealed extensive ecosystems of tracking and targeted advertising (e.g., [9, 28, 30, 61, 62, 64, 76, 84, 97, 98]), which users often find to be creepy and privacy-invasive (e.g., [29, 96, 100, 101]). The specific *content* of ads can also cause direct or indirect harms to consumers, ranging from material harms in the extreme (e.g., scams [1, 34, 72], malware [65, 74, 104, 105], and discriminatory advertising [3, 57]) to simply annoying techniques that disrupt the user experience (e.g., animated banner ads [16, 38, 45]).

In this work, we focus specifically on this last category of concerns, studying people's perceptions of problematic or "bad" user-visible *content* in modern web-based ads. Driving this exploration is the observation that problematic content in modern web ads can be more subtle than flashing banner ads and outright scams. Recent anecdotes and studies suggest high volumes and a wide range of potentially problematic content, including "clickbait", advertorials or endorsements with poor disclosure practices, low-quality content farms, and deceptively formatted "native" ads designed to imitate the style of the hosting page [4, 7, 22, 39, 52, 63, 68, 71, 75, 90, 93, 103, 106]. While researchers and the popular press have drawn attention to these types of ad content, we lack a systematic understanding of how web users perceive these types of ads on the modern web in general. What makes an ad "bad", in the eyes of today's web users? What are people's perceptions and mental models of ads with arguably problematic content like "clickbait", which falls in a grey area between scams and poorly designed annoying ads? What exactly is it that causes people to dislike (or like) an ad or class of ads? For future regulation and research attempting to classify, measure, and/or improve the quality of the ads ecosystem, where exactly should the line be drawn?

We argue that such a systematic understanding of what makes an ad "bad" — grounded in the perceptions of a range of web users, not expert regulators, advertisers, or researchers — is crucial for two reasons. First, while some ads can clearly be considered "bad", like outright scams, and others can be considered "benign", like honest ads for legitimate products, there is a gray area where it is more nuanced and difficult to cleanly classify. For example, "clickbait" ads for tabloid-style celebrity news articles may not cross the line for causing material harms to consumers, but may annoy many users and use misleading techniques. While the U.S. Federal Trade Commission currently concerns itself with explicitly harmful ads like scams and deceptive disclosures [18, 33, 63], whether and how to address "clickbait" and other distasteful content is more nuanced. As part of our work, we seek to identify ads that do not violate current regulations and policies, but do harm user experiences, in order to inform improvements such as policy changes or the development of automated solutions. Second, research interested

Figure 1: An overview of our work and contributions.

in measuring, classifying, and experimenting on bad online ads will bene t from having detailed de nitions and labeled examples of bad ads, grounded in real users' perceptions and opinions. For example, our prior work measuring the prevalence of problematic ads on the web used a researcher-created codebook of potentially problematic ad content; that codebook was not directly grounded in broader user experiences and perceptions [106].

Research Questions. In this paper, our goal is thus to systematically elicit and study what kinds of online ads people dislike, and the reasons why they dislike them, focusing speci cally on the user-visible content of those ads (rather than the underlying technical mechanisms for ad targeting and delivery). We have two primary research questions:

(1) RQ1 De ning bad in ads : What are the di erent types of negative (and positive) reactions that people have to online ads that they see? In other words, why do people dislike (or like) online ads?
(2) RQ2 Identifying and characterizing bad ads : What speci c kinds of content and tactics in online ads cause people to have negative reactions? In other words, which ads do people dislike (or like)?

While ads appear in many places online including in social media feeds and mobile apps we focus speci cally on third party programmatic advertising on the web[2], commonly found on news, media, and other content websites. Unlike more vertically integrated social media platforms, the programmatic ad ecosystem is complex and diverse, with many di erent stakeholders and potential points of policy (non-)enforcement, including advertisers, supply-side and demand-side platforms, and the websites hosting the ads themselves. A bene t of our focus on web ads is that the public nature of the web allows us to crawl and collect ads across a wide range of websites, without needing to rely on explicit ad transparency platforms (which may be limited or incomplete [26, 87]) or mobile app data collection (which is more technically challenging). We expect that many of our ndings will translate to ads in other contexts (e.g., social media, mobile), though these di erent contexts also raise additional research questions about the interactions between the a ordances of those platforms and the types of ads that people like or dislike.

Contributions. Figure 1 shows an overview of the di erent components of our work and our resulting outputs and contributions. Speci cally, our contributions include:

(1) Based on a qualitative survey characterizing 60 participants' attitudes towards the content and techniques found in modern online web ads, we distill a taxonomy of 15 reasons why people dislike (and like) ads on the web, such as untrustworthy , clickbait , ugly / bad style , and boring (Section 3, answering RQ1).
(2) Using this taxonomy, we generate a dataset of 500 ads sampled randomly from a crawl of popular websites, labeled with 12,972 opinion labels from 1025 people (Section 4, towards answering RQ2). This dataset is available in the paper's supplemental materials[1].
(3) Combining participant opinion labels with researcher content labels of these 500 ads, and using unsupervised learning techniques, we identify and characterize classes of ad content and techniques that users react negatively to, such as clickbait native ads, distasteful content, deceptive and scammy content, and politicized ads (Section 4, answering RQ2).

Our ndings serve as a foundation for policy and research on problematic online advertising: for regulators, advertisers, and ad platforms, we provide evidence on which types of ads are most detrimental to user experience and consumer welfare, and for researchers, we provide a user-centric framework for de ning problematic ad content, enabling future research on the online advertising ecosystem.

2 BACKGROUND, RELATED WORK, AND MOTIVATION

2.1 Background and Related Work

Broadly speaking, related work has studied (1) problematic content and techniques in ads directly and/or (2) people's perception of ads. Our work is inspired by (1) and expands on (2). In this section, we break down prior work based on types of concerns with ads.

Computer Security Risks and Discrimination. Online ads have often been leveraged for malicious and harmful purposes. For example, prior work in computer security has studied the use of ads to spread malware, clickfraud, and phishing attacks (e.g., [65, 74, 82, 104, 105]). Researchers have also surfaced concerns about how ads may be targeted at users in potentially discriminatory ways (e.g., [3, 57]), such as by (intentionally or unintentionally) serving ads for certain employment opportunities disproportionately to certain demographic groups.

Deceptive Ads. Another class of problematic ads is those which are explicitly deceptive either in terms of the claims that they make, or in their appearance as advertisements at all. Prior work studying deceptive advertising predates web ads (i.e., print and TV ads), showing, for instance, that false information in ads can be e ectively refuted later only under certain conditions (e.g., [14, 50, 51]), that people infer false claims not directly stated in ads and misattribute claims to incorrect sources (e.g., [44, 48, 79, 85]), and that people's awareness of speci c deceptive ads can harm their attitudes towards those brands [49] as well as towards advertising in general [19, 23].

[1] Dataset also available at https://github.com/eric-zeng/chi-bad-ads-data

More recently on the web, there has been signi cant concern about native advertisements which are designed to blend into the primary content of the hosting website (e.g., sponsored search results or social media posts, or ads that look like articles on news websites). Signi cant prior work across disciplines suggests that most users do poorly at identifying such ads (e.g., [4, 9, 47, 53, 53, 59, 89, 102, 103]) though people may do better after more experience [53], or with di erent disclosure designs (e.g., [47, 103]). Deceptive ads may a ect user behavior even when identi ed [89]. Prior work suggests that native ads can reduce user perceptions of the credibility of the hosting site, even if the ads are rated as high quality in isolation [22]. Recent work has also raised concerns about unclear a liate marketing and/or endorsements on social media [71, 93, 108]. Beyond outright scams, much of the U.S. Federal Trade Commission's recent enforcement surrounding web ads has focused on ad disclosures for such native ads of various types [31, 33, 63].

Annoying and Disruptive Ads. Even when ads are not explicitly malicious or causing material harms, many users still dislike them. Traditionally, a common reason that people dislike web ads is that they are annoying and disruptive either in general or due to their speci c designs leading in part to the development and widespread adoption of ad blockers [5, 69, 81]. Prior work has studied and summarized design features of ads that lead to perceived or measured reductions in the user experience, including ads that are animated, too large, or pop up [12, 27, 38, 86]. The impacts of these issues include increased cognitive load, feelings of irritation among users, and reduced trust in the hosting websites and in advertising or advertisers [12, 17, 107].

Clickbait and Other Low-Quality Content. The concerns around annoying and disruptive ads in the previous paragraph stem largely from the design of ads. In addition, many modern web ads contain low-quality content that walks a ne line between good and bad ads. Recent anecdotal and scienti c evidence suggests that there is a wide range of problematic, distasteful, and misleading content in online ads (or on the websites to which they lead) including low-quality clickbait , content farms, misleading or deceptive claims, mis/disinformation, and voter suppression [8, 9, 24, 36, 37, 52, 52, 54, 56, 68, 75, 78, 90, 94, 95].

Our previous work measured the prevalence of these types of ads on news and misinformation sites, nding that large fractions of ads on both types of sites contained content that was potentially problematic (based on a researcher-created de nition of problematic) [106]. Though anecdotes suggest that users recognize and dislike such ads (for example, a recent qualitative study of French-speaking Twitter discussions surfaced user criticisms of social media ads using terms such as Fail , Clickbait , and Cringe [39]), user perceptions of these types of ads which seem not to directly violate the policies of ad providers or regulators today have not been systematically studied.

Further a eld but related are broader discussions of dark patterns [15], e.g., on websites [70] and in mobile apps (e.g., [1, 43, 77]), though none of these works considered web ads.

Ad Targeting and Privacy. Finally, a unique aspect of the online advertising ecosystem is the ability to track and target speci c individuals. Signi cant prior work has revealed and measured the privacy implications of these tracking and targeting capabilities (e.g. [28, 30, 61, 62, 64, 76, 84, 97, 98]), which end users may nd creepy , insu ciently transparent, or otherwise distasteful [6, 29, 96, 100, 101]. We note that such ad targeting may increase the impacts of problematic ad content if such content is delivered to particularly susceptible users (e.g., [83]).

2.2 Motivation

We identify several key gaps in prior work that we aim to address. First, studies of user perceptions of problematic ad content in the HCI community have focused largely on more traditional design issues (e.g., animated or explicitly deceptive ads), rather than the broader and less well-de ned range of clickbait and other techniques prevalent on the modern web. Second, research on the potential harms of online advertising in the computer security and privacy community primarily focuses on ad targeting, distribution, and malware, rather than the user-facing content of the ads. Finally, many anecdotes or measurement studies of potentially problematic content in ads rely on researcher-created de nitions of what is problematic, rather than being grounded in user perceptions. Are there types of problematic ad content that bother and harm users, but have not been addressed in prior measurement studies or in the policies of regulators and ad companies? And what exactly makes a bad ad bad? In this work, we aim to bridge these gaps through a user-centric analysis of ad content, eliciting user perceptions of a wide range of ads collected from the modern web and characterizing which attributes of an ad's content contribute to negative user reactions.

3 SURVEY 1: WHY DO PEOPLE DISLIKE ADS?

Towards answering our rst research question, we conducted a qualitative survey to elicit a detailed set of reasons for what people like or dislike about the content of modern online ads. The resulting taxonomy enables future studies that classify, measure, and experiment on bad online ads, including the second part of this paper (Section 4).

Though our primary research questions are around reasons that people dislike ads, we also collect data about reasons they may like ads. This is for two reasons: rst, we expect that there are ads that users genuinely like, and that a user may both like and dislike parts of an ad, so we aim to surface the full spectrum of users' opinions. Second, online ads are fundamental to supporting content and services on the modern web, and we aim for our work to ultimately improve the user experience of ads, not necessarily to banish ads entirely.

3.1 Survey 1 Methodology

3.1.1 Survey Protocol.

We curated a set of 30 ads found on the web (described below in Section 3.1.2). We showed each participant 4 randomly selected ads, and collected:

- Their overall opinion of the ad (5-point Likert scale).
- What they liked and disliked about it (free response).
- What they liked and disliked about similar ads if they remember them (free response).