

Privacy on Adult Websites

Ibrahim Altaweel,¹ Maximillian Hils,² and Chris Jay Hoofnagle³
September 2016

Abstract

As it does in polite conversation, pornography goes unmentioned in policy discussions. This paper begins a conversation about this major use of the web, one that is sensitive and could lead to embarrassment and harm to users if publicized. In countries where pornography is illegal, tracking of these behaviors could have profound consequences for users. Viewing such material is legal in the US, yet authorities may wish to avoid the topic of protecting its consumers.

We document and discuss the user tracking dynamics on the most popular adult-oriented websites (N=11). Tracking dynamics are different on adult sites than other popular sites. There are relatively fewer third-party tracking companies involved and fewer cookies than on comparably popular sites. However, we found that Google trackers (Google Analytics and/or DoubleClick) were present on almost all the sites and that search terms were often leaked in plaintext to third parties and sometimes encoded in cookies. Finally, the dominance of video in pornography could explain the presence of Flash on almost half the sites. We found Flash being used to read HTTP cookie values, but we did not find any evidence of Flash cookies respawning.

¹ Student, University of California, Santa Cruz; Web Security Engineer at Good Research LLC. Altaweel's work has been supported by NSF.

² Graduate Student, University of Münster; Visiting Scholar, UC Berkeley School of Information; Mitmproxy Core Developer. His open-source developments have been supported by Google and other technology-industry companies.

³ Adjunct Full Professor, UC Berkeley School of Information and School of Law. Investigator Hoofnagle's work has been supported by NSF, and by technology-industry companies, including Microsoft, Google, and Nokia. He is of counsel to Gunderson Dettmer, LLP, a firm that concentrates on representing emerging technology companies. We wrote this paper on our own initiative, without external financial support. However, we have significant technology donations from Palantir and Mezzobit. We thank Nathaniel Good for feedback and advice on this project.

Introduction

In this paper, we analyze web tracking on the most popular adult websites in the United States. This paper focuses on just eleven (N=11) websites—every adult-oriented website ranked in the top 500 US sites by Alexa.

It is important to analyze these sites because pornography appears to be a major use of the web. Even though statistics on the amount of the web that is pornographic are not very reliable,⁴ some adult-oriented websites clearly have a large amount of traffic, based on their relative rankings among the most frequented sites in the US. According to Alexa, the most visited site in our study is on par with BuzzFeed.com in popularity. The least visited site in our study is still more popular than Vox, Disney Go, PBS, and Mit.edu.

These rankings indicate that there is a great deal of pornography consumption online, yet, our society has strong expressed preferences for condemnation of and even prohibition of pornography. In May 2016, the Gallup Poll found that 61% of Americans think that pornography is morally wrong.⁵ This figure is remarkably durable: Gallup has found similar numbers going back to 2011. About a third of Americans favor laws to prohibit distribution of pornography to adults,⁶ and this sentiment appears to have support even among 18–24 year old respondents.⁷ Of course, such restrictions almost always run

⁴ See discussion in Kate Darling, *IP Without IP? A Study of the Online Adult Entertainment Industry*, 17 STAN. TECH. L. REV. 655 (2014); Shira Tarrant, *THE PORNOGRAPHY INDUSTRY: WHAT EVERYONE NEEDS TO KNOW* (Oxford Univ. Press 2016).

⁵ Gallup Organization. Gallup Poll, May, 2016 [survey question]. USGALLUP.052616.R21Q. Gallup Organization [producer]. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research, iPOLL [distributor], accessed Sep-17-2016. This survey result and others were obtained from searches of the iPOLL Databank and other resources provided by the Roper Center for Public Opinion Research.

⁶ National Opinion Research Center, University of Chicago. General Social Survey 2014, Mar, 2014 [survey question]. USNORC.GSS14B.Q102. National Opinion Research Center, University of Chicago [producer]. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research, iPOLL [distributor], accessed Sep-17-2016.

⁷ Berkley Center for Religion, Peace, and World Affairs at Georgetown University. Millennial Values Survey, Mar, 2012 [survey question]. USPRRI.12MILVAL.R22D. Public Religion Research Institute [producer]. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research, iPOLL [distributor], accessed Sep-17-2016 (a study of over 2,000 18–24 year old respondents found that 30% strongly

afoul of the First Amendment. However, that is not the case in other countries, where pornography is filtered or where its possession could be punished by the criminal law. Just this month, BBC News reported that Russian media regulator Roskomnadzor ordered the nation's ISPs to filter certain popular pornography sites.⁸

The moral disapprobation of pornography is so strong that it seems to also cause people to deny that they use the internet to visit adult websites. As late as 2005, the Pew Internet & American Life Project found that 87% of internet users claimed not to view adult websites.⁹ Thus, it is obvious that there are many users of adult websites who do not want others to know about their pornography consumption. Consider that in the abstract, we know that married couples have intercourse, and we even arrange public policy to encourage childbearing, yet direct or indirect evidence of intercourse is embarrassing and most people are discreet about it.

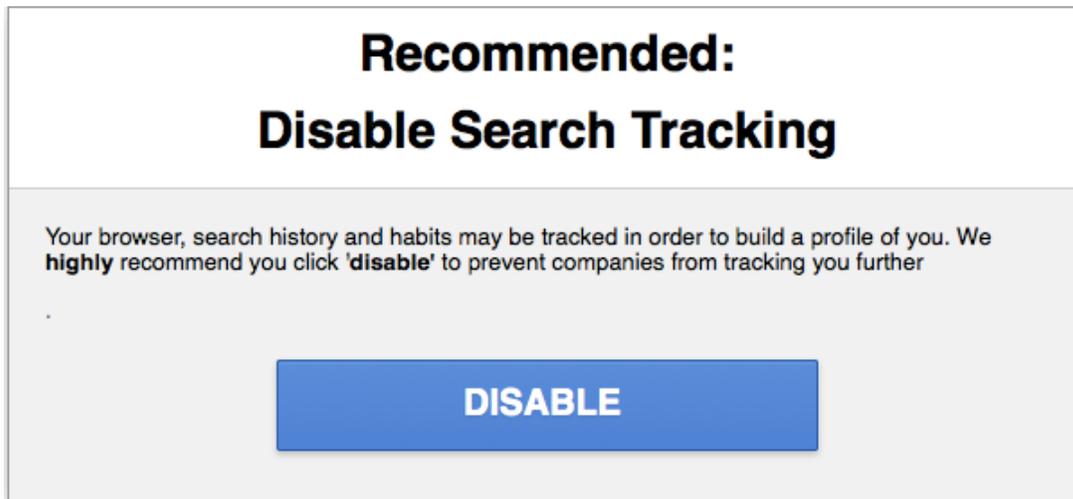


Figure 1: In our testing we found that some sites gave users privacy advice. This popup led to a search company that promised more privacy in queries.

favor, and 27% favor “Making it more difficult to access pornography on the Internet”).

⁸ BBC News, *Russia extends porn site ban*, Sept. 16, 2016.

⁹ Pew Internet & American Life Project. Pew Internet & American Life Project Poll, May, 2005 [survey question]. USPSRA.070605.WEBA12. Princeton Survey Research Associates International [producer]. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research, iPOLL [distributor], accessed Sep-17-2016.

Similarly, consumption of pornography can be embarrassing in itself as it can reveal sexual preferences or fantasies. However, the academic literature is thin on consumers' privacy expectations when consuming adult material online as well as the technical tracking that adult sites employ.¹⁰

There is reason for concern. Revelations of pornography consumption can have effects on people in profound ways. Professor Andrew Gilden recounts several examples of how pursuing online sexual fantasy has influenced real-world legal relationships, such as in custody battles, divorce proceedings, and where fantasy is used as propensity evidence in criminal trials.¹¹ The data can even be used for fishing expeditions. Recall that ten years ago, the Department of Justice issued subpoenas to several search engines in order to obtain user data about the prevalence of pornography online.¹² Google laudably fought the subpoena, but at the same time, the company would not characterize the search terms involved as "private" information.

Pornography consumption can also be used to infer other facts about an individual that could be used to extort or embarrass a person. One need only look to the fallout from the recent Ashley Madison breach, where tens of millions of users of the website for extramarital affairs was leaked online.¹³ The breach involved a public leak, where cybercriminals could attempt to extort large numbers of website members as well. However, private leakage and tracking also have a logical chain of harm. We found that nine of the websites we visited had a Google tracking script (DoubleClick or Google Analytics) and that seven

¹⁰ There is a literature on child pornography, feminist perspectives on pornography, access to pornography by children, the problem of "revenge" or non-consensual pornography, and the rights to view online pornography. On this last topic, see Jennifer Kinsley, *Sexual Privacy in the Internet Age: How Substantive Due Process Protects Online Obscenity*, 16 *Vanderbilt Journal of Entertainment & Technology Law* 103 (2013).

¹¹ Andrew Gilden, *Punishing Sexual Fantasy*, 58 *WILLIAM AND MARY L. REV.* ____ (forthcoming 2016) ("The Internet provides unprecedented opportunities to indulge in nearly all forms of sexual fantasy, and its relative anonymity can disinhibit discussions about even the most taboo topics. As a result, the transcript of a conversation on a fetish website might go into specific, painstaking detail about an elaborate kidnapping and cannibalism plot, and in the courtroom this transcript might look—without context—comparable to a wiretapped conversation about drug or weapons trafficking.").

¹² Neil Richards, *INTELLECTUAL PRIVACY: RETHINKING CIVIL LIBERTIES IN THE DIGITAL AGE* (2015).

¹³ Kristen v. *Brown, Scared, dead, relieved: How the Ashley Madison hack changed its victims' lives*, *FUSION* (Dec. 9, 2015).

leaked search terms to third parties and/or coded pornographic search terms into cookies. Some of these parties, particular Google, could trivially and secretly re-identify these users by relying on data collection from other sites. One could imagine the uses—for instance, would you sue a company for privacy violations if, in discovery, the company suggested that it could use its extensive activity logs to elucidate one’s web use at a public trial?

For these reasons, we focus here on the kinds of tracking and other privacy-relevant technical activity on websites featuring pornography.

Methods

In previous studies, we employed OpenWPM to perform large-scale crawls of the web to count HTTP, Flash, and HTML5 cookies.¹⁴ For this study, we performed both manual and automated crawls of all eleven adult websites that appear in the Alexa Top 500 US-ranked websites. For the manual and automated crawls, we directly typed in the target URL and selected three links that pointed to resources on the same domain. However, In the manual crawl, we did not click on videos.

The manual crawls were performed on Firefox in a clean state and documented with mitmproxy.¹⁵ For the automated crawls, we used Mezzobit, a cloud-based crawling platform that assesses privacy, usability, and website performance. The two methods reveal different data because mitmproxy tracks all connections during the browsing session, including popups and other resources that are loaded, while Mezzobit focuses more on the target URL. We compared our manual and automated crawl data, performed spot checks using different browsers, ran a separate analysis using Netograph,¹⁶ and supplemented these comparisons with Palantir Contour for link and statistical analysis.

Results and Discussion

Here we detail the high-level tracking dynamics on our sample of adult websites.

Third party tracking

Tracking dynamics are different on adult websites than popular non-pornographic sites. There are a relatively small number of third-party trackers

¹⁴ Ibrahim Altaweel, Nathan Good & Chris Jay Hoofnagle, *Web Privacy Census*, TECHNOLOGY SCIENCE. 2015121502 (Dec. 15, 2015), <http://techscience.org/a/2015121502>

¹⁵ A. Cortesi, M. Hils, T. Kriechbaumer, and contributors. mitmproxy 0.17. <https://mitmproxy.org/>

¹⁶ A. Cortesi, netograph, <http://netograph.io/>.

present that appear to specialize in pornographic ads. Adult sites generally lack “social buttons,” and just one site had a Facebook tracking script. This is remarkable because in 2015, we found Facebook on over half of top 1,000 most popular websites. AddThis and Twitter buttons were present on a small number of sites.

Mainstream behavioral advertisers are also not present on adult sites, aside from Google Analytics and DoubleClick. When using Mezzobit to analyze tracking on adult sites, the median site in our small sample of adult websites sent data to four third parties, with the highest sending data to ten other third parties and the lowest to just one third party. On average, sites made 25 separate communications with these third parties. However, it is important to note that Mezzobit does not measure all popup sites spawned by the adult sites.

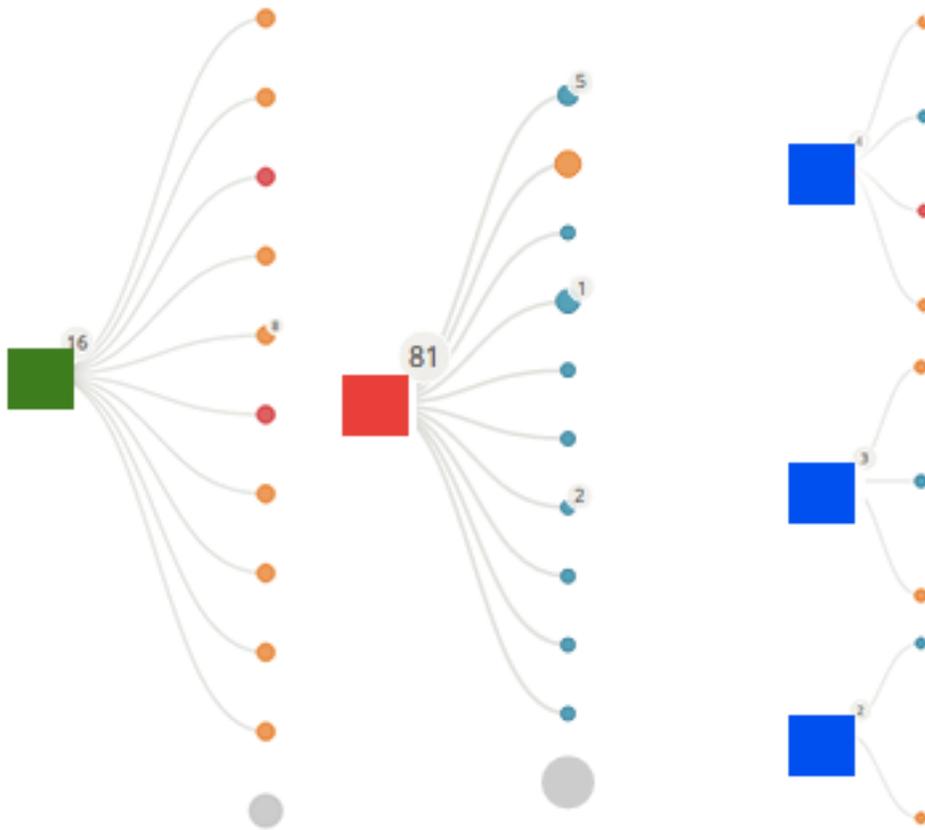


Figure 2: Mezzobit nicely illustrates website communication with third parties. Compare a popular medical information website (green square) that has 33 third party vendors and website comparably popular as the adult sites we analyzed (red square) that has relationships with 20. The blue squares represent typical popular adult sites--they have far fewer relationships with third parties, on average, four.

Turning to our analysis using mitmproxy, we find dramatically more third party tracking than average, but upon inspection, this increase is due to two websites that opened popup windows for the same third-party news website,

where there was much tracking (in one round of testing, the news website instantiated 448 third party cookies from 135 different third-party hosts).

What explains this absence of third party trackers on the adult sites themselves? One hypothesis is that pornography as a subject matter is simply too personal and creepy to track. In her survey of privacy policies in seven markets, Professor Florencia Marotta-Wurgler found that adult website privacy policies were more likely to comply with Federal Trade Commission recommendations in several respects than other, non-adult sites. Specifically, she found highly-ranked adult sites (n=17) to be more restrained in data collection and sharing, and to have shorter, yet more rights-protective privacy notices.¹⁷ The market for pornography could simply demand that adult sites are more private.

At the same time, if adult sites are too creepy to track, it stands to reason that the market would also pressure medical sites to limit tracking. However, as Figure 2 shows, a top 500 medical website sponsored by respected mainstream medical institutions has far more third-party tracking than any adult website we analyzed, as well as over 30 cookies.

Another more powerful explanation is that adult websites have low utility for non-pornography advertising. It could be that preferences for pornography simply do not have relevance for targeting non-pornographic ads. Reputation is an issue as well—mainstream advertisers probably do not want their products displayed next to pornographic content. Furthermore, as Dr. Kate Darling explains, there is a lack of trust in pornographic websites because of historical problems of unauthorized charges and malware.¹⁸ Medical websites on the other hand are more trusted, and offer more opportunities for mention of specific products.

Finally, adult websites compete with free resources, as observed by Professor Benjamin Edelman, and such free resources may be better from the consumer perspective because of concerns about fraud and monitoring.¹⁹

¹⁷ Florencia Marotta-Wurgler, *Understanding Privacy Policies: Content, Self-Regulation, and Markets*, NYU Law and Economics Research Paper No. 16-18 (Jan. 2016) (“... despite being several hundred words shorter than average, adult sites’ policies provide considerably more detailed notice of many critical privacy practices, and this translates to substantive protection for the consumer, as disclosures reveal little collection.”)

¹⁸ Kate Darling, *IP Without IP? A Study of the Online Adult Entertainment Industry*, 17 *STAN. TECH. L. REV.* 655 (2014).

¹⁹ Benjamin Edelman, *Red Light States: Who Buys Online Adult Entertainment*, 23(1) *JOURNAL OF ECONOMIC PERSPECTIVES* 209-220 (2009)

Obtrusive and obvious tracking might cause users to turn to peer-to-peer compilations of pornography free from cruft.

We did not examine browser fingerprinting, but a testing platform we used, Mezzobit, contains an automated fingerprint risk estimate score. According to its analysis, three sites were more likely than not fingerprinting browsers.

HTTP Cookies

Just 44 pages on eleven adult websites generated almost 1,100 cookies, yet we found that there are fewer cookies on adult websites than comparably popular non-adult websites. Our 2015 crawl found an average of 135 cookies (with a median of 91 cookies) on the top 1,000 most popular websites. Adult sites have fewer—the median is 19, but the average is 95. The high average is a reflection of the two sites that popped up a third party website. When those two sites are excluded, the average lowers to 33, because one of the two tracking-intensive sites had 465 cookies while and the other had 288.

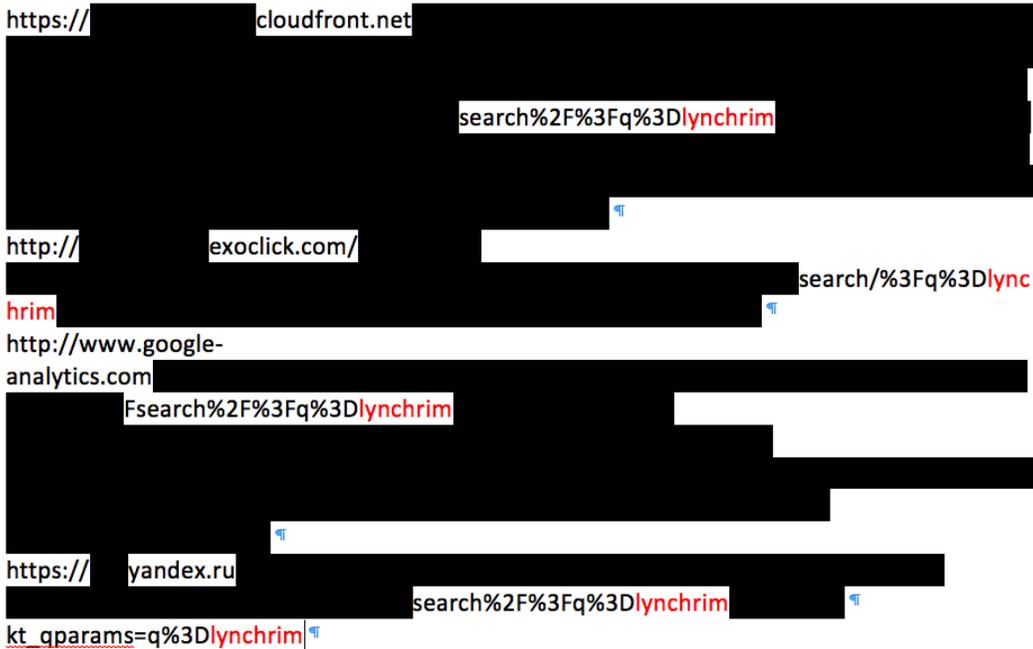
Websites had an average of ten (with a median of 11) first party cookies, 84 (median 8) third party cookies, and these third-party cookies were served on average by 25 (median 4) hosts. All of the third-party summary statistics are influenced by the two sites with many cookies, and as a result, the median values are more useful here.

Flash Cookies

We detected Flash on five of the websites. In most cases, Flash was being used to read HTTP cookie values, usually from the same domain. We found no evidence that Flash was being used to reinstate, or “respawn,” deleted HTTP cookies.

Search Term Leakage

We found that seven sites “leaked” search terms “in the clear.”²⁰ That is, if a user visited a site and performed a search, the search query was transmitted to third parties in plain text.



The image shows a network traffic capture with several entries. The first entry is a URL from a cloudfront.net domain containing the search term 'lynchrim' encoded in a URL parameter. The second entry is a cookie from exoclick.com containing the search term 'lynchrim' in plain text. The third entry is a URL from google-analytics.com containing the search term 'lynchrim' encoded in a URL parameter. The fourth entry is a URL from yandex.ru containing the search term 'lynchrim' encoded in a URL parameter. The search terms are highlighted in red text in the original image.

Figure 3: The search term "lynchrim" is leaked in URLs (red text) to Google and Russia-based Yandex, and sometimes encoded in plain text in cookies.

Additionally, the search term was often encoded into a cookie in plain text. Recipients of search terms included Google (both Analytics and DoubleClick), Russia-based Yandex, and other marketing and ad tech services. In addition to search terms, “category” tags were often encoded in plain text, meaning that a click on a specific interest (“blonde,” “trans,” and so on) were also transmitted in plain text rather than as a code (e.g. category “38273”).

HTTPS

In retail stores, pornography and sex toys are typically sheathed in a brown paper bag before leaving the store. Purchasers can pay cash and use the product in the privacy of their home. Protecting one’s privacy was as straightforward as being discreet in a context that ordinary people understand.

Turning to the internet, the online equivalent of the brown paper bag is HTTP over SSL. This technology would protect the pornographic content “in transit” from the source to the home.

²⁰ Bala Krishnamurthy, Konstantin Naryshkin & Craig Wills, *Privacy leakage vs. Protection measures: The growing disconnect*, 11(3) IEEE SECUR PRIV. 14 (2011).

We found surprising dynamics concerning HTTP over SSL. One might assume that adult websites would use HTTPS in order to limit ISP and other monitoring. However, only two of the eleven sites we tested used HTTPS by default for content delivery, and one provided HTTPS at the user's election. The remaining eight either would not load or would forward the user to HTTP if the website URL was entered with HTTPS manually.

Additionally, we found that various individual communications used HTTPS on almost every site, but more likely than not, the communications secured by HTTPS were those of third party trackers and advertising delivery companies. That is, on these sites, we found that tracking efforts and ad delivery were sometimes transmitted over HTTPS, while adult content was delivered over HTTP.

We think this lack of HTTPS is an important privacy problem that users may not understand. Many intermediaries, be it the WLAN operator or intelligence agencies, can view preference and even second-by-second decisions about consumption. The user may feel private, because one can view these sites in seclusion, protected by the walls of the home. In reality, the user is consuming the product before a one-way mirror.

Local Storage

We did not encounter any use of HTML5 local or session storage.

Conclusion

In this brief paper we explored the kinds and of amount of tracking on popular adult websites.

Adult websites have a smaller universe of trackers than popular non-adult websites. This surprising lower level of tracking is best explained by the marketability of adult websites. Mainstream advertisers do not want their content placed next to pornography, nor are they likely to be able derive actionable marketing intelligence from users' specific preferences for adult material.

However, several privacy risks are present: search terms and category tags, which may reveal sexual fantasy, are leaked in most cases in the clear and to third parties. Furthermore, just a handful of sites use HTTPS, leaving full URL strings visible for monitoring by others.

The academic literature is thin on privacy expectations in porn consumption. Our work complements Professor Marotta-Wurgler's on privacy policies, in that it adds technical analysis of adult websites' functioning that are consonant with her findings of adult sites' stated policies.

In a world with criminal prosecution for pornography possession, of large-scale leaks of private information, and of growing extortion attempts based on

personal information collected online, our results point to the need for careful consideration of whether consumption of pornography, a popular, yet unsympathetic web activity, deserves attention from consumer protection authorities.