

# Indications of Child Sexual Abuse Revealed in App Store Reviews

Brian Levine<sup>1</sup>, Jagath Jai Kumar<sup>1</sup>, Hany Farid<sup>2</sup>, Edward Dixon<sup>3</sup>, Eloghosa Ikponmwoba<sup>3</sup>

<sup>1</sup> *College of Information and Computer Sciences, Univ. of Massachusetts, Amherst*

<sup>2</sup> *Electrical Engineering & Computer Sciences, Univ. of California, Berkeley*

<sup>3</sup> *Rigr AI*

## 1 Introduction

As of 2019, 84% of teenagers possessed their own phone, with more than 50% possessing a phone by age 11 [1]. Apple’s and Google’s app stores are the anchor points for all major apps that can be downloaded to iOS and Android, with each company taking a 15–30% cut of all app purchases. Despite being accessible by so many children, the Apple and Google app stores contain little information for children who need help avoiding sexual exploitation.

Many sexual exploitation crimes against children are based on connections formed via apps vetted and distributed on Apple’s and Google’s app stores [2–4], leading to extreme and life-long harm [5, 6]. Some apps allow perpetrators to meet minors and begin a grooming process where an inappropriate sexual relationship is normalized. Too often, child sexual abuse material (CSAM) is recorded and shared online for years [7].

In aggregate, user-generated reviews can be a valuable source of information regarding the dangers posed by an app. For example, a 2019 study [8] revealed that Apple’s App Store received over 1,500 complaints against just six apps: Monkey, Yubo, ChatLive, Chat for Strangers, Skout, and Holla. All six apps are still available on Google’s Play Store; Monkey, ChatLive, and Chat for Strangers are no longer available on Apple’s app store. Meet24, FastMeet, and Meet4U were removed from the Apple App Store for dangers posed to children, though the removals came only after public comments by the US Federal Trade Commission [9]. Google took no action against the apps, and they are still available on the Google Play Store.

Shown in Table 1 are representative examples of reviews highlighting child safety. For many parents, even a single

review like this might be sufficient for avoiding an app.

Our analysis of 289 social-interaction apps reveals a much larger number of dangerous apps on Apple’s and Google’s app store. Our analysis employs a combination of machine learning and manual review to analyze over 3.2 million user-generated reviews spanning a decade. Across 100 apps, we found over 4,300 reports of sexual abuse, predators, exchange of CSAM, and inappropriate sexual interactions between minors and adults. For just these apps, the number of reports grew to approximately 80 per month, and appear to be unaddressed by Apple and Google.

We describe our methodology for exposing these dangerous apps along with the construction of a website for hosting the results of our analysis. It is our hope that this analysis and website will make parents, children, and policy makers aware of the dangers posed by certain apps, and be a wake up call to Apple and Google to address the dangerous apps on their store.

## 2 Methods

To increase the visibility of reports of child sexual abuse in social-interaction based apps, we scraped reviews from the Apple and Google stores, searching for reviews highlighting potential dangers of child sexual exploitation.

We categorized a review as positive if it suggests any of the following.

- CSAM is shared/offered.
- Children are subject to sexual requests by adults.
- Children are acting sexually or sharing sexual content.
- Adults are acting inappropriately to other adults (e.g., selling nudes/cams, sexual harassment/abuse).

We categorized all other reviews as negative. We recognize that our definitions are relatively conservative and some parents may consider other in-app behaviors equally or more dangerous.

Given our limited computational resources and Apple’s throttling of review downloads, we could not scrape all app

<p><b>Hoop</b> (Sep 21, 2020) This app has a serious child porn problem— I got through this app and every 10 posts is and advertisement for child pornography. I have reported like 30 of these accounts and they still keep coming. Real messed up</p>
<p><b>Hoop</b> (Jun 16, 2020) I like the app for getting a couple good new friends but there is alot of sexbots and ChildPorn is being advertised by certain accounts. This needs to be fixed before I think of coming back at all.</p>
<p><b>MeetMe: Chat &amp; Meet New People</b> (Jan 17, 2021) Wow... so for the last week I was doing an experiment for my college, I used an old friend s picture and went on here posing as teenage girl... unreal how many people asked for sex, wanted nude pics and even when I said I'm underage, 'age is only a number'" eventually the account was blocked but man I never experienced something like that before. Nudes, sex talk, it was horrible one star"</p>
<p><b>MeetMe: Chat &amp; Meet New People</b> (Oct 28, 2020) This app is harboring and protecting child and women sex traffickers. I saw a woman being forced to webcam on their live services in a completely cement room with only a bed. I beg the state department and FBI to investigate this app</p>
<p><b>Whisper</b> (Jan 27, 2021) grew up using this app and got groomed by paedophiles very often. im almost 18 now an its been happening since i was 12</p>
<p><b>Whisper</b> (Jan 1, 2021) There is no moderation on this app and there are men being obviously abusive towards women and no way to remove them from groups. It feels very unsafe and when people are threatening women and others with all sorts of violence and no way to do anything about it. There are also a bunch of predators looking for underage victims that even when reported dont get banned. This is not a SAFE anonymous sharing app like it was in the beginning.</p>
<p><b>Gacha Life</b> (Dec 17, 2020) This game is filled with pedophiles</p>
<p><b>Gacha Life</b> (Dec 27, 2020) The Game overall is good, graphics, quality, characters, etc. However the children on YouTube, Twitter, other socials have been posting above 13+ content, whether it's nude people, rape/insensitive topics. One day I was scrolling through my tiktok and found a very disturbing video, where there were was a character nude and was in 'heat'". Which then lead me to browsing more about it, turns out there are multiple videos of 13+ content. Please parents be aware of what your kids download."</p>
<p><b>Plenty of Fish</b> (Feb 6, 2021) Sites run by paedophiles and they endorse rapist. Congratulation Plenty of Flops your admin team are a bunch of brainless fk tards you have blood on ur hands..</p>
<p><b>Plenty of Fish</b> (Oct 12, 2020) Website/app allows prostitution. They can say they don't but there is profile after profile of young women with their butts out saying clearly on their page they charge money for sex AND THIS SITE WILL DO NOTHING ABOUT IT. Sounds like a pimp to me. Time to get rid of them. Report after report on their site and they do nothing. Try to upload a picture of your dog it gets taken down but a 19 year old uploads a picture of her ass and it's fine and her page even says money for meet ups.</p>
<p><b>Skout</b> (Feb 8, 2021) Excessive amount of scams and illegal activity. Easily out of 100 profiles only 1 may be a real person who isn't begging for money, prostitution, drugs, pedophilia etc. They do absolutely nothing to stop it, they also ban your account if you submit negative feedback.. This app absolutely needs to be taken off the store and developers investigated.</p>
<p><b>Skout</b> (Jan 18, 2021) They do absolutely nothing to curb fake profiles, scams, pedophiles, prostitutes and drugs.</p>

Table 1: Representative app-store reviews highlighting child safety concerns.

store reviews. We started with a subset of apps based on previously stated concerns (e.g., Hargove et al. [3]). We added to this subset based on Apple's recommendations. In particular, if an app had reviews with words on a custom keyword list (e.g., "child porn" or "pedo"), then we also scraped apps

recommended by Apple as related. We also included a large selection of popular apps unrelated to minors (e.g., LinkedIn). Overall, this yielded 289 apps in total.

Consisting of three basic steps, our analysis employs a combination of automated and manual review, as follows:

1. We scraped reviews for the 289 apps from the Apple and Google stores (most apps appeared in both stores), yielding a total of 3,246,581 reviews. The reviews spanned a decade, but the vast majority of reviews are recent.
2. We manually annotated 6,400 reviews as indicating dangers of child sexual exploitation or not using the criteria above.
3. We trained a classifier on the manually annotated reviews, described below.
4. We automatically classified the unlabelled scraped reviews.
5. We manually reviewed 6,482 reviews classified as positive to filter out false positives.

We used three reviewers in Step 2 and four reviewers in Step 4, who were all trained to help increase inter-reviewer accuracy.

Our model is based on a BERT [10] encoder pre-trained (in English) on Wikipedia and BooksCorpus. This encoder yields a compact embedding used as input to a neural-network classifier with three fully connected layers. The primary challenge in training and evaluating a classifier is the low prevalence of comments matching the two highest-severity classes in our ontology (our manual labeling campaign yielded only 994 high-consensus training samples). Our classifier achieved an  $F_1$  score of 0.71 on a validation set. This classifier is not sufficient to automatically classify all reviews, but is an effective triage for subsequent manual review.

Overall, we found 2,366 additional positives from our review of 6,400 positives. The performance of the classifier was worse than in earlier testing. An exploration of the false positives revealed that our training samples were likely not from a sufficiently broad set of apps. If building on this work, we also suggest using the largest available pre-trained model.

### 3 Results

Out of the 289 apps: 42 had 10 or more reviews indicating abuse; 100 had at least 1 such review; and 189 had no such reviews. Shown in Figure 1 are the total number of surfaced reviews indicating child exploitation reported between 2011 and 2021 for the Apple (top) and Google (bottom) stores. These results show a troubling increase in the number of reports of abuse over the past decade. This total count is almost certainly a lower bound on the actual number of reports given our limited number of scraped apps. As shown in the figure, abuse of minors is reported for apps with a variety of age ratings.

Shown in Figure 2 are the total number of surfaced reviews for the 12 apps with the most number of reports indicating child exploitation. Although the number of apps vary from year to year, they are consistently present across many years. The consistency of these reports makes it difficult to understand why the app makers and app stores do not respond more aggressively to threats to children.

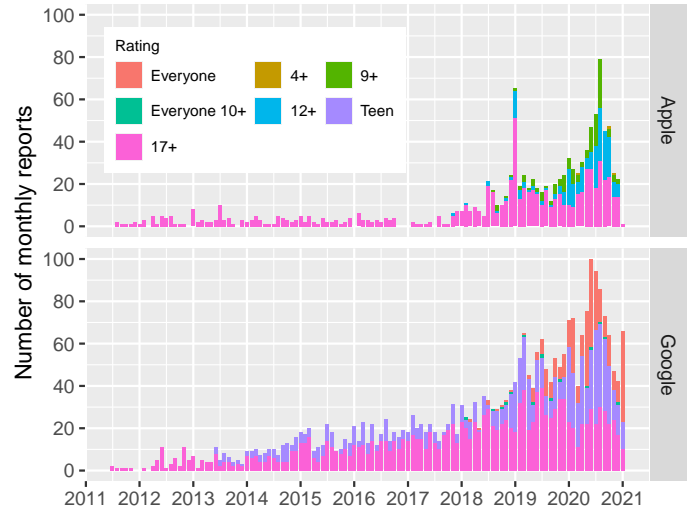


Figure 1: Monthly reports of identified danger in each app store. The values in the plots are stacked by age rating.

As a fraction of all written reviews per app, we find that, when present, reports of danger to children are small (1–3%). This infrequency further highlights the difficulty in organically finding problematic reviews, and the value of the type of analysis performed here.

While we, of course, do not have the ability to evaluate the veracity of each report of danger, we believe that in aggregate, multiple reports should, at a minimum, be of interest to parents and child-safety organizations. These reports should alert the app’s hosts, Apple and Google, (and the Federal Trade Commission) to launch an investigation into potentially problematic apps.

We made all scraped and classified reviews available at <http://appdanger.org>. We intend to keep the website updated with newly scraped reviews indicating abuse as we find them.

### 4 Discussion and Recommendations

App Stores exert a great deal of control over who can publish to their stores and what can be published [11]. And even if they didn’t, we would contend that Apple and Google have a responsibility to act on user-generated reports of abuse on their app stores, for apps hosted on their operating systems, and from which they profit.

Given the relative ease that we — outside of the store ecosystem — were able to surface thousands of troubling reports of child sexual exploitation, there is no reason why the stores themselves should not be aware of these reports.

The extent to which Apple and Google respond to user-generated reports is unclear. After public criticism for promoting hate, for example, both took action against Parler demanding better moderation of the content on its app [12]. We are unaware, however, if similar actions are consistently taken



Figure 2: Quarterly reports of danger per quarter for a representative set of 12 apps. The parenthetical value associated with each app is the total number of problematic apps.

against other apps that promote illegal activity or violate the stores’ terms of service. And, given, the large number of reports of child exploitation over hundreds of apps spanning over a decade, it seems likely that user-generated reports are not being consistently analyzed by Apple or Google.

In addition to serving as better care takers, we also contend that the app stores should make it easier to surface problematic reviews so that parents can make more informed decisions about the appropriateness and safety of apps.

Apple and Google should release quarterly public transparency reports detailing the number of app reviews highlighting concerns of child sexual abuse. They should also require that app developers who sell on their stores to similarly publish transparency reports detailing the number of reports of child sexual abuse they received directly from users, and the number of CyberTips sent to law enforcement. This information should be made available alongside the apps and easily surfaced. Apps that are not looking for abuse, or cannot look for abuse due to their use of end-to-end encryption, should also be made apparent to parents so they are not lulled into a false sense of security.

## Acknowledgements

This work was funded in part by Oak Foundation. We are grateful to Peace Ekeinde, Alyssa Lanter, Authority Michael Odili, Onyiriuka Ngozi Palmer, and Jerry Zeng for their help with labeling.

## References

[1] A. Kamenetz, “It’s a smartphone life: More than half of u.s. children now have one.” National Public Radio, October 31 2019.

[2] U.S. Dept. of Justice, “The National Strategy for Child Exploitation Prevention and Interdiction: A Report to Congress.” <https://www.justice.gov/psc/file/842411/download>, April 2016.

[3] D. Hargrove, R. Avitabile, N. Hsu, M. Dean, and T. Jones, “11 lesser-known apps that experts say could expose your child to sex trafficking.” NBC San Diego, August 2020.

[4] Federal Bureau of Investigation, “Sextortion: An online threat to kids and teens.” <https://www.fbi.gov/scams-and-safety/common-scams-and-crimes/sextortion>.

[5] M. Cutajar, P. Mullen, J. Oglloff, S. Thomas, D. Wells, and J. Spataro, “Psychopathology in a large cohort of sexually abused children followed up to 43 years,” *Child Abuse & Neglect*, vol. 34, no. 11, pp. 813–822, 2010.

[6] G. Pérez-Fuentes, M. Olfson, L. Villegas, C. Morcillo, S. Wang, and C. Blanco, “Prevalence and correlates of child sexual abuse: a national study,” *Comprehensive Psychiatry*, vol. 54, no. 1, pp. 16–27, 2013.

[7] G. Bissias, B. N. Levine, M. Liberatore, B. Lynn, J. Moore, H. Wallach, and J. Wolak, “Characterization of Contact Offenders and Child Exploitation Material Trafficking on Five Peer-to-Peer Networks,” *Child Abuse & Neglect*, vol. 52, pp. 185–199, 2016.

[8] R. Albergotti and A. Johri, “Apple says its App Store is ‘a safe and trusted place.’ We found 1,500 reports of unwanted sexual behavior on six apps, some targeting minors.” *Washington Post*, November 2019.

[9] Federal Trade Commission, “App Stores Remove Three Dating Apps After FTC Warns Operator about Potential COPPA, FTC Act Violations.” <https://www.ftc.gov>, May 6, 2019.

[10] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: pre-training of deep bidirectional transformers for language understanding,” 2018.

[11] Reuters, “Epic Games Asks Court to Prevent What It Describes as Apple’s ‘Retaliation’.” <https://www.reuters.com/article/us-apple-epic-games-idUSKBN25W0HE>, September 2020.

[12] J. Nicas and D. Alba, “Amazon, Apple and Google Cut Off Parler, an App That Drew Trump Supporters.” *New York Times*, January 2021.