Understanding Targeted Video-Ads in Children's Content

Flavio Figueiredo
Universidade Federal de Minas Gerais

Felipe Giori
Universidade Federal de Minas Gerais

Guilherme Soares
Universidade Federal de Minas Gerais

Mariana Arantes
Universidade Federal de Minas Gerais

Jussara M. Almeida
Universidade Federal de Minas Gerais

Fabricio Benevenuto
Universidade Federal de Minas Gerais

ABSTRACT

As the volume of online video entertainment via streaming increases, ever so more are users targeted by online advertisement algorithms. Nevertheless, this rise in targeting and revenue does not come without any concerns. That is, even though the online advertising business model has is very successful, nowadays, rising societal concerns regarding the ethics and extent to which such algorithms agree with the laws of different countries are also present. Motivated by the dichotomy above, we here explore how targeted video-ads meet the regulatory policies regarding children advertising in Brazil and Canada. To perform our study, we create synthetic user personas that watch YouTube videos daily. Our personas are tailored to stream children's content while controlling for several variables (e.g., gender, country, and type of content streamed). With the data gathered, our analyses reveal statistical evidence of algorithmic targeting in videos geared towards children. Also, some of the advertised products (e.g., alcoholic beverages and fast-food) go directly against the regulations of the studied countries. With advertisements being matched to users by machine learning algorithms, it is impossible to state whether regulations are not followed on purpose (e.g., advertisers gaming the system). Nevertheless, our findings and discussion do raise a flag that regulations may not be sufficient, and content providers may still need to audit systems to meet the regulations.

CCS CONCEPTS

• **Information systems** → **Social advertising**; • **Social and professional topics** → *Children*.

KEYWORDS

video streaming, advertising, ethics and law, children

ACM Reference Format:
Flavio Figueiredo, Felipe Giori, Guilherme Soares, Mariana Arantes, Jussara M. Almeida, and Fabricio Benevenuto. 2020. Understanding Targeted Video-Ads in Children's Content. In *Proceedings of the 31st ACM Conference on Hypertext and Social Media (HT '20), July 13–15, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3372923.3404787

HT '20, July 13–15, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-7098-1/20/07…$15.00
https://doi.org/10.1145/3372923.3404787

1 INTRODUCTION

In this age of online marketing in which we live in, human choices are usually complemented or guided by algorithms. Take, for instance, the moment whenever a user logs on to a social website such as YouTube. At this time, both recommendation and online advertisement algorithms target the user. Whereas the former aims at suggesting content that entertains and keeps the user online, the ultimate goal of the latter is to sell a product/service or an idea. That is, although advertisements may be entertaining, they exist to subsidize the Internet economy by maximizing profits for content providers, marketers, and brands.

Even though the ultimate goal of advertising may be one of maximizing profits, different laws regulate the online advertising space. These laws will counterpoint the goal of profit maximization. For instance, Internet services, such as YouTube, are subject to the regulatory terms, guidelines, and laws of individual countries, which may limit the ads shown to the users. An interesting setting arises when children access video streaming services. Different countries, being Brazil and Canada two examples, have policies that aim to limit the kind of content that is advertised to children both online and offline. One tool employed to guarantee that children are not subject to online ads is a limit on the minimum age required to create user accounts[1]. However, the effectiveness of such actions is questionable since it is well known that, regardless of such limits, social media usage is strong among infants and teenagers [25, 30].

We emphasize that even specialized children's apps, like YouTube Kids[2], ultimately rely on advertising for a sustainable business model[3]. Such apps present a recent attempt to filter out unacceptable content. In the end, regulations limiting advertisements to children create a conflict for content providers. That is, underage audiences are significant targets for different brands and services [26, 31]. However, advertising to children is not universally socially, or even lawfully, acceptable [39].

To provide more transparency on how services follow such regulations, we here investigate how video advertisements on YouTube target children's content consumers via online ad-auctions. Our analysis explores the *potential* of real children being targeted by online advertising. Our study is guided by the following two research questions: **RQ1 (1-a)** Are children's content consumers being targeted by video-ads on YouTube? **(1-b)** If so, does user behavior (based on different personas) affect the advertisement shown to the user? **RQ2 (2-a)** Which content, brands, and products target children's content consumers on YouTube? **(2-b)** Are their ads possibly contradicting official country-wide guidelines?

[1]https://support.google.com/accounts/answer/1350409. URLs last accessed in June 2020.
[2]http://kids.youtube.com
[3]https://support.google.com/youtube/answer/6168681

In order to investigate those questions, we created synthetic user personas [11] that consume YouTube videos daily. Personas correspond to Google accounts whose preferences are defined via browsing behavior. Specifically, our personas browsed YouTube videos focusing on children's content from popular YouTube channels. To mimic both users with the youngest age allowed by the system (13-year-old) and older users who consume children's videos, we created different accounts with different settings. The latter aims at capturing children who may use a home computer accessing YouTube with their parents' accounts or simply parents who stream videos to their kids. While watching videos, our personas log the video advertisements shown to them. Synthetic personas are valuable tools as they allow us to control not only preferences via browsing but also other variables such as age and country via sign-up information. In particular, our personas were deployed on a cloud provider and consume videos in two different countries, our case studies, Brazil and Canada.

Our results on RQ1 unveil that, statistically speaking, the streamed videos are targeted for personalized advertisements based on the persona profiles. That is, personas with different browsing traits were exposed to different video ads. This effect occurs even when to personas watch the same video. To tackle RQ2 we analyzed the brands and products that advertise products to personas mimicking children behavior. The products sold to them vary from popular toys, clothes, fast-food, and even alcoholic beverages. We further show how some video advertisements explore multimedia content and sell products that go against country-wide regulations.

2 BACKGROUND

In this section, we review related work on online ads, video ads, and children's behavior when exposed to marketing (Section 2.1). We then briefly discuss advertisement guidelines in the two countries of our study, Brazil, and Canada (Section 2.2).

2.1 Prior Studies

Advertising to children is a complex and controversial subject. Even when the content of an ad follows regulations [39], there still exists empirical evidence that the message of the ad may affect the behavior of children [10, 26, 31, 35]. One example of such findings is the work of Aktacs et al. [10]. Here, the authors studied the presence of well-known brand characters/mascots and how they raise the awareness of brands in children. In their results, the authors show how children with as little as 3 years of age are already able to recognize brands. Similar findings have been presented by Oates et al. [24] and Valkenburg and Buijzen [34]. Complimentary, Opree [26] used a causal model to show that advertisement exposure indeed increases materialistic values (consumerism) in infants of age 8 to 11. Such findings are interesting as it shows how advertisements are able to guide infants to different social values. Even though such studies looked into advertisements in an offline setting, they produced cautionary findings which may translate to online spaces.

In the context of online advertising, most prior efforts focused on online bidding algorithms and how to improve the success of ad campaigns [13, 23, 32, 36, 40]. Some studies have also discussed the negative experience of users when exposed to annoying and intrusive ads [17, 22, 28, 38]. Yet, in general, they have not looked into how such experiences relate to country-wide guidelines and restrictions, as we here do.

Only recently, there have been some investigations on fairness and transparency issues of online advertising [5, 9, 14, 33, 38]. In particular, the work of Araujo et al. [9] characterized the presence of brands on videos geared towards children on YouTube.

Focusing specifically on online *video advertising*, we are aware of only a few, recent, papers that have studied YouTube video advertisements. In particular, Arantes et al. [6–8] studied user consumption of YouTube video ads based on web traffic and monetization analysis as well as user centric studies. Other authors have looked into caching and traffic properties of video-ads on mobile devices [3]. Our work complements these previous endeavors by looking into aspects regarding regulation and potential guideline infringements.

Finally, several efforts have looked into advertising using synthetic personas as we do. In particular we cite the work of Carrascosa et al. [11] and Le et al. [21] have employed the use of personas to understand search engine advertising. These endeavors are the basis for our justification in using personas as this kind of study is able to unveil advertising biases [21]. These two efforts however do not focus on video advertising as we here do.

2.2 Regulations and Guidelines

Regulations regarding advertisements are a non-trivial matter. Usually, multiple laws, codes of conduct, and guidelines determine whether an ad can is an infringement. For instance, most countries have laws against discrimination of any form (online or not). In this sense, discrimination laws should impact what can be exhibited in an ad both online and offline. Thus, one must take several sources into account in order to understand what kind of content can be exhibited, in particular to children (in a legal sense). We here summarize some of the guidelines regulating advertisement to children in two specific countries: Brazil and Canada. We selected these two countries as case studies due to our computational resources. Our experiment was deployed on a Azure cloud. On this environment, we were able to create one unique IP address per persona country. Moreover, both Brazil and Canada have clear regulations regarding children's advertising.

In legal terms, advertisements in Brazil are regulated by several different counsels and institutions. The most important ones are the Federal Constitution[4], the Consumer Protection Code[5], and the Code of Advertisement Regulations[6]. Associations, such as the Brazilian Association of Advertising Agencies (ABAP)[7], also provide documents with summarized guidelines for advertisement agencies. The ABAP guidelines are presented below. As stated, these guidelines apply to children, whose formal definition in Brazil is someone who is under 13 years of age[8].

(1) Imperative verbs (e.g., buy, purchase) should not be employed in advertisements;
(2) Advertising children's products (e.g., toys) is banned on open television;

[4]http://english.tse.jus.br/arquivos/federal-constitution
[5]http://www.planalto.gov.br/ccivil_03/LEIS/L8078.htm
[6]http://www.conar.org.br
[7]http://www.abap.com.br/
[8]Minors over 13 and under 18 are legally considered teenagers, whereas anyone over 18 is considered an adult.

(3) No ad should employ well-known characters (e.g., cartoon characters from famous TV shows);
(4) Ads should not induce fear or any other ill feeling;
(5) Ads should encourage healthy eating habits;
(6) Similarly, advertisements should not promote products which replace daily meals;
(7) Ads should not discriminate in relation to race, gender or religious beliefs.

Similarly to Brazil, advertisements in Canada are regulated by different counsels and institutions, such as the Canadian Radio-television and Telecommunications Commission (CRTC)[9], Industry Canada[10], and Health Canada[11]. To develop our work, we present an overview on Canadian guidelines based on two documents from the Canadian Code of Advertising Standards, the Children's code in particular[12], as well as documents from Health Canada[13].

One interesting point in Canadian documents is that the Children's Code, for instance, determines that a child is someone with less than 12 years of age. This definition is similar to the Brazilian one (13 years). To reduce unhealthy eating habits across several ages, health Canada argues that even teenagers (someone whose age is less than 17 or 18 years of age, depending on the Canadian province) need restrictions with regards to junk food ads. Overall, regulations in Canada are similar to Brazil, stating that (when appropriate, we have shortened the wording or summarized the presented guidelines):

(1) No children's advertising may employ any device or technique that attempts to transmit messages below the threshold of normal awareness;
(2) Puppets, persons and characters (including cartoon characters) well-known to children and/or featured on children's programs must not be used to endorse or personally promote products, premiums or services;
(3) Ads must not directly urge children to purchase or urge them to ask their parents to make inquiries or purchases;
(4) Children's advertising must not encourage or portray a range of values that are inconsistent with the moral, ethical or legal standards of contemporary Canadian society;
(5) Ads may not imply that possession or use of a product makes the owner superior or that without it the child will be open to ridicule or contempt.

These last two guidelines are further detailed in terms of food products. The code states that: "Advertising of food products should not discourage or disparage healthy lifestyle choices or the consumption of fruits or vegetables, or other foods recommended for increased consumption in Canada's Food Guide, and Health Canada's nutrition policies." Finally, it is interesting that both countries share similar guidelines focused on: healthy eating, reduced exploitation of well-known characters, and materialistic values.

[9]https://crtc.gc.ca
[10]https://www.ic.gc.ca/
[11]https://www.canada.ca/en/health-canada.html
[12]http://www.adstandards.com/en/clearance/childrens/broadcastcodeforadvertisingtochildren.aspx
[13]https://www.canada.ca/en/health-canada/services/publications/food-nutrition/restricting-marketing-to-kids-what-we-heard.html

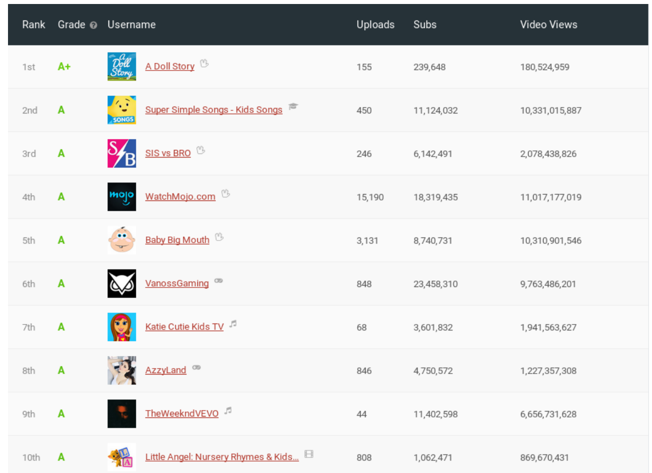

Rank	Grade	Username	Uploads	Subs	Video Views
1st	A+	A Doll Story	155	239,648	180,524,959
2nd	A	Super Simple Songs - Kids Songs	450	11,124,032	10,331,015,887
3rd	A	SIS vs BRO	246	6,142,491	2,078,438,826
4th	A	WatchMojo.com	15,190	18,319,435	11,017,177,019
5th	A	Baby Big Mouth	3,131	8,740,731	10,310,901,546
6th	A	VanossGaming	848	23,458,310	9,763,486,201
7th	A	Katie Cutie Kids TV	68	3,601,832	1,941,563,627
8th	A	AzzyLand	846	4,750,572	1,227,357,308
9th	A	TheWeekndVEVO	44	11,402,598	6,656,731,628
10th	A	Little Angel: Nursery Rhymes & Kids...	808	1,062,471	869,670,431

Figure 1: Social Blade's Ranking for Canadian YouTube

3 METHODOLOGY

We start by defining two terms used throughout the paper. A **video content** is defined as a YouTube video a user is able to stream using a browser. In contrast, a **video advertisement**, or **video ad**, consists of videos streamed to the user without his/hers explicit request in the form of an advertisement. Our work focuses on video ads which are paired with video contents through online bidding.

3.1 Persona Based Crawling

In order to collect both video ads and contents we developed synthetic users (e.g., personas) that watched YouTube videos. Within the scope of this paper, a persona is implemented via a bot that watches YouTube videos. Such bots were developed by automatizing a Firefox[14] browser via Selenium[15]. Our code was executed on Azure Cloud using virtual machines deployed in Brazil and Canada.

In marketing terms, a persona is defined as a set of user characteristics which may be exploited for advertising purposes [2]. To create marketing personas, researchers explore a wide range of tools and source datasets [18]. However, marketing personas are usually created offline (here to be interpreted as the opposite of real-time). Online advertising changes this perspective as user profiles are created by browsing behavior [11]. In this sense, our bots simulate marketing personas, thus the same name, by setting up novel Google accounts and exploring browsing history as well as account features. Our accounts were set up having residence in the country where the bot was deployed and we also controlled for age variables. Prior to execution, there was no browsing history in such accounts.

To *train personas* automatically, in other words to create preference profiles for our bots compatible with children's content consumers, we crawled Social Blade[16] to determine the top children channels in Brazil and Canada. Social Blade is a website that ranks YouTube channels based on a combination of several popularity scores. As an example, Figure 1 shows a snapshot of Social

[14]https://www.firefox.org
[15]https://www.seleniumhq.org
[16]https://socialblade.com/

Table 1: Children's Channels in Brazil/Canada. Statistics from Social Blade collected on March 26th 2018.

Brazil	#Subs.	#Views
Galinha Pintadinha	10.41M	7.91B
Turma da Mônica	6.51M	5.8B
TotoyKids	8.22M	4.37B
Patati Patatá	2.57M	1.92B
Planeta das Gêmeas	6.51M	1.64B
Erlania e Valentina	6.04M	1.3B
O Reino das Crianças	1.93M	1.25B
Kids Fun	5.64M	1.23B
Bia Lobo	1.5M	404M
Totoykids Explorer	785K	328M
Canada		
Super Simple Songs - Kids Songs	9.15M	8.95B
Big Bady Mouth	7.56M	7.68B
Katie Cutie Kids TV	2.5M	1.51B
Awesome Toys Collectors	839K	603M
Little Angel: Nursery Rhymes	616K	514M
Playtime4kidz	814K	433M
The Kiboomers - Kids Music	392K	388M
HZHtube Kids Fun	2.86M	313M
Barney	228K	198M
FamousTubeKIDS	710K	89M

Blade's ranking for Canadian YouTube channels (taken on August 17^{th} 2018). The website ranks channels using a proprietary algorithm exploring metrics such as numbers of subscribers (users who receive updates from the channel), video uploads and video views. Using such metrics, Social Blade assigns an overall influence score or rank in Figure 1. Although the algorithm is proprietary and closed source, our goal in using the service was simply to identify popular children channels (we do not make use of any other feature of Social Blade). In particular, we explored a set of ten channels per country, each with a strong focus on uploads targeting children.

To identify children channels, two coders (authors) browsed Social Blade and coded channels as either: Children's Content or Not Children's Content. Coders browsed channels following the website's overall ranking (see Figure 1) from top to bottom. Moreover, the coding was performed until both coders agreed on ten channels per country (Brazil and Canada). We emphasize that the coding was performed simultaneously by both coders and there was no disagreement in any channel.

Table 1 presents the names of our selected channels along with the numbers of subscribers and views at the time of selection. Note that the selected channels attract from hundreds of thousands to millions of subscribers, indicating their significant popularity.

Given the selected channels, personas were trained by having the bots watch the channels on a regular fashion (more below). Thus, our code streamed video contents and relied on YouTube to create the preference profile. We also considered the case of users who do not only browse children's contents (e.g., parents who stream content to their children but also request content of their own interest), and developed personas that would also watch popular music videos. That is, with 50% chance, these personas would randomly watch a popular music video. Other control variables in our study are persona age, country and skipping behavior (i.e., whether the persona skips video ads it is exposed to or not).

As discussed, Google allowed us to determine age and country when setting up accounts. The age of created user accounts was set to either 13 (minimum allowed) or 40 years of age. We refer to them as *child* and *adult* personas, respectively. It is important to point out that even though we create *adult* personas, they aim at capturing parents who stream videos to their children, or similarly children who share devices with parents at home. The country of the accounts was chosen to be either Brazil (BR) or Canada (CA). Similar to the content variable, skipping behavior was controlled by setting one group of personas to skip video ads, interrupting their streaming and jumping to the video content, as soon as possible. A second group would stream ads till completion.

In total, 16 personas, generated by the combinations of age (child/adult), type of content streamed (children/mixed content), skipping behavior (skip/never skip) and country (BR/CA), were deployed on Azure. Given our control variables, we are able to infer whether there are differences in the video ads streamed across personas which may unveil evidence of targeting, as discussed in the next section. We refer to each persona according to the variable configurations. That is, BR-Child-Children-Skips indicates a Brazilian persona (BR), under 13 years of age (Child), who streams children videos (Children), and always skips video ads (Skip).

We developed our personas to be the least intrusive with YouTube system possible. That is, synthetic personas were mere bystanders that **only** watched videos. Also, we attempted to mimic common human daily patterns by executing each persona for one hour in three different time periods: 9h to 10h (morning), 15h to 16h (afternoon), and 19h to 20h (night). Our personas watched videos from March 26^{th} to May 25^{th} 2018. During each hour, personas would stream random videos from the selected channels, as many as possible[17]. This leads to repeated streams, a frequent behavior on social websites as YouTube [4].

While streaming, personas would be subject to video advertising on YouTube. For every video content and video ad encountered, the persona would log the id and collect the metadata for each content using YouTube's API[18]. One particular persona in our study (Brazilian child streaming mixed content with no skipping behavior) did not execute properly as Google would frequently log-out this persona. To have a balanced dataset, this single persona was redeployed from October 3^{rd} to November 10^{th} 2018. In the next subsection we present an overview of the collected data.

We programmed our personas to follow patterns uncovered by previous work. In particular, personas browse videos on peak YouTube browsing times as discussed by Arantes et al. [6, 7]. Previous work also on how children use search engines [15, 19] concluded that children prefer to browse links instead of using search queries. This is the main behavior of our personas, as they click on videos from subscribed channels and do not engage in search activities. However, being synthetic in nature, it is impossible for our personas to capture the complex patterns of human (child or

[17] After each stream, a persona would interrupt its behavior if the one-hour window had expired, and continue issuing requests otherwise. Thus, each such window could exceed one hour depending on the durations of the video contents and ads streamed.
[18] https://developers.google.com/youtube/v3

Table 2: Overview of Content Streamed by Personas.

Brazilian Personas	# Video Content	# Exhibitions Content	# Video Ads	# Exhibitions Ads	# Channels
BR-Child-Mixed-Skips	310	1712	259	766	159
BR-Child-Children-Skips	268	1964	288	924	194
BR-Child-Mixed-NoSkip	259	1643	269	489	170
BR-Child-Children-NoSkip	268	1993	285	750	176
BR-Adult-Mixed-Skips	304	1658	338	1050	214
BR-Adult-Children-Skips	268	1982	384	1284	237
BR-Adult-Mixed-NoSkip	303	1652	211	632	124
BR-Adult-Children-NoSkip	268	1937	258	828	155
Canadian Personas					
CA-Child-Mixed-Skips	319	1343	239	1017	223
CA-Child-Children-Skips	308	1303	277	1410	240
CA-Child-Mixed-NoSkip	318	1261	271	1374	267
CA-Child-Children-NoSkip	308	1229	275	1472	266
CA-Adult-Mixed-Skips	307	1034	250	672	168
CA-Adult-Children-Skips	291	778	240	992	191
CA-Adult-Mixed-NoSkip	292	968	231	638	177
CA-Adult-Children-NoSkip	286	744	236	916	171

adult) behavior. This limitation does not affect our statistical evidence (discussed in Section 4) that our personas are being tracked and targeted. Overall, we can state that our study unveils evidence of targeted advertisements when users stream children's content.

We now describe our dataset. Table 2 shows the numbers of unique video contents, unique video ads and exhibitions (capturing the number of streams) streamed by each of the 16 personas. In total, our personas performed from 744 to 1,993 requests to video contents, streaming from 259 up to 277 unique contents. Since each persona would stream video contents for one hour at a time, this variation in persona behavior stems from the different durations of video contents and ads. In our dataset, video content duration ranges from 31 to 4,893 seconds (over 1h and 20 minutes), with an average of 442 seconds and standard deviation of 640 seconds. Similarly, the durations of video ads range from 10 to 4,463 seconds (over 1h and 15 minutes), with average and standard deviation equal to 44 and 81 seconds, respectively.

Table 2 also shows the number of different channels associated with the video ads each persona was exposed to. We note that on YouTube any video content may be advertised, thus becoming a video ad. That is, video ads are simply videos where owners (channels) created a marketing campaign to publicize them. In this sense, each ad channel can be mapped to a potential brand or content producer. Our personas were targeted by at least 124 such promoters. Finally, over 200 different video ads were streamed, with some personas streaming video ads over 1,000 times. Such numbers show how common advertisements are on YouTube. On average 68% of video content exhibitions had paired video ads.

3.2 Coding

To answer our research questions we performed an open coding on our dataset [20] with the participation of volunteers. The objective of this coding was to manually identify the types of content being advertised to the personas. Our final coding questionnaire is shown in Table 3. This questionnaire was created to capture various aspects of video ads such as: the product and brand being advertised (Q2-3), the content of the ad (Q4), the sponsor of the ad (Q5), and factors related to how the ad may affect children (Q6-9). Before addressing these questions, the volunteers were asked to indicate whether the video ad was still online (Q1). If this was not the case, coders would follow on to the next video ad.

Question Q1 is factual (yes/no answer) as it does not depend on the opinions of the coders. For addressing Q2 and Q3 we employed an open coding [20] approach: coders could write the answer they felt was more appropriate in free text form. The answers to these questions were later compared and standardized. Q4-Q9 are closed questions, i.e., coders had to select one out of pre-defined possible answers (shown in the table). For Q4 in particular, we employed a refined set of categories using the YouTube 8m dataset [1].

To perform our analysis, we employed a three step coding approach. We began our work with a small sample (50 video ads). Three coders (authors) worked together on this sample, discussing and determining the possible closed categories for questions Q2-Q9. The final questionnaire shown in Table 3 was developed based on the results of this initial step, which were discarded afterwards.

On a second step, we used another small sample of video ads to validate coder agreement: we selected 126 random video ads and used them to measure the inter-rate agreement between coders. Video ads were selected in a uniform manner with each persona having roughly 8 video ads selected. This initial coding was performed as follows. Two coders watched video-ads and answered the questions in Table 3. After both coders finished their work, we measured the Fleiss' Kappa [16] (κ) agreement score. This score varies from -1 to 1, with negative values indicating disagreement and positive ones indicating agreements. As a guideline, $\kappa > 0.4$ captures a reasonable agreement, whereas $\kappa > 0.75$ serves as evidence of very strong agreement [16]. Table 3 shows the agreement scores for each question. As we can see, all scores are above 0.63 and statistically significant. Next, coders discussed their answers as to standardize their results (e.g., lower case v. upper case names) for the open questions. Also, coders found that in some cases the

Table 3: Coding Questionnaire. All of the κ values are statistically significant with $p < 0.01$

ID	Question	Type	κ
Q1	Has this ad been deleted?	Factual	-
Q2	Name the brand of the ad	Open	0.90
Q3	Name the product of the ad	Open	0.63
Q4	Categorize video ad (see Table 4)	Closed	0.81
Q5	Who sponsored the ad? (commercial brand, public agency, social influencer)	Closed	0.78
Q6	Does this ad employ multimedia content in it's discourse to attract the attention of children? (Yes/No)	Closed	0.68
Q7	Do any well known children character appear on the ad? (Yes/No)	Closed	0.65
Q8	Does this ad advertise a product geared to children? (Yes/No)	Closed	0.76
Q9	Based on your personal beliefs, do you think the ad is unsuitable for children? (Yes/No)	Closed	0.79

Table 4: List of Video Categories

Arts & Entertainment	Auto & Vehicles
Beauty & Fitness	Business & Industrial
Clothing	Computers & Electronics
Finance	Food
Games	Health & Hygiene
Hobbies & Leisure	Home & Garden
Internet & Telecom	Jobs
Law & Government	People & Society
Pets & Animals	Real Estate
Science	Shopping
Sports	Transportation
Travel	Another Channel

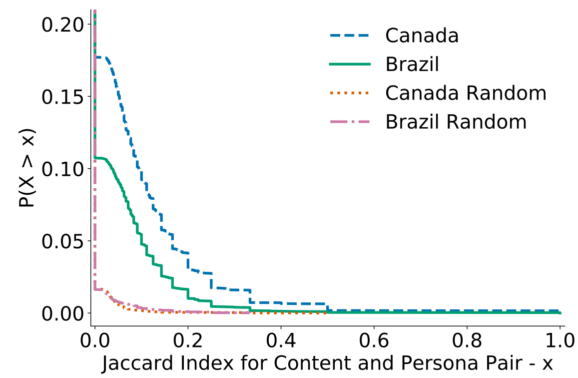

Figure 2: CCDF for Jaccard Index

YouTube 8m categories were either too broad or narrow After some discussion, they created a refined set of categories for Q4 (Table 4).

Finally, the two independent coders were asked to code at least 600 (300 each) video ads. We scheduled video ads to coders in a uniform manner as to preserve a roughly equal number of ads shown to each persona. We also ordered the video ads by popularity (in # of exhibitions). Each coder then worked independently until coding at least 300 video ads. In the end, we reached 663 coded video ads, out of which 23 had been deleted by YouTube (Q1). These video ads account for 78% (11,992 out of 15,214) of all exhibitions in the original data. Considering only coded video ads, each persona had on average 765 exhibitions of 168 unique ads. Also, the vast majority (11,226) of the exhibitions were of ads from commercial brands. There were very few cases of social influencers (other channels) and public agencies exhibiting video ads. This suggests that brands are the ones most often exploring YouTube's advertising ecosystem at the time of our data collection. The coded video ads came from 397 different brands (Q2), selling 311 products (Q3) in total.

Before doing presenting our results we justify our coding methodology. We here followed a standard coding approach that was refined in three steps. Note that most of the questions are either closed or objective in nature Q2-Q3 (open and objective) Q4-Q8 (closed and objective). While Q9 is subjective, we only use the answers of this question for anecdotal evidence. That is, we describe the video ads deemed unsuitable by coders but we do not state in any sense that this is are general finding nor that they reflect societal opinions. In particular, we further explore these videos using factual properties (e.g., alcoholic products or game ratings). Finally, we employed adequate statistical tests [16] in consonance with our sample sizes to show that coders agreed on most questions.

4 EVIDENCE OF BEHAVIORAL TARGETING

We now tackle our first research question. *Are children's content consumers being targeted by video-ads on YouTube? If so, does user behavior (based on different personas) affect the ad pairings?*

Initially, we analyzed the answers given to questions Q6, Q7 and Q8 with respect to the exhibitions of coded video ads (11,226). We found that 1,242 of the exhibitions employed multimedia content to attract children's attention (Q6), 242 had some well-known children character (Q7), and 688 advertised products targeting children (Q8). Ads selling children's products included: toys, clothes and meals. The majority of the other video ads sold a wide variety of products such as Internet connections, cars, or travel packages. Even the numbers focused on children's product or multimedia elements targeting children may seem small, they serve as evidence that marketers are exploring YouTube to target consumers of children's content (our personas).

Next, we analyzed whether personas are subject to some kind of targeting. To do so, we initially explored whether different personas are subject to different video ads when streaming the same video content. Thus, we measured the Jaccard Index of the video ad ids which were paired for the same video content for every pair of personas. To define the index, let $\mathcal{A}_i$ be the set of ads for some persona i and $\mathcal{A}_j$ be a similar set for persona j with $i \neq j$. The Jaccard index $J(i,j) = \frac{|\mathcal{A}_i \cap \mathcal{A}_j|}{|\mathcal{A}_i \cup \mathcal{A}_j|}$ captures the fraction of ads similar for both personas. This index was measured controlling for video ads paired with the same video content. Naturally, our analysis is limited to video contents streamed by at least two personas. Our analysis is performed in a stratified manner for each country.

In Figure 2 we show the complementary cumulative probability distribution (CCDF) of the Jaccard index. From the figure, we can see that both for Brazil and Canada, the fraction (y-axis) of Jaccard indexes above 0 (x-axis) is less than 20%. In this sense, for the majority, over 80%, of the cases there is no similarity in video ad ids streamed to different personas when streaming the same content. In fact, less that 2.5% of Jaccard indexes are above 0.4. Finally, the Jaccard index for Canada is usually superior when compared to Brazil, indicating more overlap in ads for Canadian personas.

We can view two explanations for such results. The first is that YouTube is targeting our personas. The second is that the Jaccard is small due to the large amount of video ads being streamed. Thus, random (absent) targeting also leads to such small numbers. To discard this second hypothesis, we also measured the Jaccard index for the same video content considering random video ad pairings. These results are also shown on the figure for both countries. Notice that for random pairings, the Jaccard index is significantly lower than for our real data. This difference can raises the argument that the second hypothesis is false. Thus, we have some evidence of targeting being taken place for personas streaming children's content. This result is particularly interesting given that we conditioned, controlled, on video content. Other variables from our personas are the ones being used for video ad pairings.

To further explore such evidence, we employed a χ^2-test [37] on the responses of Q4. The test is used to assess whether the contents of the video ads, captured by their categories (response to Q4), differ across personas. The null hypothesis of the test states that the distributions of video ad categories of two personas are the same. Rejecting this hypothesis provides evidence that some types of video-ads are more often shown to certain personas, that is, behavioral targeting occurs.

We started by comparing the contents of video ads shown to Adult (age over 40) and Child (age 13) personas. Due to language differences, all of our tests are applied for each country separately, as it is naturally expected for video ads in Brazil and Canada to differ. The number of streams for each category of video ads of all Adult personas is thus compared to the same distribution for Child personas. For fair analysis, we only executed the test considering categories with at least 5 streams for both types of personas.

In this first test, the χ^2-statistic was 420 ($p < 0.001$) for Brazil and 79 ($p < 0.001$) for Canada, which implies that we should reject the null hypothesis This result is interesting given that, in this particular setting, only the age differs across personas. That is, even with the Adult group streaming similar videos to the Child group, the (categories of) video ads shown to both are statistically different, thus serving as evidence of targeting towards both adults and children. The latter raises a flag of concern, whereas the former, though expected, may also be an issue since some of our adult personas aimed to capture the setting of a parent streaming children's content. We repeated the test by considering only personas that focused solely on children content, in order to remove any possible effect caused by the accesses to random music videos (mixed content). However, once again we can reject the null hypothesis (χ^2 = 109 for Brazil χ^2 = 66 for Canada, both with $p < 0.001$).

To look further into the issue, we decided to analyze every persona against all others. That is, we aimed to uncover evidence of targeting by looking into all of our variables. To do so, we tested whether the category distribution of each persona is the same as the overall category distribution computed over all video ads shown to the other personas (every persona excluding the one of interest). As we are performing multiple tests (16), we correct p-values using the Benjamini-Hochberg approach [37]. Once again, we found that all personas are statistically different from the others ($p < 0.001$). Given that personas share the same country as well as subscribed channels, and stream videos in a similar pattern, we view this result as evidence that the considered variables, namely age, skipping behavior and content watched, are being taken into account by YouTube so as to target video ads towards the personas.

Table 5: Accuracy of Adult/Child Classification – Brazil

	Precision	Recall	F1-Score	Support
Multinomial Naive Bayes				
Adult	**0.75**	**0.70**	**0.72**	3,768 (0.57)
Child	**0.64**	**0.71**	**0.67**	2,907 (0.43)
SVM				
Adult	**0.76**	**0.68**	**0.72**	3,768 (0.57)
Child	**0.64**	**0.73**	**0.68**	2,907 (0.43)

Our final analyses aimed to assess the extent to which children, or at least children personas, are targeted by YouTube. To do so, we explored the extent to which one can automatically predict whether a given video ad will be shown to a given persona. That is, we developed a classifier for child and adult personas. This classifier is based solely on the video ad's content children/adults.

We used the video ad metadata gathered from the YouTube's API to train binary classifiers. We experimented with both Multinomial Naive Bayes (NB) and Support Vector Machines (SVM) [37] classifiers. The explanatory variables, or features, for each classifier is composed of the textual content shown in the Title and Description of the video ad. Given that such features are textual, we encoded them using a *one hot* approach (a binary vector where 1 indicates the presence of a term). We also added the moment during the day a video ad was to be streamed as a categorical feature (morning, afternoon and evening).

We focus on features extracted from the metadata of the video ads as they are the common factor across personas. That is, we avoid adding features that are related to our interpretation of the video ads (features extracted from the responses of the questionnaire). Similarly, our control variables also correlate with the content of video ads (as discussed previously), thus such variables could also lead to over-fitting. Finally, the response variable, or class, is set to one when the persona simulates a child and zero otherwise. Since none of our coded variables are used, we ran our experiments on the complete dataset (Table 2). After filtering out video ads without any API data (deleted or private), we were left with 6,675 exhibitions to Brazilian personas and 8,448 to Canadian ones.

Our results are based on a 5-fold cross-validation experimental design. In this approach, the dataset is split in 5 folds of equal sizes. For each classifier, 3 folds are used for training, one for validation (parameter tuning) and one for testing (held out). We point out that the Multinomial Naive Bayes classifier does not require any parameter tuning. For SVM we used a Linear kernel for the SVM. We tuned the SVM cost (C) through a grid search

Table 6: Accuracy of Adult/Child Classification - Canada

	Multinomial Naive Bayes			
	Precision	Recall	F1-Score	Support
Adult	**0.45**	**0.55**	**0.50**	3,196 (0.38)
Child	**0.68**	**0.59**	**0.63**	5,252 (0.62)
	SVM			
Adult	**0.43**	0.44	**0.43**	3,196 (0.38)
Child	**0.65**	**0.65**	**0.65**	5,252 (0.62)

$C = 0.001, 0.01, 0.1, 1, 10, 100, 1000$ and selected the one with best result in the validation set. To compute statistical significance we repeated each 5-fold experiment 20 times. Given that each 5-fold leads to an average scores (over the 5 held out sets), our results are presented as an average of 20 repetitions of such 5-fold experiments.

We note that our dataset have duplicate entries, as some video ads are exhibited more than once. Indeed, 23% of the video ads in our dataset have repeated exhibitions, and one video ad has 262 exhibitions (maximum). Such duplicates may negatively impact the performance of the classifiers, especially when the same video ad (same features) is exhibited to both classes. This lowers the classifier capability of distinguishing between the classes. We decided to keep such instances, and analyze the performance of the classifiers in light of this factor. We note that we are still able to uncover evidence of targeting even when keeping them.

Tables 5 and 6 show classification results for Brazil and Canada, respectively. The table presents average precision, recall and F1-score per class over the 20 executions. The precision for a class is the fraction of predictions for that class which were correct. In contrast, recall is the fraction of correct predictions out of all instances of the class. The F1-score is trade-off score capturing the harmonic mean of precision and recall. Each metric varies from 0 to 1 (max).

To uncover statistical evidence of targeting, we resorted to a uniformly random classifier, which simply picks Adult and Child labels per stream with 50% chance each. The idea is that if targeting is occurring, our classifiers, which exploit the video ad metadata, should perform significantly better than random chance. Also, with half of our personas belonging to each class, the random classifier represents a system where no targeting occurs. A random classifier has an expected recall of 0.5 and a precision proportional to the number of instances in the class (column support shown in the two tables). We compared precision, recall and F1 results of each classifier with those of a random classifier using a t-test [37]. We found that most results of our classifiers are indeed statistically higher than random chance ($p < 0.001$ shown in bold on the tables). This suggests that signals of the personas are being used for targeting.

The only result that was below random chance was recall of Adult personas by SVM for Canada. We emphasize that our features are somewhat limited, exploring only the textual content of video advertisements. While this may explain the lower results for the Canadian dataset, the results for Brazil are quite high and significantly above random chance in every metric. Indeed, for Brazil, the lowest improvement over random chance is 33% (0.75 vs 0.57 of Precision on Child personas). Such results may imply that the content of the video ads had a stronger importance in targeting

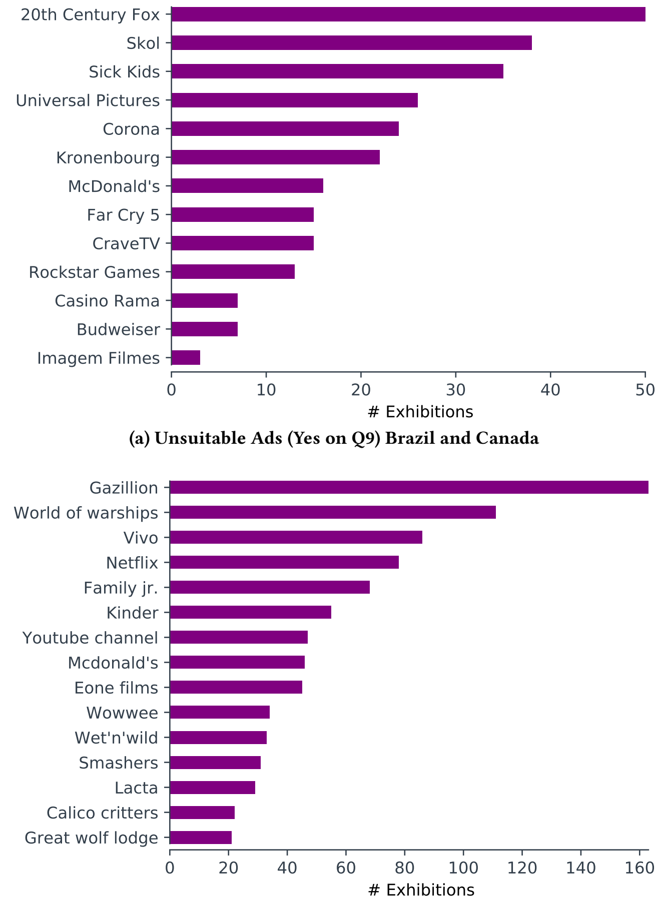

(a) Unsuitable Ads (Yes on Q9) Brazil and Canada

(b) Child Oriented (Yes on Q6, Q7 or Q8) Brazil and Canada

Figure 3: Brands with answer Yes on either Q6, Q7, Q8, Q9.

specific personas in Brazil, whereas other features (which we do not capture) may play a larger role in Canada.

In sum, regarding our first research question, we can conclude that children content consumers are indeed being target. The results of our coding indicate that advertisers are not only streaming video ads geared towards children, but also exploring multimedia content that targets children to sell products. Our Jaccard index, χ^2, and classification results unveil evidence that video ads being exhibited are not random. In our analysis, we controlled for age, country and video content, all reaching similar results. The last experiment, our classification analysis, also unveils some evidence of children accounts being specifically targeted (scores above random chance).

5 ADVERTISEMENTS AND REGULATIONS

We now present our results on RQ2: *Which content, brands and products target children content consumers?*

We begin our discussion by showing in Figure 3 the brands that had at least one exhibition that received a *Yes* response for Q6, Q7, Q8, and Q9. To investigate video ads with content unsuitable to children (in the opinion of coders), we show the result of Q9 in Figure 3(a), whereas to investigate video ads geared towards children we combine the results for Q6, Q7 and Q8 in Figure 3(b). Each figure shows the number of exhibitions for each brand. Recall that, in our dataset, each brand usually sells a single product. Thus, the figures show the brand name only (we shall discuss products in the text). Due to space constraints, the plots show the most popular (in # of exhibitions in our dataset) brands only.

Our discussion on RQ2 is mostly descriptive. The figures show the number of exhibitions combining all personas. Thus, we shall discuss brands that are either unsuitable (Q9) or stream ads geared towards children to our personas (i.e., Yes on Q6 to Q8).

Before discussing our findings, recall that it is impossible to set up accounts with less than 13 years of age. In other words, no account is legally considered as child in neither Brazil nor Canada. However, as discussed in our introduction, children are known to be major consumers of content on YouTube regardless of age limits imposed by the system [12, 25]. Here, we look into potentially controversial (unsuitable or possibly conflicting regulations) content that may be exhibited when streaming children's content.

5.1 Unsuitable Advertisements

We now describe the ads deemed as unsuitable to children by coders. We point out that the coders' opinions may not reflect overall values of society nor reflect different interpretations of regulations. Instead, we aim here at discussing controversial and unexpected observations, such as beer and mature game/movie brands targeting our personas, based on factual properties (e.g., game ratings).

From Figure 3(a), we initially point out to alcoholic beverages (beer) brands whose ads appeared in our datasets. A list of these brands are: Skol, Corona, Kronenbourg, and Budweiser. We also found that movie trailers (20th Century Fox, Universal Pictures, Imagem Filmes and CraveTV) seem to be targeting our personas. When checking the Internet Movie Database (IMDB), we found that most of these movies (with one exception) are rated as mature (17+). A similar finding occurs with video games (Red Dead Redemption 2 and Far Cry 5), where games are rated as mature according to the Internet Games Database (IGDB).

We also found evidence of fast food adverts targeting our personas. In particular, McDonald's Brazil was marketing *happy meals*[19]. While it is not our focus to discuss the nutrition values of happy meals, the previously discussed guidelines for both Brazil and Canada make explicit statements on fast food and products.

5.2 Child Oriented Advertisements

When watching the video advertisements, coders found several multimedia elements commonly exploited that may attract the attention of children (Q6). The two coders pointed out to elements such as: (1) animations; (2) children within the advertisement interacting with the product; (3) children with the advertisement interacting with others; (4) ads selling toys and children's games; (5) ads selling foods bundled with toys (e.g., surprise gifts in cereals and chocolates); (6) famous characters; and, (7) children's songs.

We now discuss some of the brands and products that explicitly geared their advertisements towards children (Figure 3(b)). We begin by discussing some of these brands/products that may contradict healthy eating guidelines. We have already mentioned McDonald's as possibly marketing an unsuitable product. Here we also found evidence of video ads selling chocolates/candy (Kinder and Lacta). Though the latter are not directly advertising replacements to healthy meals, candy and sweet desserts are usually viewed as junk-food[20]. Brands explicitly targeting other less controversial food products to children were cereal brands.

On both Brazil and Canada, our personas were target of toy adverts from companies such as Gazzilion, WowWee, Kinder and Shashers shown in the figure, as well as other less frequent brands like Monica Toy and Lego. As we have discussed, Brazil explicitly bans selling toys on open television. We argue, that despite not being categorized as open television, YouTube does share some similarities such as being immensely popular and free to access. Lego, also employed well known children's characters from Marvel Comics. Another popular brand to make use of this practice was Riachuelo (clothing store in Brazil), advertising products with characters from DC Comics and Harry Potter. We also found other YouTube channels (YouTube channel on the figure) that employ variations of well known children's characters to attract viewership. Other authors have found that some uploaders explicitly exploit well known characters for malicious purposes [27].

The extent to which the above examples may explicitly break laws (in the judicial sense) is highly debatable. Nevertheless, all of these examples do appear to conflict, in one way or another, with the aforementioned countries' regulations. Our goal with this discussion was to uncover evidence of such practices on YouTube. Together with our results on RQ1, our findings raise a flag on how the system is being used to target products to children's videos.

6 CONCLUDING DISCUSSION

We now discuss some implications of our work. Our findings unveil important concerns and troubling findings that arise when children are exposed to YouTube as entertainment. One of the most important findings of our research is the evidence found of algorithmic targeting accounts that mostly stream children's content. We also discuss how the content and product of such ads may conflict with regulations of children's advertising.

As we have noted, YouTube does not fall into the category of open television that is usually regulated. However, when we consider the immense popularity of the system together with our findings, it may be the case that regulations need to the developed for children's advertising on online streaming services.

It is important to point out that it is impossible to state (with our experiments) whether marketers are explicitly using YouTube to target children's videos. One possibility, which may also explain our findings, are rogue machine learning model which are targeting videos geared towards children spuriously. That is, the system does not allow marketers to target children explicitly. However, trained models still enable such practice since they may not be able to distinguish children. Whatever the reason, YouTube has recently stated that it aims to limit advertisements in children's content[21].

[19]http://YouTube.com/watch?v=ct0cW1HjAtI

[20]https://en.wikipedia.org/wiki/Junk_food

[21]https://www.bloomberg.com/news/articles/2019-08-20/youtube-plans-to-end-targeted-ads-to-kids-to-comply-with-ftc

Even if YouTube limits advertising towards children's content in the future, we argue for the need to audit online ad matching algorithms. That is, if algorithms are spuriously targeting children nowadays, this practice may continue towards other sensitive demographics even after limiting to children. We argue that content providers such as YouTube may offer more transparency on why video ads are targeted to such accounts. Transparency measures can help both end-users and marketers to understand why some ads pair with specific content.

As we have discussed in the text, human behavior is naturally complex, and we cannot state that our personas capture all of the intricacies of adult or child behavior online. However, we coded our personas to follow patterns of video browsing [6–8], as well as conventional patterns of children when browsing the web [15, 19]. That is, our personas consume content via browsing (subscriptions) and not via search engines. It is known that content providers like Google to track multiple apps. By solely focusing on YouTube, we show that YouTube is effectively tracking our personas via the profile page on Google accounts. Thus, exploring more complex personas through multiple apps and websites is left as future work.

Another limitation is that our work focuses solely on browsing via YouTube's desktop applications. Children often use both desktop and mobile apps to stream videos [12, 25, 29], and some only have access to specialized apps such as Youtube Kids. As mentioned, those specialized apps still rely on advertising. Moreover one might argue that their use is not yet widespread[22]

ACKNOWLEDGEMENTS

We thank Yuri dos Santos for comments. Funding was provided by the authors' individual grants from CNPq (421884/2018-5), CAPES, and Fapemig. Computational resources were provided by the Microsoft Azure for Data Science Research Award (CRM:0740801).

REFERENCES

[1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. 2016. Youtube-8m: A large-scale video classification benchmark. *arXiv:1609.08675* (2016).

[2] Tamara Adlin and John Pruitt. 2010. *The essential persona lifecycle: Your guide to building and using personas.* Morgan Kaufmann.

[3] Maria Carpen Amarie, Ioannis Pefkianakis, and Henrik Lundgren. 2014. Towards Understanding Mobile Video Ads. In *Ubicomp*.

[4] Ashton Anderson, Ravi Kumar, Andrew Tomkins, and Sergei Vassilvitskii. 2014. The dynamics of repeat consumption. In *WWW*.

[5] Athanasios Andreou, Giridhari Venkatadri, Oana Goga, Krishna Gummadi, Patrick Loiseau, and Alan Mislove. 2018. Investigating ad transparency mechanisms in social media: A case study of Facebook's explanations. In *NDSS*.

[6] Mariana Arantes, Flavio Figueiredo, and Jussara M Almeida. 2016. Understanding video-ad consumption on YouTube: a measurement study on user behavior, popularity, and content properties. In *Websci*.

[7] Mariana Arantes, Flavio Figueiredo, and Jussara M Almeida. 2017. Towards Understanding the Consumption of Video-Ads on YouTube. *The Journal of Web Science* 3, 1 (2017).

[8] Mariana Arantes, Flavio Figueiredo, Jussara M Almeida, and Raquel O Prates. 2018. An Investigation of User Actions and Experiences when Exposed to YouTube Video Ads. In *Webmedia*.

[9] Camila Souza Araújo, Gabriel Magno, Wagner Meira, Virgilio Almeida, Pedro Hartung, and Danilo Doneda. 2017. Characterizing videos, audience and advertising in Youtube channels for kids. In *SOCINFO*.

[10] Yasare A. Arnas, Isil Tas, and Ogul Gurgah. 2016. The development of brand awareness in young children: how do young children recognize brands? *International Journal of Consumer Studies* 40, 5 (2016).

[11] Juan Miguel Carrascosa, Jakub Mikians, Ruben Cuevas, Vijay Erramilli, and Nikolaos Laoutaris. 2015. I always feel like somebody's watching me: measuring online behavioural advertising. In *CONEXT*.

[12] CGI-Brasil. 2017. Pesquisa sobre o uso da Internet por crianças e adolescentes no Brasil.

[13] Jilin Chen, Eben M Haber, Ruogu Kang, Gary Hsieh, and Jalal Mahmud. 2015. Making Use of Derived Personality: The Case of Social Media Ad Targeting.. In *ICWSM*.

[14] Amit Datta, Anupam Datta, Jael Makagon, Deirdre K Mulligan, and Michael Carl Tschantz. 2018. Discrimination in Online Advertising: A Multidisciplinary Inquiry. In *FAT*.

[15] Sergio Duarte Torres, Ingmar Weber, and Djoerd Hiemstra. 2014. Analysis of search and browsing behavior of young users on the web. *ACM Transactions on the Web (TWEB)* 8, 2 (2014), 7.

[16] Joseph L Fleiss, Bruce Levin, and Myunghee Cho Paik. 2013. *Statistical methods for rates and proportions.* Wiley & Sons.

[17] Umar Iqbal, Zubair Shafiq, and Zhiyun Qian. 2017. The ad wars: retrospective measurement and analysis of anti-adblock filter lists. In *IMC*.

[18] Soon-Gyo Jung, Jisun An, Haewoon Kwak, Moeed Ahmad, Lene Nielsen, and Bernard J Jansen. 2017. Persona generation from aggregated social media data. In *CHI*.

[19] Andrew Large, Valerie Nesset, and Jamshid Beheshti. 2008. Children as information seekers: what researchers tell us. *New Review of Children's Literature and Librarianship* 14, 2 (2008), 121–140.

[20] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2010. *Research methods in human-computer interaction.* Wiley & Sons.

[21] Huyen Le, Raven Maragh, Brian Ekdale, Andrew High, Timothy Havens, and Zubair Shafiq. 2019. Measuring Political Personalization of Google News Search. In *WWW*.

[22] Matthew Malloy, Mark McNamara, Aaron Cahn, and Paul Barford. 2016. Ad blockers: Global prevalence and impact. In *IMC*.

[23] Yan Mei, Weiliang Zhao, and Jian Yang. 2017. Maximizing the Effectiveness of Advertising Campaigns on Twitter. In *BigData*.

[24] Caroline Oates, Mark Blades, and Barrie Gunter. 2002. Children and television advertising: When do they understand persuasive intent? *Journal of Consumer Behaviour: An International Research Review* 1, 3 (2002), 238–245.

[25] UK Ofcom. 2017. Children and parents: Media use and attitudes report.

[26] Suzanna J Opree, Moniek Buijzen, Eva A van Reijmersdal, and Patti M Valkenburg. 2014. Childrenś advertising exposure, advertised product desire, and materialism: A longitudinal study. *Communication Research* 41, 5 (2014).

[27] Kostantinos Papadamou, Antonis Papasavva, Savvas Zannettou, Jeremy Blackburn, Nicolas Kourtellis, Ilias Leontiadis, Gianluca Stringhini, and Michael Sirivianos. 2019. Disturbed YouTube for Kids: Characterizing and Detecting Disturbing Content on YouTube. *arXiv preprint arXiv:1901.07046* (2019).

[28] Enric Pujol, Oliver Hohlfeld, and Anja Feldmann. 2015. Annoyed users: Ads and ad-block usage in the wild. In *IMC*.

[29] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, and Serge Egelman. 2018. "Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale. *Proceedings on Privacy Enhancing Technologies* 2018, 3 (2018), 63–83.

[30] Vicky Rideout. 2016. Measuring time spent with media: The Common Sense census of media use by US 8-to 18-year-olds. *Journal of Children and Media* 10, 1 (2016).

[31] Juliet B Schor. 2014. *Born to buy: The commercialized child and the new consumer cult.* Simon and Schuster.

[32] Sergej Sizov and Sarah Piller. 2016. Economic value of web keyword advertising campaigns. In *Websci*.

[33] Till Speicher, Muhammad Ali, Giridhari Venkatadri, Filipe Nunes Ribeiro, George Arvanitakis, Fabrício Benevenuto, Krishna P Gummadi, Patrick Loiseau, and Alan Mislove. 2018. Potential for Discrimination in Online Targeted Advertising. In *FAT*.

[34] Patti M Valkenburg and Moniek Buijzen. 2005. Identifying determinants of young children's brand awareness: Television, parents, and peers. *Journal of Applied Developmental Psychology* 26, 4 (2005), 456–468.

[35] Dieneke Van de Sompel and Iris Vermeir. 2016. The influence of source attractiveness on self-perception and advertising effectiveness for 6-to 7-year-old children. *International Journal of Consumer Studies* 40, 5 (2016).

[36] Pengyuan Wang, Yechao Liu, Marsha Meytlis, Han-Yun Tsao, Jian Yang, and Pei Huang. 2014. An efficient framework for online advertising effectiveness measurement and comparison. In *WSDM*.

[37] Larry Wasserman. 2013. *All of statistics: a concise course in statistical inference.* Springer Science & Business Media.

[38] Stephen B Wicker and Kolbeinn Karlsson. 2017. Internet advertising: technology, ethics, and a serious difference of opinion. *Commun. ACM* 60, 10 (2017), 70–79.

[39] Wikipedia, The Free Encyclopedia. 2018. Advertising to Children. https://en.wikipedia.org/wiki/Advertising_to_children

[40] Jinyoung Yeo, Sungchul Kim, Eunyee Koh, Seung-won Hwang, and Nedim Lipka. 2017. Predicting Online Purchase Conversion for Retargeting. In *WSDM*.

[22] At the time of writing, the YouTube Android app had 25M downloads on the Play Store (http://play.google.com), whereas the YouTube Kids app had close to 580k.