

Is This An Ad?: Automatically Disclosing Online Endorsements On YouTube With AdIntuition

Michael Swart
Princeton University
michaelswart@gmail.com

Ylana Lopez
Princeton University
ylana.lopez96@gmail.com

Arunesh Mathur
Princeton University
amathur@cs.princeton.edu

Marshini Chetty
University of Chicago
marshini@uchicago.edu

ABSTRACT

Undisclosed online endorsements on social media can be misleading to users who may not know when viewed content contains advertisements. Despite federal regulations requiring content creators to disclose online endorsements, studies suggest that less than 10% do so in practice. To overcome this issue, we need knowledge of how to best detect online endorsements, knowledge about how prevalent online endorsements are in the wild, and ways to design systems to automatically disclose advertising content to viewers. To that end, we designed, implemented, and evaluated a tool called AdIntuition which automatically discloses when YouTube videos contain affiliate marketing, a type of social media endorsement. We evaluated AdIntuition with 783 users using a survey, field deployment, and diary study. We discuss our findings and recommendations for future measurements of and tools to detect and alert users about affiliate marketing content.

Author Keywords

social media, browser extension, advertisements, influencer

CCS Concepts

•Human-centered computing → Human computer interaction (HCI);

INTRODUCTION

Online endorsements are a form of advertising that help social media influencers to monetize their content [16]. These influencers are paid because of the perception that they are able to shape the opinion of their followers on a daily basis. For instance, Kylie Jenner, a social media influencer, reportedly made \$26.5 Million from just 53 Instagram advertisements [6]. The high fees that brands are willing to pay for online endorsements show how lucrative they can be for content creators [39]. When content creators do form a connection to

a brand, they are required by the Federal Trade Commission (FTC) to disclose this relationship on the platforms they use for endorsements [11]. However, a recent study suggests less than 10% actually do so on YouTube and Pinterest [21] and this is problematic because viewers may be misled by these disguised advertisements.

The consequence of undisclosed endorsements may be benign: a viewer may not realize that an influencer's endorsement of a product is inauthentic. More extreme consequences include financial loss such as when the content creator is being deceptive [23] or the product being marketed has deceptive practices [8]. Arguably, influencers' 'disguised advertisements' are a form of 'dark pattern' [4], a choice to leave out information in order to lead a viewer down a decision-making path for the benefit of the influencer. Undisclosed advertising content could even be classified as misinformation, causing viewers to falsely believe the content is unbiased, which can be particularly egregious if in aid of a political agenda [30].

Although some platforms such as YouTube do allow content creators to self-indicate a sponsorship, for instance, checking a box to indicate 'Paid Promotion' [13], it is unclear how often content creators use these features. Yet, studies have shown having these types of disclosures can help users identify advertisements and form more critical attitudes towards the brand being promoted [9, 34, 3, 37]. The question then is how can we automatically detect and communicate to a user when content contains online endorsements so that users are informed about the content they are consuming? We set out to address this issue by focusing on one type of online endorsement, affiliate marketing, on the YouTube platform. In affiliate marketing, influencers with a brand relationship are paid for sales or referrals generated from their content consumers. For example, a YouTuber may earn a commission if a video viewer clicks on a brand-generated Uniform Resource Locator (URL) provided to the YouTuber that links to a brand's product or website. These links serve as a source of ground truth for automatically detecting affiliate marketing content which makes automatic disclosure feasible. In our work, we posed the following research questions:

- How can we best detect and measure affiliate marketing content automatically and in real-time?

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '20, April 25–30, 2020, Honolulu, HI, USA.

© 2020 Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-6708-0/20/04 ...\$15.00.

<http://dx.doi.org/10.1145/3313831.3376178>

- How can we design and implement automatic ad disclosures for affiliate marketing content?
- How do users react to real-time automatic ad disclosures ?

To answer these questions, we used the only publicly available existing data set of known affiliate marketing link patterns [19] and a set of 0.5 million YouTube videos compiled by Mathur et al. [21] to detect affiliate marketing content from YouTube video descriptions. We used the data set of videos to identify two additional features for detecting affiliate marketing content automatically, 1) url parameters called Uniform Tracking Modules (UTM) [14] in links in video descriptions and 2) customized coupon codes [24] used in video descriptions. We built a coupon code classifier and detection module for these features in YouTube videos in real-time. Next, we integrated these modules into a browser extension called *AdIntuition* that automatically detects and discloses affiliate marketing content to users using a banner displayed above a video and by highlighting links and coupon codes in video descriptions.

We evaluated the efficacy of our interface using three studies. In Study 1, we performed a survey of 300 Amazon Mechanical Turkers to evaluate as part of the *AdIntuition* design process to test if the banner was noticeable to users and if they understood the disclosure. In Study 2, we conducted a diary study with 11 users over two weeks to examine how users react to automatic advertising disclosures and in Study 3, we conducted a field deployment with 472 users to evaluate how well our coupon code classifier and feature detection modules performed with videos real users watch and to learn more about how real users encounter affiliate marketing content. Based on our studies, we have the following main findings:

- Users were able to more easily identify advertising content with *AdIntuition* disclosures than without them.
- Users became more reflective on content creators and content when automatic disclosures were presented to them in real-time
- Affiliate marketing content on YouTube can be detected in real-time using the known affiliate link patterns from Mathur et al. [21] but also by detecting the additional features we identified, UTM parameters and coupon codes, in video descriptions.
- Users encountered affiliate marketing content very differently depending on whether the content creators they view use affiliate marketing in the first place and how frequently they view videos from the same creators

Our contributions include 1) a coupon code classifier and feature detection module to detect affiliate marketing videos in real time; 2) evidence users notice and understand *AdIntuition*'s automatic disclosures; 3) evidence of automatic disclosures' effect on users, and 4) evidence that *AdIntuition* works well with real videos in the wild. We have made *AdIntuition* publicly available as a Chrome and Firefox extension with source code on and a list of commonly occurring UTM parameters associated with affiliate marketing on GitHub¹. Based

¹<https://github.com/mike-swart/AdIntuition/>

on our findings, we make recommendations for future work including how our tool can be extended to detect, measure, and automatically disclose advertising content on other platforms in real-time including blogs, websites, and other social media platforms. *AdIntuition* also serves as example of how to combat misleading online content and can inform the design of similar tools to inform consumers about dark patterns, misinformation, and disinformation.

BACKGROUND AND RELATED WORK

We describe disguised advertisements, affiliate marketing, disclosures, and related work.

Disguised Advertisements

Many studies examine deceptive and manipulative practices online that steer users into making decisions for the benefit of the service provider. These practices, known as 'dark patterns' encompass disguised advertisements, which at the interface level are meant to trick the user into clicking on them to force engagement [4, 15, 20]. Arguably, advertising on social media where there is no disclosure is also a form of a disguised advertisement. In our work, we provide a tool for informing users of disguised advertisements on YouTube. Tangentially, many researchers are examining online misinformation and disinformation, where users are deliberating spreading fake or false information, sometimes in aid of a political agenda [30]. Our work is also related to these studies in that *AdIntuition* informs users of misleading advertising content so that users can assess the content and any bias accordingly. We also have to take care that *AdIntuition* itself is not misleading.

Types of Advertising on Social Media

Social media advertising takes two forms. Platform-based advertising refers to ads on social media platforms that can be purchased by merchants or sellers. Examples of such advertising include Facebook ads [10] and Twitter ads [32]. Endorsement-based advertising or online endorsement by contrast refers to ads on social media platforms in which merchants or sellers engage with specific users—influencers or endorsers—to advertise their products to other users. Online endorsements occur in three forms [38]: sponsored advertisements, product sampling, and affiliate marketing. In **sponsored advertising**, a merchant or seller pays a fee to an influencer or endorser in return for endorsing their product. In **product sampling**, merchants or sellers send free products to influencers to get them to endorse and promote the product to their followers. **Affiliate marketing**—the focus of this paper—an influencer or endorser earns fees based on the sales that they generate for the product [7].

Affiliate Marketing

In affiliate marketing, first, merchants and influencers register with an affiliate marketing company, which mediates their relationship. Next, influencers drive sales to the merchant through the affiliate marketing company. This is done by placing customized URLs or website links called affiliate URLs, or coupon codes that are published by the affiliate marketing company in influencers' content. These custom links or coupon codes e.g., '*CODE FOR 5% OFF: MELON*', enable the affiliate marketing company to track sales from customers

[24]. UTM query parameters in URLs are analytics tracking parameters that are used to quantify audience characteristics and to track referrals [14]. Finally, for each tracked sale, the merchant pays the influencer a portion of the sale through the affiliate marketing company. There are few studies of affiliate marketing, notably, in the security community focused on detecting cases where content creators defraud affiliate marketing programs, e.g., through setting fraudulent cookies along the tracking chain of links [28, 29, 5] or by typosquatting [18]. In our paper, we focus on how to automatically disclose in real-time when affiliate marketing content is present.

Disclosures in Online Endorsements

According to the FTC’s endorsement guidelines [11], advertisers are required to disclose their relationships with merchants to consumers so that they can recognize and assess the content in an informed way. These disclosures need to comply with the *clear and conspicuous* standard so they are identifiable to consumers. Specifically, affiliate marketing disclosures need to be placed close to the endorsement and the URLs included by the content creator. Using the text *affiliate link* as a disclosure statement is insufficient; instead, using an explanatory phrase such as *I get commissions for purchases made through links in this post* is encouraged.

Although the HCI community has studied social media for many years (e.g., [35, 12, 36]), only a few studies focus on understanding how users process and deal with advertising content. For instance, there are studies of online behavioral advertising and ad-blocking tools [33, 22]. Outside of the HCI community, researchers have examined whether influencers make advertising disclosures and whether users notice and understand these disclosures [3, 37]. These studies have found that using the text *paid ad* was effective in helping users identify sponsored Instagram post [9] and that users became more resistant to bloggers’ endorsements when bloggers disclosed their sponsored content.

One prior study focuses directly on measuring the prevalence of affiliate marketing content and how users perceive affiliate marketing ad disclosures. Collecting a dataset of 500K YouTube videos and ~1M Pinterest pins, Mathur et al. [21] discovered that 90% of affiliate links are undisclosed on these platforms and that those that did disclose used what the authors call *affiliate link* style disclosures. They highlighted the need for tools that alert users about endorsement-based advertising. We build directly on this work to create automatic detection techniques for disclosing affiliate marketing content to users.

Automatically Identifying Advertisements

Automatically disclosing online endorsements in real-time is an under-studied phenomenon with most tools focused on identifying and blocking malicious or privacy-infringing advertisements on the web [22]. Ad-blocking extensions like Adblock [2] and Adblock Plus [1] maintain *filter* lists of ad servers and user interface elements that correspond to advertisements, which they use to block those advertisements. Other tools detect online behavioral ads using lightweight computer vision and image processing techniques [31]. However, no such parallel—to the best of our knowledge—exists for

endorsement-based advertisements. Traditionally, it has been hard to identify these advertisements because we did not understand their identifying features until the Mathur et al. [21] study. Mathur et al. [21] note that users do not always notice or understand disclosures and that most videos do not have a disclosure. Also, Mathur et al.’s dataset consists of a random sampling of videos which may not reflect a set of videos real users may watch. We built on their work to create our AdIntuition browser extension which can immediately analyze any video in real-time for affiliate marketing content using features from this work and additional ones we identified. Our work also contributes by alerting users to this content automatically and showing how users react to automatic ad disclosures for YouTube videos. Finally, in our work we collect data on and test our tool with actual videos watched by real users.

METHODS

In this section, we describe our detection techniques for affiliate marketing content on YouTube. We then describe the AdIntuition tool design and implementation. We also describe the user studies conducted to evaluate AdIntuition, a survey of 300 users on Amazon Mechanical Turk, a diary study with 11 users, and a field deployment with 472 users.

Automatically Detecting Affiliate Marketing Content

We focus on detecting affiliate marketing since affiliate marketing links, provided by a brand, are ground truth for a brand-content creator relationship. Other types of sponsored content such as native advertising do not have a similar ground truth to detect and verify brand-relationships automatically. Our goal is to expand on the set of signals that can increase confidence in detecting affiliate marketing content, focusing specifically on UTM query parameters and coupon codes as additional signals. First, we examined which UTM query parameters and coupon code patterns are associated with affiliate marketing content to verify whether these features could serve as a ground truth in detection. We used the data set of 515,999 YouTube videos from Mathur et al. [21] and found that 1.2% of all the videos in the original data set contained either a known affiliate link pattern, UTM query parameter in a URL, or coupon code. We performed a manual inspection of these features in the video descriptions. UTM query parameters that appeared frequently included “utm_source=”, “utm_term=”, “campaignid=”, “utm_campaign=”, “utm_content=”, “aff_id=”, and “utm_medium=”. These often contain textual values related to affiliate marketing, such as ‘aff’ or ‘affiliate’.

URL/UTM Parameter Detection Module In AdIntuition

For each link found in a YouTube video’s description, our detection module in AdIntuition checks whether the link, including any intermediate redirected link(s), matches against the known list of affiliate marketing patterns [19] or contains the UTM parameters of interest. If the link or any link in the redirect chain is found to contain one of these two features, then the relevant URL is highlighted in the video description to show that it is an affiliate marketing link.

Coupon Code Clustering and Classification

We also built a classifier to detect coupon codes in video textual descriptions.

Finding Coupon Codes In Existing Data Set: Starting with the 515,999 videos in the Mathur et al. [21] dataset, we found that 174,885 had descriptions that were in English. We tokenized these descriptions into 1,139,880 individual sentences. To shrink the number of inputs to our clustering algorithm, we generated a unique set of potential coupon codes for each video based off of the content creator’s channel name. For example, Casey Neistat is a famous YouTuber whose channel name is ‘CaseyNeistat’. Our dictionary generator checked the description of his videos for his full channel title, ‘CaseyNeistat’, individual components of his channel title split on lowercase to uppercase transitions or spaces, “Casey” and “Neistat”, and then the initials of his channel title components, “CN”.

Using the dictionary generating technique, we found that 170,473 sentences contained a match. Upon closer inspection, we noted that many of these sentences did not contain coupon codes because content creators tend to use the same username across social media platforms. For example, Casey Neistat’s twitter page is @CaseyNeistat and his Instagram name is @caseyneistat, both of which correspond to a case-insensitive search term that was generated in the list of possible coupon codes. We manually labelled 129 sentences as containing coupon codes but only 31.8% of these coupon codes were related to the channel name.

Clustering Sentences With Coupon Codes: We filtered the data set to sentences that might contain a coupon code using our coupon code generator. Using a bag-of-words representation of all of the sentences, we computed the cosine distance between each representation and all of the other sentences’ representations. We used the cosine distance so that long sentences and short sentences would not necessarily be farther apart from each other, as can be the case with Euclidean distance calculations. The clustering algorithm automatically chooses a threshold to cluster the data into distinct sets. Using this algorithm, we clustered the sentences into 253 clusters. The first author manually inspected each cluster to note which contained most of the sentences with coupon codes which we could use to train the classifier. Figure 1 shows the funnel for clustering and classification of videos descriptions into sentences and then sentences selection. The words and weights shown in this figure are most (positive valence) or least predictive (negative valence) of a sentence containing a coupon code e.g., “checkout”.

Building A Classifier To Find Coupon Codes In New Data: With a cluster of 177 sentences containing coupon codes, we then built a classifier that we could use to determine new coupon codes. We sampled 1,000 other random sentences in addition to the 177 that were in the coupon code cluster to add negative examples to our training dataset. Next, we used a bag-of-words representation for each sentence and assigned each sentence a value of 1 if it contained a coupon code and 0 otherwise. We performed three-fold cross validation to train, validate, and test the classifier [27] as described in the next subsection. For each round of training, validating, and testing, we created a vocabulary based off of the sentences in the training set and then fit the classifier to the bag-of-words

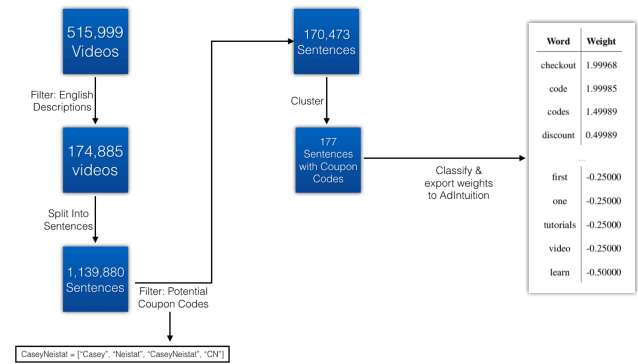


Figure 1. Clustering and Classification Funnel




Figure 2. ROC Curve for the Coupon Code Classifier

representations of those sentences. We then evaluated and tuned the parameters that the classifier used with the validation set. Finally, we quantified the performance of the classifier using the testing set. The performance was evaluated using r-squared score, F1 Score, and ROC curve shown below.

Evaluating The Classifier With Existing Data: Testing different classifiers, we found that Support Vector Classifier (SVC) worked best instead of a Random Forest Classifier or a Decision Tree Classifier, both of which were also tested. We used the F1 score metric to evaluate two important aspects of the classifier: precision and recall. Precision refers to how well the classifier performs at finding sentences that actually contain a coupon code and how well it avoids false positives. Recall refers to how well the classifier is at identifying a large number of the sentences that are known to have a coupon code in them. The classifier had an F1 score of 0.992, which meant that it performed reasonably well given that the maximum score possible is 1. The r-squared value characterizes how well the model’s division of the sets matches up with the data. The SVC model’s r-squared value was 0.976, and the ROC Curve is shown in Figure 2.

Since the classifier used a bag-of-words representation, words were given weights for how important they were in finding coupon codes. Words such as “first” and “video” do not correlate with the presence of a coupon code, whereas “code” correlates strongly. Though trained on sentences that contained coupon codes that were related to channel names, the classification focuses on the words that surround the coupon

code, not the coupon code itself. Out of 129 sentences that we manually labeled to have coupon codes, only 31.8% of them corresponded to the channel name. Meanwhile, the classifier was able to correctly identify 93.8% of the sentences with coupon code (true positive rate), many of which were not related to the channel name.

Integrating The Coupon Code Classifier Into AdIntuition:

Once the classifier was trained, the vocabulary words and associated weights were exported into JavaScript code that is used in the browser extension. In the AdIntuition extension, when a YouTube video loads, the description is parsed into sentences which are transformed into a bag-of-words vector in the same way as before. This vector is then multiplied by the weights vector to get a score, which can be used to determine if a sentence has a coupon code. We used the GridSearchCV library to optimize parameters for the SVC with their suggested threshold value of 1.0 [26].

Designing AdIntuition Automatic Disclosures

We integrated the detection techniques above into our AdIntuition browser extension for Chrome and Firefox, to automatically detect and disclose affiliate marketing content on YouTube videos. We had the following design goals building on findings from prior work [21]:

1. Display a noticeable automatic disclosure (since users often do not notice ones in the video description)
2. Use clear explanatory disclosures (since these are most understood by users)
3. Highlight affiliate marketing links and coupon codes in video descriptions (to aid with alerting users)
4. Distinguish between known affiliate marketing links and ‘suspected’ affiliate marketing content based on UTM parameters and coupon codes (so users are not misled by false positives in our detection techniques)

To meet these goals, AdIntuition displays an unobtrusive but detectable banner directly above a video when it loads if the extension detects any features of affiliate marketing content. We used three colors to distinguish between clear versus known sponsorships to avoid being misleading. Users can click the AdIntuition extension for the color key. For known affiliate marketing links, the banner is pink and states: *‘This video contains affiliate marketing content. The creator may make a commission if you click on the highlighted portions of the description’*. For content AdIntuition flags based on UTM parameters or coupon codes alone, we display a yellow or orange banner respectively and change the wording to ‘contains suspected’ affiliate marketing content to reflect our confidence in the classification. Finally, AdIntuition highlights affiliate marketing links, links with UTM parameters we have flagged, and coupon codes in the video description, also using color to distinguish which flag was set.

Instrumenting AdIntuition To Log Data

We instrumented AdIntuition to log events of interest to an Amazon Web Services (AWS) DynamoDB linked to an API

Gateway to help improve our detection techniques and measure real world usage. Using the API Gateway, AdIntuition logs data from each user to the central database. For each user, we generate a random user id to distinguish between different AdIntuition extension users. This random user id allows us to promote anonymity in the data and minimize risk to individual users. To preserve privacy, we do not log any video watched that is not flagged as affiliate marketing content.

AdIntuition collects information for the following events:

- **add_username**: this event is logged when the user first downloads the extension and when the user id was created. User ids are generated once and then used for the rest of the time that the user uses the extension.
- **vid_watch**: the user watched a video. Video id is logged only if the video was flagged as containing affiliate marketing content.
- **utm**: a URL with UTM parameters of interest was found in the video description. The video id and matched URL are logged.
- **aff**: a URL with a known affiliate link was found in the video description. The video id and matched URL are logged.
- **coupon_code**: a sentence containing a coupon code was found in the video description. The video id and sentence are logged.

AdIntuition Evaluation

We conducted three user studies, a survey, a diary study, and a field deployment to evaluate how users interact with AdIntuition, all of which were approved by our institution’s Institutional Review Board (IRB). In this paper, we only report on findings relevant to AdIntuition usage.

Study 1: Evaluating AdIntuition Interface

As part of the design process, to evaluate if the AdIntuition banner was noticeable and understandable to users, we designed an online survey. This survey contained 17 questions based on Mathur et al.’s user study of how users interpret disclosures on YouTube [21]. We recruited 300 participants on Amazon’s Mechanical Turk (MTurk) service to participate in the survey between April 24th and May 1st, 2019. Participants were required to be at least 18 years old, be in the United States, and have an MTurk score of 95% approval or higher. They were paid \$1.25 for a maximum of 15 minutes of work, which was calculated using the minimum federal hourly wage in the U.S. of \$7.25 and dividing it by 6.

Participants were asked to watch one of three YouTube videos, each of which had a known affiliate marketing link pattern and a product being marketed. For instance, one creator described a new cinnamon donut flavored cookie and provided a link to buy the cookies in the video description. All participants were split randomly into two groups. Users in the control group were shown one of the YouTube videos with no AdIntuition banner. The other group was shown one of the YouTube videos with an automatic pink disclosure and highlighted link as it would appear in AdIntuition. Note, we only tested the pink

Group	Video 1	Video 2	Video 3	Total
Control	62	39	48	149
Treatment	50	50	52	151

Table 1. Number Of Participants In Each Condition (Study 2)

banner to ensure users noticed and understood our design in a controlled setting. We had 6 total conditions with a control and treatment group for each of three videos. The survey first asked about users social media usage. Next, participants were asked to watch their randomly assigned YouTube video and to read the video's description. Participants were unable to continue the survey until they had stayed in this section for at least 3 minutes to ensure they watched the video. Participants were asked to describe the video that they watched, rate their opinion of video content, and to provide reasons for their answers. After this section, participants were shown the product that was described in the video and asked their impression of the product. They were then asked to rate how likely they felt a relationship between the content creator and the organization selling the item existed and why. Following this step, participants were told that the video did contain affiliate marketing content and asked if a banner or a highlighted link would or did assist them in determining the relationship. Finally, we collected demographic information.

Analysis: We conducted a descriptive analysis on the survey data. In addition, we performed inductive thematic analysis [25] on the open-ended answers using a codebook of 8 codes that we created after reading participant responses and team meetings. We met regularly to discuss themes. Examples of codes include 'influence of the AdIntuition banner' and 'knowledge of YouTube sponsorship trends'. Each coded survey response was reviewed by at least two team members and then we wrote summaries of emerging themes for each code. After multiple team meetings, we reached consensus on the main study themes. Study 2 participants are denoted with the identifier *AMT* and a participant id.

Participant Demographics: 300 participants were randomly assigned to a group and had valid responses. Table 1 shows the number of participants in each video and group combination. The median age for all groups was 37 except for the control group for video 1 which was 40. 94% of users in each group reported using YouTube several times a week, 75% also said at least once a day. Across all 300 participants, 45% were female, 54% were male, and 1% did not disclose gender. Ages ranged from 23 to 73 years old with a median of 37 years old. 97.3% of participants used YouTube at least once a week, with 57.7% reporting using it several times a day.

Study 2: Evaluating AdIntuition User Experience With Two Week Diary Study

To evaluate the user experience of AdIntuition, we conducted a diary study [17] in July-August 2019. This study allowed us to interview users about their real-world encounters with and perceptions of AdIntuition banners with all three colors and their effects on user opinions. We recruited 12 participants through online mailing lists and social media postings. We recruited participants who were regular YouTube users over 18 who lived in the United States and had access to a device and browser compatible with the extension (Google Chrome

or Firefox). We defined a 'regular' YouTube user as anyone who watches videos on YouTube multiple times a week.

The diary study was divided into 3 parts: pre-study interview with installation of the tool, diary log period, and post-study interview. Each interview lasted up to 30 minutes. The interviews were conducted either in-person at our institution's campus or through video/audio call (Skype or Google Hangouts). All interviews were audio-taped. During the pre-study interview, we gathered baseline data: how often/what YouTube videos users watched, if they ever encountered a video they thought was sponsored before, why they thought that, and how it affected them. After the interview, the researcher helped the participants install the tool onto the participant's personal device, showed them how to use it, and verified it was working from the user_id generated by the extension. After this interview, participants were asked to complete diary log entries in a paper or electronic diary for the next 10 days. Each day, participants were asked to log information on the videos they watched on YouTube that contained an AdIntuition disclosure banner or note if they did not see any videos with these banners. Each diary log asked for video title and channel name, a rating of level of surprise at seeing a banner above the video indicating affiliate marketing content on a Likert scale, and reasons for the rating.

During this diary period, participants were instructed to watch videos as usual. To incentivize logging, we sent participants daily email reminders and compensated them \$2 for each day that contained at least 1 diary entry (in addition to the \$25 for participants in the pre and post study interviews for a total of \$45 for 10 days maximum.) In the post-study interview, participants were asked about their experiences with AdIntuition, the tool design, and their opinions on affiliate marketing on YouTube. Participants were also shown how to uninstall the tool but reminded that they could continue using the tool if they wanted beyond the study. If participants did not encounter any affiliate marketing content during the diary logging period, we showed them an affiliate marketing video with an AdIntuition disclosure banner before asking them the questions.

Analysis: All diary logs were converted to a digital format in an Excel spreadsheet for analysis. We transcribed the interviews and developed two codebooks based on the interview guide and diary instructions which we refined after reviewing several transcripts, diary logs, and team discussions. We used the same analysis process as Study 1 to perform inductive thematic analysis [25] on the transcripts and diary logs. Two co-authors went through and coded the transcripts, met weekly with the team, and refined the themes based on points of agreement. We had a total of 80 codes (70 for interviews, 10 for diary logs). Example codes for interviews included 'affiliate marketing knowledge', 'attitudes towards affiliate marketing disclosures', and 'thoughts on the AdIntuition design'. Example codes for the diaries included 'types of videos watched' and 'reasons for feeling surprised to see the banner above a video'. We denote Study 2 interview quotes with *P* and quotes from participants' diary logs with *PDL* and the participant id.

Participant Demographics: Only 11 participants completed the diary study (7 female, 4 male) for a minimum of 10 days.

The age range was 19 to 71, with more than half the participants being under 35. Participants had a variety of occupations including undergraduate and graduate students, housewives, a librarian, an administrator, a manager, a director, and a retiree. All but one participant reported watching YouTube videos daily with 1 saying they watched videos several times a week. Over the study period, 9 participants completed 108 diary log entries with a median of 10 entries per participant. 3/9 came across the banners every day and 6/9 saw the banners only on some days. Participants who saw banners during the study period saw the following types of videos: how to/educational, entertainment, vlogging/video personalities, sports/video games, product reviews, style, and music. The remaining 2 participants did not come across any videos with AdIntuition banners during the study even though they watched multiple videos a day. Therefore these participants did not have any diary logs.

Study 3: AdIntuition Field Deployment With Real Users

To evaluate how effective our classifier's efficacy with videos real users watch in real-time, we conducted a field deployment. In this study, we were unable to interview participants since we did not collect personally identifiable information to enable a large deployment. We recruited users to test out a tool 'to identify misleading ads on YouTube' through a blog on our institutional website, Twitter, and an article in the popular media. To preserve user privacy, we created a privacy policy detailing everything the tool collects. We also built functionality for users to download a report of all of the data that AdIntuition collects for them and to delete entries in that record. Finally, we added options for users to opt into or out of data logging. Overall, we had a total of 472 downloads on the Chrome and Firefox stores for the period June 5th-August 26th, 2019.

Analysis: We assumed users who had just downloaded AdIntuition would be interested to see how it worked and therefore would have the potential to skew our analysis. For any metric that had to do with individual usage of the extension, we removed all users who used AdIntuition for 1 day only and considered users as an active user only if they used AdIntuition for at least 1 additional day beyond the day of download. In our analysis, we focus only on these *active* users. We also removed all AdIntuition diary study users from our analysis. We only report on videos with affiliate marketing content which we verified by manual inspection and we *exclude false positives* in all graphs and data presented.

Limitations

Our survey study evaluated a simulation of the AdIntuition interface so the results may differ if we used experience sampling with our real world deployment instead (which also presents challenges of asking users to respond in context). Further, our diary log study is limited by self-report and the sample size. Finally, to preserve privacy in the field deployment, we only collected video ids for videos AdIntuition flagged as containing affiliate marketing content. Thus, we are unable to calculate false negatives for our detection techniques.

FINDINGS

Study 1 and 2: User Experiences With AdIntuition

We report on the findings across Study 1, the survey, and Study 2, the diary study.




Figure 3. Perceived Likelihood Of Content-Creator-Organization Relationship In AdIntuition Survey




Figure 4. Reasons For Perceived Content-Creator-Organization Relationship In AdIntuition Survey

AdIntuition Helps Users Identify Affiliate Marketing Content

Our first study suggested that AdIntuition's interface helped participants to identify affiliate marketing content. As shown in Figure 3, the experimental group (with AdIntuition disclosure banner and highlighted affiliate link) was much more likely to perceive a relationship between a content creator and the organization selling merchandise in the video than in the control group which did not see an AdIntuition banner. For instance, 71.4% in the experimental group that saw the AdIntuition disclosure banner above the video and the highlighted affiliate marketing link reported that this relationship was likely or extremely likely compared to 56.3% of the participants in the control group who did not see any AdIntuition banner when they watched a video. We calculated an odds ratio of 1.941 with $p < 0.00628$ suggesting this result was not owing to random chance.

Figure 4 summarizes the reasons participants provided for their answers. Notably, 18.18% of participants in the experimental condition explicitly reported their reason for believing a relationship is likely is because of the AdIntuition banner. For instance, participant AMT196 stated *'There was a message above the video stating that it included affiliate links. More than likely this was a paid sponsorship by Chips Ahoy'* (None reported this in the control condition since they saw no banner). More participants, 17.88%, also expressed that they were unsure or needed more information to answer the question in the control condition versus 9.74% in the experimental condition. Interestingly, a similar number of participants mentioned other *implicit* indications of an affiliate marketing relationship in both conditions (32.45% control vs 28.57% experimental) such as noticing links in the video description for the products (e.g., *'Given that a link pointed to this specific item that*

appeared to be used in the video, I'm guessing the presenter is related to DJI' (AMT348), the behavior of the content creator, or video characteristics or expectations of marketing content on YouTube (e.g., 'It's a pretty common thing among influencers' (AMT245).

Disclosures Most Unexpected If Implicit Signals Are Absent

In our second study, across all users, 26% of the total diary logs indicated that participants were surprised that a particular video had affiliate marketing, in 43% of the logs participants indicated they were not surprised to see an ad disclosure, and 31% of the entries were rated as a neutral. Often, participants logged that they were surprised to see the AdIntuition banner when the product was not aligned with the type of video or video creator. For example, one participant described being very surprised to watch a video where there was advertisement for something that contradicted the entire premise of the video: "I didn't expect a food ad from a water fasting video, found it counterintuitive" (P10DL6) Other reasons for being surprised included: unclear product promotion, the sponsorship not being typical of the channel or influencer, the channel having a small following, or containing little to no indication of a sponsorship in the video as summarized by Participant P3 (DL4): "I was surprised because I didn't notice the sponsorship at all while watching the video. It was very subtle".

For some videos, participants did not report feeling surprised to see they contained affiliate marketing content. Their reported reasons in the interviews and diary logs largely matched the implicit indicators mentioned by participants in Study 1. In these cases, participants mentioned noticing a video was sponsored, or more implicit indicators, such as the influencer talking very favorably about a specific product in comparison to other products, talking about product very positively or for a significant amount of time of the video, the frequency of mentioning a product, a clear focus on the product, product placements disrupting the natural flow of the video, and talking 'commercial-like'. An example diary log summarizing these reasons include: "She regularly promotes her merch that she sells (sometimes to raise money for charity). She also just released her own line of nail polish" (P1DL26). In these logs, participants reported feeling neutral when the sponsorship made sense. Notably, the majority of participants did not understand the differences between the different color banners.

Disclosures Made Users More Reflective On Content Creators

In Study 2, we were also curious to see if participants noticed any changes in their behavior while using AdIntuition. Less than half of the participants reported no change in their YouTube watching behaviors as a result of using the extension. However, 6 of 11 participants found that using AdIntuition made them think more about sponsorships on YouTube and made them more aware and perceptive of them. Some participants even told us that they found themselves actively searching for additional videos to see which videos would have a banner, trying to see which product would be promoted in the video, and how the influencer would promote the product. As an example, participant P2 summarized: "I did notice a difference in the way I viewed certain videos. As I was watching them I was definitely more sort of clued in, so looking for any




Figure 5. Total Affiliate Marketing Videos By Feature Type

clues as to what they were advertising. [In] some videos [it] is really obvious but some videos where it may be is more subtle then quite often, I found myself trying to watch to figure out where they were advertising stuff".

All but one participant reported some change in their attitude or perception of sponsored content as a result of using AdIntuition. Some participants expressed more negative attitudes towards videos with the AdIntuition banner, such as noting that the intent of the video was unclear or that those videos were of lesser (subjective) quality in comparison to videos without banners from the same video creators. Some participants also described negative attitude changes towards the influencers, such as a decrease in trust, questioning the influencer's true intentions behind creating specific content, and feeling disappointed in the influencer, or that 'whoever was producing the video didn't have my best interests in heart' (P11). As for the products being promoted, a few participants expressed an increase in skepticism of the legitimacy of the endorsements and true quality of the product in videos that had the AdIntuition banner. Participants also felt deceived by unexpected banners as expressed by P4: 'I felt disappointed in them even though I don't know them, but I didn't change what I watched.'

However, other participants expressed positive attitude changes such as excitement when coming across a video with a banner, feeling more intrigued by the video, and gaining new insights into the business aspects of YouTube content creators. This is captured in a quote by Participant P2: "The banner made me realize that it's much more of a business than just purely people just having fun. But I don't think it made me think any better or worse of the people who are making them". These changes contrasted with pre-study sentiments which were more neutral overall regarding sponsored content.

Study 3: Performance Of Automated Detection Methods

Next, we report findings from the field deployment, Study 3.

Multiple Features Aid Affiliate Marketing Automatic Detection

In total, 472 AdIntuition users saw a total of 60,835 videos over the deployment period of 82 days. Of the total number of videos watched across all users, (some of which were watched more than once), 6071 were true positives and contained one of the three indicators of affiliate marketing content. Of this total

utm_source	utm_medium	utm_content
youtube (1424)	paid (326)	cta-link (94)
affiliate (360)	social (278)	[no value] (51)
open.spotify.com (75)	youtube (189)	banggoodtv (15)
blogger (51)	affiliate (188)	description (15)
youtube.com (48)	yt+main (186)	en (14)
impact_radius (40)	yt main (115)	dhalucard (14)
yt (40)	referral (113)	tw (12)
afc-linus+media+group (36)	video (97)	linustechtips (10)
vlogger (31)	description (84)	descrip (9)
refersion (31)	open (75)	youtube (9)

Table 2. Most Common UTM Parameter Values For Various UTM Query Parameters With Counts of Occurrence

number of affiliate marketing videos watched across users, there were 4494 total unique videos. Figure 5 shows the breakdown of which features these unique affiliate marketing videos were flagged on. Note, this shows cases where a video was flagged on more than one feature.² Of the total flagged videos (including overlap between videos with more than one feature flagged), 71.7% had at least one link with a known affiliate link pattern, 14.33% had at least one link from a domain confirmed to use UTM parameters in their affiliate marketing campaigns, and 29.02% had at least one coupon code in the video description.

UTM Parameters: In the unique affiliate marketing videos data set, we found the most commonly occurring values for each of the UTM parameters that we logged and then manually inspected this list as shown in Table 2. We then flagged the ones that appeared to relate to affiliate marketing terms such as “paid” (the most common utm_medium value), “affiliate” (the second most common utm_source and fourth most common utm_medium value), and “referral” (the seventh most common utm_medium value). We also added terms surfaced in the known affiliate marketing links list [19]. This list included the following values that we then searched for: ‘=aff’, ‘=infl’, ‘aff_id=’, ‘aff=’, ‘sponsor’, ‘=paid’, ‘=ref’, ‘ref_id=’, ‘promotion’, ‘ref=’, ‘referral’, ‘affiliate’, and ‘influencer’. Of the total unique links that were logged to contain UTM parameters, 33.24% contained a UTM parameter with one of these tags. These links were tagged 7386 times in total, some of them just once and some multiple times depending on the number of times a video was watched by AdIntuition users.

We filtered down the total list of 929 domains containing UTM parameters we flagged to 28 domains, each of which had at least 10 unique links present in the data set. We then manually checked whether each of these 28 domains on this list had affiliate programs by visiting their home pages to confirm if they offered such programs or in some cases, signing up on affiliate marketing company programs claiming to do affiliate marketing on behalf of these sites. 21/28 domains had affiliate marketing programs. Of the 7 remaining domains, 5 were unclear, and 2 had affiliate marketing programs that were no longer active. As a lower bound, we were thus able to verify that on these frequently occurring 28 domains, AdIntuition

²Only our coupon code classifier is predictive in that it can handle inputs it has not seen before based on pre-determined weights calculated in training, therefore we cannot report on predictive power for each attribute flagged.




Figure 6. Affiliate Marketing Encounters Over Days Active




Figure 7. Individual User Encounters With Affiliate Marketing

had a true positive rate of 75% based on a UTM parameter feature alone. Less frequently occurring domains may also have affiliate marketing programs but further manual verification is required. We have made this list of commonly occurring UTM parameters and affiliate marketing domains available on GitHub to extend the known affiliate marketing patterns made available in [21].

Coupon Codes: AdIntuition flagged a total of 2539 coupon codes in the videos that had affiliate marketing content. We manually reviewed these codes and found that our classifier had falsely tagged 470 codes and that 43 codes were unclear as to whether it was a coupon code. We had a true positive rate of 81.17%, suggesting our classifier had a slightly worse true positive rate on real world data.

Affiliate Marketing Encounters Depend On Viewing Behaviors During Study 3, the median user used the extension for 4 days and saw 2.7 videos per day. The range of total videos watched per user per day between 1 and 123. The median video views per a user was 11. Across all active users, the median percentage of affiliate marketing content seen was 7.55% of videos. As shown in Figure 6, which maps the total affiliate marketing content active users encountered, AdIntuition users who have used the extension for longer periods see slightly fewer videos with affiliate marketing than those who have only used it for a few days. However, as users interact with the extension more, the prevalence increases.

We also plotted the percentage of total affiliate marketing videos encountered by each active user over time. This is illustrated in Figure 7, where each bubble on the graph represents

Content Creator	No Of Videos Watched Per User [User1, User2,...]
Sweet Anita	[91]
iFL TV	[43]
Tim Pool	[35]
FNG	[22]
The Alpha Male Strategies Show	[22]
Cereal Entrepreneur - Jordan Steen	[18]
LetsPlay	[18]
SAM THE COOKING GUY	[17]
Unbox Therapy	[16]
Movieclips	[15, 37]
Linus Tech Tips	[15, 14, 12, 12, 12]
Simply Nailogical	[15]
jade darmawangsa	[15]

Table 3. Most Common Content Creators And No Of Unique Videos Watched From That Creator By Various Users

one user. Noticeably, the number of users who have not seen any affiliate marketing videos (0% prevalence) stops after 11 days of AdIntuition usage. Finally, we examined the most commonly occurring content creators with more than 15 unique videos in our data set and found that the majority of them were only watched by a single AdIntuition user (As shown in Table 3, in an extreme case, 1 user watched 91 videos from ‘Sweet Anita’). Our data demonstrates that users encounter affiliate marketing content very differently over time, likely owing to their viewing preferences and whether content creators they view are affiliate marketers.

DISCUSSION

We make the following recommendations based on our results.

Ethics And Automatically Disclosing Advertising Content

Creating automatic disclosure tools for flagging disguised advertisements or misleading online content in general can cross an ethical line if there are false positives which sway consumer opinions. For example, in our survey and diary study, participants sometimes viewed content more negatively when a disclosure was present. In AdIntuition, to avoid being misleading we used color and wording to convey the confidence in our classification. We need further research to understand how best to inform users about misleading content such as advertisements without having them form overly negative opinions or habituate to automatically generated disclosures. In this way, we can ensure that automated disclosure tools do not themselves perpetuate misleading information or dark patterns in cases of false positives.

Future work could examine combining our approach with crowd-sourcing to have secondary checks and balances on the information provided. Future studies could also work on incorporating a reporting tool for users to mark incorrectly flagged content. These tools should also allow influencers to view their own content with these tools and similarly report inaccurate messages. Finally, future work could also examine the effects of different types of disclosures on user opinions to mitigate negative effects. The need to present information in an unbiased informative manner is one that is common to any system that aims to inform users about misleading online content such as dark patterns, disinformation, and misinformation. Future studies can build on our work to find common guidelines for informing users about misleading content in a neutral fashion for these related domains.

Automatic And Human Aided Disclosure

AdIntuition demonstrates that automatically disclosing one type of online endorsement on social media is possible. However, our work also raises the larger discussion of who is responsible for creating these disclosures about misleading online content? Browsers could integrate automatic disclosure tools into their capabilities so that users can see advertising content more easily across the web. Platforms could similarly integrate this approach or affiliate marketing companies could more strongly require that influencers disclose their relationships. We do see a place for third party automatic disclosure tools too for informing regulators about the prevalence of this type of content in the wild, how often disclosures are happening, and how users react to these disclosures.

An open question is how to maintain a tool for automatic disclosures? Much like an ad blocker [22], lists of known affiliate marketing programs are constantly shifting and changing so our data set needs to be continuously curated and expanded upon. Crowd-sourcing is one method that has been used successfully in Adblockers, this technique could similarly add value to detection and automatic disclosure tools such as AdIntuition. AdIntuition can also constantly be improved as more data is gathered with which to refine its classifier and incorporate other classifiers and detection methods.

Detecting Affiliate Marketing On Other Platforms

Our results show that we can automatically detect affiliate marketing content on YouTube using known affiliate marketing link patterns, coupon codes, and UTM parameters. Future work could extend our approach to automatically detect and disclose affiliate marketing content on other social media platforms or blogs. Our user studies suggest automatic detection and disclosures are useful since they do not rely on content creators to self disclose this information. Another area for future research concerns when to show an automatic disclosure to a user. For instance, participants in the diary study suggested seeing automatic disclosures when searching for content could be useful in certain scenarios such as searching for reviews of a product one is about to buy. Future work could investigate how user experience varies when automatic disclosures are shown at different user decision making points.

CONCLUSION

In this paper, we presented new ways to detect and measure affiliate marketing content on YouTube and a tool for automatically disclosing this content to users on the platform. We also presented findings from evaluating the tool, AdIntuition which suggest that the detection techniques are performing reasonably well and that users are able to better identify advertising content with AdIntuition’s automatic disclosures. Based on our findings, we recommend that future studies extend our work to build more robust online automatic ad detection tools to keep users informed about the content they are viewing.

ACKNOWLEDGMENTS

We thank our participants.

REFERENCES

- [1] 2019. Adblock Plus. <https://adblockplus.org>. (2019).

- [2] AdBlock. 2019. AdBlock. <https://getadblock.com>. (2019).
- [3] Sophie C Boerman, Eva A Reijmersdal, and Peter C Neijens. 2012. Sponsorship disclosure: Effects of duration on persuasion knowledge and brand responses. *Journal of Communication* 62, 6 (2012), 1047–1064.
- [4] Harry Brignull. 2018. Dark Patterns. <https://darkpatterns.org/>. (2018). Accessed March 12, 2019.
- [5] Neha Chachra, Stefan Savage, and Geoffrey M Voelker. 2015. Affiliate crookies: characterizing affiliate marketing abuse. In *Proceedings of the 2015 ACM Conference on Internet Measurement Conference*. ACM, 41–47.
- [6] Amy Chozick. 2019. Keeping Up with the Kardashian Cash Flow. *New York Times* (March 2019). <https://www.nytimes.com/2019/03/30/style/kardashians-interview.html?module=inline>
- [7] Federal Trade Commission. 2017. What’s affiliate marketing? Should I care? <https://www.consumer.ftc.gov/blog/2017/09/whats-affiliate-marketing-should-i-care>. (2017).
- [8] Federal Trade Commission. Accessed 2018. Enforcement Policy Statement on Deceptively Formatted Advertisements. https://www.ftc.gov/system/files/documents/public_statements/896923/151222deceptiveenforcement.pdf. (Accessed 2018).
- [9] Nathaniel J. Evans, Joe Phua, Jay Lim, and Hyoyeon Jun. 2017. Disclosing Instagram Influencer Advertising: The Effects of Disclosure Language on Advertising Recognition, Attitudes, and Behavioral Intent. *Journal of Interactive Advertising* 17, 2 (2017), 138–149. DOI : <http://dx.doi.org/10.1080/15252019.2017.1366885>
- [10] Facebook. 2019. Facebook Ads. <https://www.facebook.com/business/ads>. (2019).
- [11] FTC. 2018. The FTC’s Endorsement Guides: What People Are Asking. <https://www.ftc.gov/tips-advice/business-center/guidance/ftcs-endorsement-guides-what-people-are-asking>. (2018).
- [12] Eric Gilbert and Karrie Karahalios. 2009. Predicting tie strength with social media. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 211–220.
- [13] Google. 2018. Paid product placements and endorsements. <https://support.google.com/youtube/answer/154235?hl=en>. (2018).
- [14] Google. 2019. Collect campaign data with custom URLs. <https://support.google.com/analytics/answer/1033863>. (2019).
- [15] Colin M. Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L. Toombs. 2018. The Dark (Patterns) Side of UX Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI ’18)*. ACM, New York, NY, USA, Article 534, 14 pages. DOI : <http://dx.doi.org/10.1145/3173574.3174108>
- [16] V Kumar and Rohan Mirchandani. 2012. Increasing the ROI of social media marketing. *MIT sloan management review* 54, 1 (2012), 55.
- [17] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2010. *Research Methods in Human-Computer Interaction*. Wiley Publishing.
- [18] Victor Le Pochat, Tom Van Goethem, and Wouter Joosen. 2019. A Smörgåsbord of Typos: Exploring International Keyboard Layout Typosquatting. In *Proceedings of the 4th International Workshop on Traffic Measurements for Cybersecurity*. IEEE Computer Society Press.
- [19] Arunesh Mathur. 2018. Known Affiliate Marketing Link Patterns. https://github.com/aruneshmathur/affiliate-marketing-disclosures/blob/master/affiliate_markting_links.txt. (2018).
- [20] Arunesh Mathur, Gunes Acar, Michael J. Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 81 (Nov. 2019), 32 pages. DOI : <http://dx.doi.org/10.1145/3359183>
- [21] Arunesh Mathur, Arvind Narayanan, and Marshini Chetty. 2018a. Endorsements on Social Media: An Empirical Study of Affiliate Marketing Disclosures on YouTube and Pinterest. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 119 (Nov. 2018), 26 pages. DOI : <http://dx.doi.org/10.1145/3274388>
- [22] Arunesh Mathur, Jessica Vitak, Arvind Narayanan, and Marshini Chetty. 2018b. Characterizing the Use of Browser-Based Blocking Extensions To Prevent Online Tracking. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. USENIX Association, Baltimore, MD, 103–116. <https://www.usenix.org/conference/soups2018/presentation/mathur>
- [23] NYTimes. 2017. Fyre Festival, a Luxury Music Weekend, Crumbles in the Bahamas. <https://www.nytimes.com/2017/04/28/arts/music/fyre-festival-ja-rule-bahamas.html>. (2017).
- [24] Richard L Oliver and Mikhael Shor. 2003. Digital redemption of coupons: Satisfying and dissatisfying effects of promotion codes. *Journal of Product & Brand Management* 12, 2 (2003), 121–134.
- [25] Johnny Saldaña. 2013. *The Coding Manual for Qualitative Researchers* (2nd ed ed.). SAGE, Los Angeles.
- [26] scikit learn developers. 2007-2019a. *sklearn.model_selection.GridSearchCV*.

- [27] scikit learn developers. 2007-2019b. *sklearn.model_selection.KFold*.
- [28] Peter Snyder and Chris Kanich. 2015. No Please, After You: Detecting Fraud in Affiliate Marketing Networks.. In *WEIS*.
- [29] Peter Snyder and Chris Kanich. 2016. Characterizing fraud and its ramifications in affiliate marketing networks. *Journal of Cybersecurity* 2, 1 (2016), 71–81.
- [30] Kate Starbird, Ahmer Arif, and Tom Wilson. 2019. Disinformation As Collaborative Work: Surfacing the Participatory Nature of Strategic Information Operations. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 127 (Nov. 2019), 26 pages. DOI: <http://dx.doi.org/10.1145/3359229>
- [31] Grant Storey, Dillon Reisman, Jonathan Mayer, and Arvind Narayanan. 2017. The future of ad blocking: An analytical framework and new techniques. *arXiv preprint arXiv:1705.08568* (2017).
- [32] Twitter. 2019. What are Promoted Tweets? <https://business.twitter.com/en/help/overview/what-are-promoted-tweets.html>. (2019).
- [33] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *proceedings of the eighth symposium on usable privacy and security*. ACM, 4.
- [34] Eva A van Reijmersdal, Marieke L Fransen, Guda van Noort, Suzanna J Oprea, Lisa Vandeberg, Sanne Reusch, Floor van Lieshout, and Sophie C Boerman. 2016. Effects of disclosing sponsored content in blogs: How the use of resistance strategies mediates effects on persuasion. *American Behavioral Scientist* 60, 12 (2016), 1458–1474.
- [35] Aditya Vashistha, Abhinav Garg, Richard Anderson, and Agha Ali Raza. 2019. Threats, Abuses, Flirting, and Blackmail: Gender Inequity in Social Media Voice Forums. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, Article 72, 13 pages. DOI: <http://dx.doi.org/10.1145/3290605.3300302>
- [36] Elham Vaziripour, Devon Howard, Jake Tyler, Mark O'Neill, Justin Wu, Kent Seamons, and Daniel Zappala. 2019. I Don'T Even Have to Bother Them!: Using Social Media to Automate the Authentication Ceremony in Secure Messaging. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, Article 93, 12 pages. DOI: <http://dx.doi.org/10.1145/3290605.3300323>
- [37] Bartosz W Wojdyski. 2016. The deceptiveness of sponsored news articles: How readers recognize and perceive native advertising. *American Behavioral Scientist* 60, 12 (2016), 1475–1491.
- [38] Katrina Wu. 2016. YouTube marketing: Legality of sponsorship and endorsements in advertising. *JL Bus. & Ethics* 22 (2016), 59.
- [39] David Zax. 2017. The War To Sell You A Mattress Is An Internet Nightmare. <https://www.fastcompany.com/3065928/sleepopolis-casper-bloggers-lawsuits-underside-of-the-mattress-wars>. (2017).