The Journal of Web Science, 2018, 4: 1–19

Towards Understanding the Consumption of Video-Ads on YouTube

Mariana Arantes[1], Flavio Figueiredo[1] and Jussara M. Almeida[1]

[1] *Universidade Federal de Minas Gerais - Brazil*

ABSTRACT

Being the most popular online video platform nowadays, YouTube is a complex ecosystem that generates billions of dollars of revenue yearly. This revenue mostly stems from online advertisements that are shown on the website. Like other social media platforms, YouTube enables any user to create and upload content, create ad-campaigns that promote advertisement content, as well as monetize channels (i.e., YouTube video uploaders) by showing ads from other channels to viewers. More importantly, any individual can watch videos for free and, in consequence, be exposed to advertisements. The mediation of these different parties that interact through ads, as well as the YouTube platform itself is done by online ad auction algorithms. In this paper, we study the aforementioned ecosystem through the use of advertisements in the form of video (video-ads). Online video-ads are a novel medium that is gaining significant t raction o n s ocial m edia p latforms l ike Y ouTube. O ur s tudy p resents insights on (1) the behavior of users when exposed to video-ads; (2) the popularity of the video-ads over time; (3) the relation between contextual advertising and the effectiveness of ads; (4) the success of ads in generating revenue; and, (5) the success of channels in attracting revenue as exposers of ads. The results here presented have practical implications for content providers, creators, channels, and YouTube viewers.

Keywords: Video Ads, YouTube, User Behavior, Popularity

ISSN 2332-4031; DOI 10.1561/106.00000011

1 Introduction

Web auctions and online advertisements are ubiquitous on the Web. Nowadays, different websites, and in particular social media platforms, provide free access to content in exchange for user attention (Ribeiro, 2014). In order to generate revenue, these platforms exchange user data and attention with advertisers that ultimately promote their brands and content to end viewers. What is more interesting is that social media platforms enables any user to play the role of: viewer, content creator or advertiser/marketer.

YouTube, the most popular social media video sharing platform nowadays, is a successful example of application that relies mainly on advertisements to generate revenue. Users can have access to a wide range of content (videos) in the website for free, while YouTube creates ad auctions to select advertisements to be exhibited to users. Despite watching videos, users in the website can also contribute creating original content and even being an advertiser that runs ad campaigns. Thus, the application is an example of new advertising market, where every type of user can benefit from ads. On one hand, the common users, viewers, can have a better experience when the ads exhibited to them are related to their interests. On the other hand, the users that also contribute to the creation of videos in the website can earn monetary shares for ads associated to their content. Finally, the users that market their videos to be advertised in the website gains from well placed ads that are able to capture the interest of the audience with YouTube itself profiting from a fraction of the successfully monetized ads.

With this complex ecosystem in mind, we here focus on a specific type of advertisement on YouTube, the ads presented to the user in the form of videos (video-ads). Our goal is thus to understand this novel ad market. Video-ads are rising as one of the most important means of online publicity (Variety, 2015; Variety, 2014; Wall-Street-Journal, 2014) and this study characterize properties of video-ads, how they are consumed by users and how the revenue is spread across content creators.

In details, we study YouTube video-ads in light of five different research questions that aim at understanding different aspects of the aforementioned ecosystem. Our five research questions (RQs) and major findings are:

RQ1: How do users consume video-ads? We start our study by analyzing the behavior of users when exposed to video-ads. Whenever a video-ad is displayed to a user on YouTube, the user may be allowed to skip the advertisement after some initial seconds, jumping directly to the requested video. Therefore, here we analyze the "skipping" behavior of users with the aim of understanding how users often react to advertisements, what is the fraction of the video-ads that are streamed before the skip and whether users tend to skip video-ads or consume them in full. We also draw insights into how effective these ads are in terms of attracting attention of users, in particular when compared to other forms of online advertising (Schneider *et al.*, 2009; Krishnan and Sitaraman, 2013).

RQ2: How does video-ad popularity evolve over time? We use the distributions of the *number of views* and *exposure time* in order to understand properties of video-ad popularity.

Since the former captures the amount of accesses to each video-ad and the latter captures the amount of time that users were exposed to its content, they have both been used as a measure of the success of ad campaings (Farahat and Bailey, 2012; Drèze and Zufryden, 2004; Ghose and Yang, 2010; Manchanda *et al.*, 2006). In RQ2, we improve our analysis of the effectiveness of video-ads on YouTube looking into how bursty the popularity evolution of video-ads is, the time it takes for a video-ad to peak in popularity, and the different profiles of popularity evolution.

RQ3: What are the relationships (if any) between a video-ad and the video-contents with which it is associated? In ad-auctions, a video-ad is paired with a piece of content (a YouTube video in our case, or simply a video-content) to be displayed to the user. In RQ3 we have two goals. First, we analyze whether more popular video-ads tend to be paired with videos that are also very popular. Secondly, we assess the extent to which video-ads that are paired with more similar content have a tendency to be more effective (popular), thus uncovering evidence of whether contextual advertising (Lacerda *et al.*, 2006) increases the effectiveness of video-ads.

RQ4: Which factors lead to video-ads being monetized? In order to provide a good value to advertisers, YouTube does not charge for every exhibition of video-ad on the website. The advertiser only pays when the user shows some level of interest in the video-ad, which is measured by the amount of time the video-ad is streamed. In RQ4, we deepen our understanding of the video-ad market on YouTube by looking at the exhibitions that generated revenue. We look into the fraction of monetized exhibitions as a whole and per video-ad. Our results here provide a look into the role of the video-ad itself on it's success (more monetized exhibitions). We also present estimate of YouTube monetization as whole.

RQ5. How successful are channels in attracting revenue? YouTube allows any user to create content an earn monetary shares for video-ad exhibitions associated with their videos. All videos published by the same content creator will belong to the same channel, which is the home page for the user account. Thus, in RQ5, by analyzing the revenue from the perspective of the channels, we are focusing on the content creators that were able to profit from video-ads associated with their contents. Here we analyze the popularity of channels and the success of content creators in profiting from ads.

Towards answering these questions we relied on the combined use of datasets from two rich sources. On one hand, we made use of a dataset of HTTP requests originating from a major university campus network in Brazil. Using this dataset, we were able to obtain the unique identifiers of a number of video-ads and the video-contents with which they were paired. We then crawled the API information on each such video-content and video-ad, as well as the public statistics datasets that contain daily time series of the number of views and the exposure time of each video-ad. The main contribution of this study is to provide an in-depth view of different properties of video-ads on YouTube, notably user consumption behavior, popularity, ad-content pairings, monetization, and successful channels. Our key findings can be summarized as follows:

- Users often skip video-ads as early as possible, and, on average, 20% after the video-ad content has been streamed. Yet, a considerable fraction (29%) of all video-ad exhibitions in our datasets are streamed until completion (RQ1).
- Video-ad popularity is heavy tailed in nature (with most ads attracting few views). Moreover, for most video-ads, popularity is concentrated on a few days, reaching its daily peak shortly after upload, although some of the very popular video-ads remain attractive for longer periods of time. Indeed, we found six different profiles of video-ad popularity evolution (RQ2).
- Video-ad and *aggregated* video-content popularity are strongly correlated, as popular contents tend do attract more users to ads, although the correlation is not as strong for individual video-contents. Yet, the content similarity between video-ad and video-contents in each ad-to-content pairing tends to be very small. Moreover, we found only weak evidence that more similar pairings tend to lead to more popular video-ads (RQ3).
- A considerable fraction of video-ads were successful in generating revenue to YouTube and content creators, although the contribution of each one in particular was small. Also, successful video-ads tend to have a short duration and the categories of the video-ads are related to their chance of achieving success (RQ4).
- A considerable number of channels were able to profit from YouTube, even though most of them were associated to video-ads just a few times. We also found that in general, the video-ads exhibited in each channel are very diverse (RQ5).

The main contribution of this study is to provide an in-depth view of different properties of video-ads on YouTube. Our findings offer a *novel*, *broad* and *timely* look into the ecosystem of video advertisements, drawing valuable insights that motivate the design of more cost-effective strategies to make online video-ads potentially more profitable. Such insights should be of interest to content producers, content providers and marketers, who financially benefit from the success of ad campaigns. Our findings are also of interest to YouTube users in general since they are subject to video-ads.

The rest of this paper is organized as follows. In Section 2 we present our discussion on related work. Next, we provide an overview of the YouTube advertisement ecosystem in Section 3. Our datasets are discussed in Section 4. Sections 5 to 9 present our findings in RQs 1 through 5 respectively, with each research question being discussed individually on a single section. Finally, Section 10 provides the implications of our findings and directions for future work.

2 Related Work

In this section we present our discussion of previous efforts related to our work. Before continuing, we point out that most previous efforts on online advertisements have not yet studied video-advertisements as we do. Only recently is it that video-ads have gained the attention of researchers and practitioners in Web science. Throughout this section, we begin our discussion on online advertisements in general. We then provide a more in-depth discussion of those efforts related to video-advertisements. Finally, we also mention previous efforts on TV commercials, since they are also presented in the video format.

2.1 Online Advertising

It is well known that advertisements are ubiquitous on today's Web economics (Abraham *et al.*, 2013; Amarie *et al.*, 2014a; Amarie *et al.*, 2014b; Carrascosa *et al.*, 2014; Carrascosa *et al.*, 2013; Farahat and Bailey, 2012; Ghose and Yang, 2010; Ghosh *et al.*, 2015; Gill *et al.*, 2013; Krishnan and Sitaraman, 2013; Lacerda *et al.*, 2006; Liu *et al.*, 2014). The profit of most social media applications and websites stem from the advertisement sector. On the earlier days of the Web, advertisements were shown to users in *banner* form (Ducoffe, 1996; Sterne, 1997). Also, ad placements on websites, that is, the choice of which banner to show on a Web page was mostly static. Over time, this situation changed with the use of smarter ad placements that can either rely on the content of the ad and page (Lacerda *et al.*, 2006), or also on contextual information about the user visiting a Web page (Gill *et al.*, 2013; Ghosh *et al.*, 2015). Video-advertisements are a novel step on this progression from banner ads. Video content is a richer, information-wise, medium and currently responsible for a large fraction of the traffic online[1]. It is natural that we currently witness a rise in video-advertisements. However, regardless of media and form, advertisements online are inherently tied to online auctions and ad networks, two topics that we now discuss.

Online auctions are a common practice on the Web to select which advertisement will be shown based on a user request. Examples of requests are search engine queries, accesses to YouTube videos or simply logging in on Facebook. At the time of the request, information about the user and the request is shared with automatic bidding bots. These bots partake in an auction in attempt to get their ad to be shown to the user. Thus, online marketers configure their bots based on aspects such as target demographic (e.g., ads targeted to certain ages and/or gender) or keywords (e.g., targeting users that query for certain products or brands) (Ghosh *et al.*, 2015; Gill *et al.*, 2013). Auctions are performed based on different bidding strategies. Chapter 9 of the Easley and Kleinberg book gives an overview on the subject (Easley and Kleinberg, 2010). Another concept on today's Web is that of ad networks. These networks are composed of different partners (e.g., companies or web sites) that connect to provide resources that are used by marketers to serve ads (i.e., partake in auctions). Next, we discuss recent efforts on auctions and ad networks that are more related to our work.

Due to the popularity of advertisements on the Web, the number of ad networks that serve ads to several websites is increasing. Since most of these services are implemented as auctions, recent research has been done in an effort to better understand them (Gill *et al.*, 2013; Liu *et al.*, 2014; Bachrach *et al.*, 2016; Zhang *et al.*, 2014). In special, Liu *et al.* (Liu *et al.*, 2014) performed a study of the online social network ad market, using data from Facebook. The authors explored suggested bid data for users from different locations and interests and they found huge differences in prices paid. Among their results, they concluded that the ad market on social networks is still not mature, presenting a lot of variability. On the other hand, Bachrach *et al.* (Bachrach *et al.*, 2016) focused on two specific types of auctions: generalized second-price (GSP) and Vickrey-Clarke-Groves (VCG) and they proposed a transition model for systems to migrate from GSP to VCG auctions, with the goal of reducing the impact of the change in the revenue. Therefore, although these works are not directly related to user behavior and ad efficacy, they study auctions that play an important rule in the selection of ads that will be displayed to users.

2.2 Video-Advertising

In contrast to the large amount of research that has been done in online advertising in general, video-advertisements have only been studied very recently (Li and Lo, 2015; Dardis *et al.*, 2016; Krishnan and Sitaraman, 2013; Amarie *et al.*, 2014b; Amarie *et al.*, 2014a).

Li *et al.* (Li and Lo, 2015) performed a user study to assess the effects of video-ad length, position, and context on brand name recognition. The authors found that different ad positions have different degrees of influence on brand name recognition and also that the duration of the advertisement has a positive impact on the recognition of the brand by the users. Even though the authors studied some properties of video-ads and their relation to the success in achieving brand recognition, the experiments were limited to a small number of ads and settings. More importantly, the authors goal is orthogonal to ours. That is, the authors focused on brand name recognition, we here focus on user behavior.

Still considering user studies, Dardis et al. (Dardis *et al.*, 2016) conducted an experiment to understand the impact of banner ads and video ads on brand recall. They performed the study within two different game settings: games created with the purpose of advertising (called advergames) and non-branded games. Among their findings, they discovered that video-ads are better than banner ads in non-branded games and also that mid-roll video-ad position is more influential. Thus, although the work was focused on a very specific scenario of video-ads in games, it provides motivation to study video-ads considering a broader view, as we do here.

Aside from user experiments, Amarie *et al.* (Amarie *et al.*, 2014b; Amarie *et al.*, 2014a) used a sample of ads to motivate caching strategies for mobile advertisement in video form. The authors characterized the following properties of a small sample (458) of video-ads shown in mobile devices: size (in bytes), duration, category and time of day when the video-ad is streamed. This work is complementary to our present effort. While the authors did look into some properties of video-ad popularity, their study is focused on a small sample of ads shown in mobile devices only. Moreover, they did not study video-ad popularity evolution, content properties of pairings, user consumption behavior and monetization, as we do here.

Stepping away from social media applications, Krishnan *et al.* (Krishnan and Sitaraman, 2013) characterized a large sample of video-ads streamed from professional content websites (e.g., NBC, CBS, CNN, Hulu, Fox News etc.) using Akamai's content distribution network (CDN). One of the results

[1] http://tubularinsights.com/2019-internet-video-traffic

reported by the authors is that video-ads have completion rates (fraction of ads that are streamed in their full length to the users) ranging from 44%, when shown after the video-content, to 96%, when shown in the middle of the video-content. They also showed that longer video-contents have higher video-ad completion rates. However, the applications analyzed by the authors did not allow users to skip the video-ad exhibition and jump to the video-content: users had to abandon watching the video-content altogether so as to stop watching the video-ad. YouTube users, on the other hand, are allowed to skip the video-ad, jumping directly to the video-content, typically after an initial exhibition period. Thus, unlike in (Krishnan and Sitaraman, 2013), we here study user consumption behavior in a broader sense, by analyzing the fraction of time users were exposed to the video-ad before skipping it. We also tackle novel aspects of video-ad consumption which were not discussed in (Krishnan and Sitaraman, 2013), notably popularity evolution and video-content to video-ad relationships.

2.3 TV Commercials

Finally, we also point out various previous efforts that have studied TV commercials. Although ads on TV are presented in video format, they are different from online video-ads. On Youtube, a video-ad is exhibited when a user requests a content, whereas on TV, various ads are exhibited in sequence, during break periods of the TV program being broadcasted. Moreover, consumers are not able to skip TV commercials. Therefore, previous efforts have studied the effectiveness of TV commercials using a wide range of metrics, for instance, attitude of consumers towards the ad, characteristics of ads perceived as informative and level of irritation (Kim *et al.*, 2017; Rubinson, 2009; Stone *et al.*, 2000; Gelb and Pickett, 1983; Aaker and Norris, 1982; Stathopoulou *et al.*, 2017). Since TV commercials have also a social impact on the society, past efforts also studied other factors beyond the effectiveness of ads (Peruta and Powers, 2017; Pelsmacker and Van den Bergh, 1999; Wells *et al.*, 1971). For instance, Peruta et. al. aimed at understanding the representativeness of gender and race on TV commercials and the impact of these commercials on children. Among all these efforts, there is one work in particular (Stathopoulou *et al.*, 2017) that is more related to our study, since the authors explored social networks in order to measure the effectiveness of TV commercials. The authors used Twitter to study the engagement of consumers with advertisements and showed that consumers are more likely to engage with a brand through the use of hashtags when the commercial incorporating the hashtag is original and unusual.

Overall, this paper complements all of these prior studies as it focuses on novel aspects of video-ad consumption. The work here presented is an extended version of our previous study (Arantes *et al.*, 2016). Here, we present an in-depth look at video-ad consumption based on five research questions that cover: user behavior; video-ad popularity over time; content similarity; and, monetization effects. Our study also provides a discussion that can be used by content providers (e.g., YouTube), content creators, and online marketers.

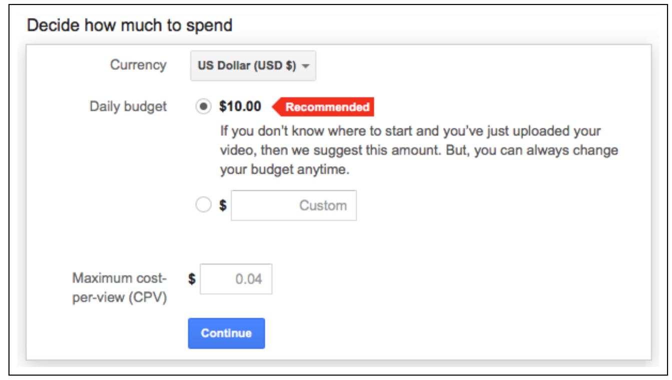

Figure 1: Defining the budget.

3 The YouTube Ecosystem

We start this section presenting an overview of the YouTube ecosystem and introducing some concepts that are used throughout this paper. YouTube is a global video-sharing website created in 2005 with the aim of allowing users to connect and communicate through videos on the Internet. Users are encouraged to watch videos, post comments, as well as publish original content. These different actions allowed for the creation of an active video based community[2]. More importantly, YouTube also allows most individuals (regular users and marketers) to upload advertisements and create advertisement campaigns. Given that most services provided by YouTube are free, the site relies on ads to generate revenue.

Several types of ads are explored by YouTube. Online marketers can choose from a set of formats and placements, ranging from banners, that are displayed to the right of the feature video, to videos that cover the entire content the user is watching. In this paper, we focus our attention to ads presented to the user in the form of a video, one of the most popular formats on YouTube. When a user requests a piece of content (a video on YouTube), an advertisement in the form of a video may be exhibited to the user. The advertisement can be displayed before, in the middle or after the streaming of the content.

The process to create an ad campaign on YouTube is straightforward. First, the advertiser needs to select the YouTube video to be used in the campaign and inform title and description of the advertisement. Next, the budget for the campaign must be defined, as presented in Figure 1. YouTube requires the advertiser to choose a daily budget and also the cost-per-view, that is, the highest price he/she is willing to pay for one exhibition of the ad. Finally, the advertiser can choose the target audience. This step is optional and YouTube allows users to be target by age, gender, interests and location, as presented in Figure 2. After the creation of the ad campaign, the advertiser has to enter account and billing information and then the ad is ready to be launched.

We use the term **video-ad** to refer to the advertisement in the form of a video and **video-content** to refer to the content requested by the user. Since a video-ad is always associated to a video-content, we call this association a **pairing**. A pairing occurs in *real time*, that is, whenever the user requests a content,

[2] https://www.youtube.com/yt/press/

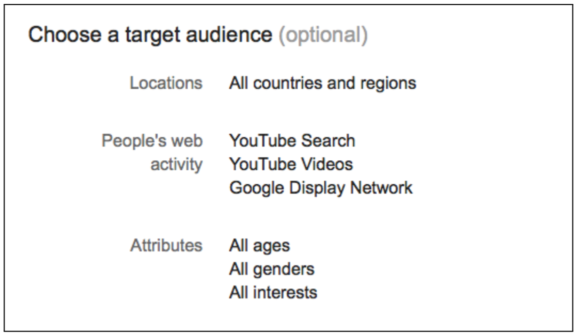

Figure 2: Targeting the audience.

one video-ad is selected to be paired with that content. Thus, the same video-content may be associated to multiple video-ads (as no video-ads at all) as response to different requests to the same content. A video-ad **exhibition**, is then, one (partial or complete) streaming of the video-ad while paired with a given video-content, and the time period during which a particular user was exposed to a video-ad exhibition is referred to as **exhibition time**. Finally, the **exposure time** of a video-ad refers to the total amount of time (all) users dedicated to streaming the given video-ad (i.e, total exhibition time).

Another important concept is that of a **monetized exhibition**. Given the exhibition time of a pairing, the video-ad may or may not be monetized. Monetization incurs in a payment from the advertiser and helps the owner of the video, a **channel**, to generate profit. Monetized exhibitions are defined by video-ads that are streamed for over 30 seconds or completely (whichever comes first)[3]. While this definition may, and has changed over time, we make use of the current policy (30 seconds) defined by YouTube. Thus our findings on monetization (see Section 8) reflect the potential profits generated by ads *if they were exhibited at the time this paper was written*. Although these policies will likely change over time, our results can be adapted to newer policies if need be. Finally, we note that the owner of the channel cannot access the revenue related to a video-ad immediately after its exhibition. Instead, he/she has to wait for a given number, usually 1,000, monetized exhibitions. The amount payed varies depending on the bids.

The selection of the best video-ad to be paired with the content is performed by YouTube. At the time the user requests the content, YouTube considers all ads that are eligible for that content (based on the target options selected by the advertisers) and chooses the best one. Selection takes into account the price the advertiser is willing to pay to exhibit the ad (called bid), and features extracted from the user (e.g., gender), video-ad and the video-content being requested. All eligible ads are competing for the same placement and YouTube runs an auction to select the winner.

Any user on YouTube can watch videos and publish content, thus any user can take the role of a viewer, a content creator or even an advertiser. Advertisers pay to run ads on the website, while content creators receive monetary shares for ads associated with their content. In this environment, content creators are motivated to publish high-quality videos in order to increase the audience and consequently, the revenue. Advertisers want to show ads that will attract the attention of users and viewers want ads that are relevant to them. Hence, these three players are important for the maintenance of the website and they can all benefit from ads.

In addition, we note that a video-ad is a video by itself on YouTube and for that reason, it may also be requested directly, without being paired with other videos. Thus, in our study, a video-ad is ultimately any video that is used as an advertisement by being paired with other video-contents in the system. In the next section, we detail our datasets.

4 Data Collection and Cleaning

In order to provide answers to our five research questions, we combined data from two rich and complementary sources. Initially, we collected HTTP requests from a university campus network to analyze user behavior when exposed to video-ads. From these requests, we filtered every video-ad to video-content pairings (both uniquely identified by system ids) that occur when video-ads are displayed in YouTube videos. This dataset was combined with the public information available from the YouTube's API[4] and statistics provided on the HTML content of the video page. Such information allowed us to analyze global properties of video-ad consumption, while still focusing on the same video-ad and video-content pairings present in our HTTP requests.

4.1 Capturing User Behavior

In order to capture user behavior in terms of how they consume video-ads on YouTube, we relied on logs of HTTP requests originating from the campus network of a major Brazilian university, with a population (including students, faculty and staff) of over 57 thousand people. Specifically, we captured the outgoing/incoming HTTP traffic from the local campus network using TSTAT (Finamore *et al.*, 2011). The tool provides us the headers, originating IP addresses, and timestamps of each request/response pair. Our goal was then to extract from these requests each video-ad to video-content pairing, as well as the exhibition time of the video-ad in each such pairing. This was a challenging task, as, in the absence of prior studies of video-ad requests to YouTube, we did not know how to identify neither the pairings nor the exhibition times in the traffic log.

Thus, we started by first manually identifying different request patterns for video-ads. We did so by browsing different YouTube videos and using network analysis tools provided by modern browsers (e.g., Firefox and Google Chrome) to assist in our investigation. We were able to identify request patterns for video-ads exhibited on: (1) the YouTube website; (2) embedded videos on different websites[5]. These requests contain the

[3] https://creatoracademy.youtube.com/page/lesson/ad-types

[4] http://developers.google.com/youtube/

[5] We also attempted to identify video-ad requests from mobile devices. However, due to the different YouTube streaming applications (e.g., Android and IOS), as well as different mobile browser request patterns, we were unable to identify a representative set of requests to cover the various means of exhibiting YouTube video-ads

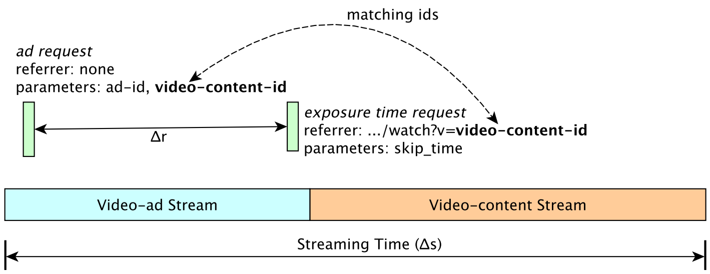

Figure 3: Matching video-ad ids to video-content ids to identify ad to content pairings.

unique YouTube identifiers of both video-ad and video-content, as exemplified below:

```
(1) ...youtube.com/api/stats/ads?
          ad_v=WVgYOaERNj4&
          content_v=-faTXv3Frc0&...
(2) ...youtube.com/yva_video?
          video_id=WVgYOaERNj4&
          content_v=-faTXv3Frc0&...
```

In requests to the YouTube's website (example (1)), the unique id of the video-ad is captured by the `ad_v` parameter. In requests for embedded video (2), it is identified by `video_id` parameter. In both cases, the video-content id is captured by the `content_v` parameter. Using only these requests, it is possible to identify all ad to content pairings that occurred inside the campus network, but *not* the video-ads' exhibition times. In order to capture this metric, we identified two other HTTP requests that are triggered when: (3) the video-ad is exhibited in full to the user; (4) the video-ad is exhibited only partially as the user skips it after a certain initial period of streaming. Examples of these two request types are shown below:

```
(3) ...doubleclick.net/pagead/conversion
          label=videoplaytime100&...
(4) ...doubleclick.net/pagead/conversion
          label=videoskipped&
          len=30&
          skip=6&...
```

In (3), the video-ad was streamed until completion (as identified by `videoplaytime100`), while in (4) the user skipped the video-ad exhibition after 6 seconds (as identified by the `skip` parameter). Notice that neither request contains any parameter that can be used to identify the ids of the video-content and the video-ad.

In order to match the *video-ad requests* (1-2) to the *exhibition time requests* (3-4), we made use of the HTTP referrer field, which captures the URL from which the user originated the HTTP request. All exhibition time requests have the page of a YouTube video-content as referrer, regardless of whether the request was triggered from YouTube's website or from an embedded video[6]. Making use of the referrer field, we were able to match the video-ad requests to the exhibition time requests using the following simple heuristic, which is illustrated in Figure 3.

Let us define $|\Delta_r|$ as the shortest *absolute*[7] time interval between a video-ad request and an exhibition time request that meets the following criteria: (a) both requests originated from the same IP address; (b) the video-content id on the referrer of the exhibition time request matches the `content_v` parameter on the video-ad request. Also, let us define Δ_s as the time the user spends streaming both the video-ad and the video-content. We consider that a successful match occurs between a video-ad and an exhibition time request that meet the above criteria whenever $|\Delta_r| < \Delta_s$. Otherwise, we discard the request as an unsuccessful match.

The heuristic would be sufficient if *network address translation (NAT)* was *not* present in the campus network, which we cannot guarantee. Due to NAT, multiple exposure time requests from the same IP may have the same video-ad request as a candidate match (i.e., with the shortest $|\Delta_r|$). We call this case a conflict. To deal with these conflicting matches, we initially consider as successful the match with the shortest $|\Delta_r|$ out of all matches in conflict. We then remove the matched video-ad and exhibition time requests from the HTTP trace, updating $|\Delta_r|$ for all other conflicts[8]. This is done by considering the next video-ad request with the shortest $|\Delta_r|$ as a match for the remaining conflicted exposure time requests. The process is repeated for every conflict.

$|\Delta_r|$ can be computed directly from the timestamps of the HTTP requests, as shown in Figure 3. Δ_s was approximated by the sum of: (1) the video-content duration (obtained from the API, as discussed below) and (2) the value of the `skip` parameter of the exhibition time request (for partial exhibitions of the video-ad) or the video-ad duration (for full exhibitions). Video-content and video-ad durations were obtained from the API (next section). Whenever the video-content or video-ad was not available in the API, we used the average value of the respective duration.

It is important to point out that, while the use of the total duration of the video-content will fail to capture the behavior of users that abandon watching the content, our goal with this heuristic is to *simply* match video-content to video-ad pairs and not to capture the amount of time the video-content was streamed. One issue that may rise with the use of the total duration is a *false positive* on our matching heuristic. However, such cases are similar to the above described conflicts, where we may falsely match a video-content to a video-ad. Nevertheless, this situation is also dealt with our conflict resolution strategy, given that we keep the match closest to when the video-content began streaming.

In our study we analyze the behavior of users from an aggregated level. That is, due to privacy ethics and NAT, the IP addresses (which are anonymized in our dataset) are used in our matching heuristic, they are not used in any of our analyses. Moreover, because of the possible presence of NAT, we only analyze user behavior in terms of individual video-ad exhibitions. One limitation of our dataset is that we do not have demographical data of every member of the academic population, and thus we are unable to study targeted ads to individual

on mobile devices. We leave this task for future work.

[6] In the cases of embedded videos, it would be expected that the referrer field in the requests in examples (3) and (4) would be equal to the URL that embedded the video. However, we found that the referrer is always a YouTube video page given that the video-player is actually hosted on `youtube.com`.

[7] We use absolute values of Δ_r as there is no guarantee that the video-ad request will precede the exhibition time request.

[8] In practice, the HTTP trace is not altered, the whole process is done in linear time by keeping track of conflicts in dictionaries.

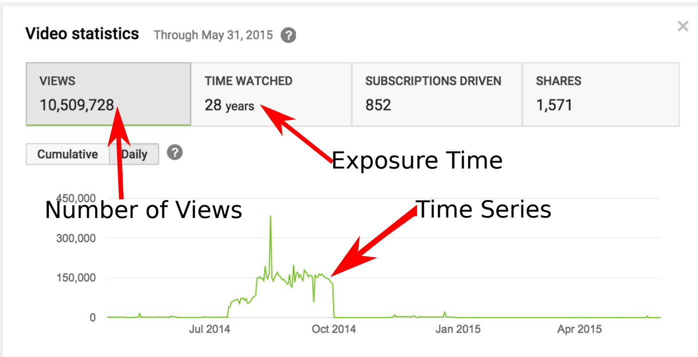

Figure 4: Public statistics data provided by YouTube.

users. However, our goal with this study is to uncover properties on the skipping behavior of users, popularity properties of video-ads, study contextual advertisements and monetization. We leave the task of analyzing personalized ads as future work. Nevertheless, we can state that based on the public campus census, the university is attended by students from all over the country, most of them are in the 20-24 age range and there is a roughly equal number of men and women.

It is also important to mention the influence of ad-blockers in our dataset. Ad-block is a type of software installed as an extension of the browser and it is used to block advertisements exhibited online. It is raising in popularity, previous efforts estimate that around 20% of users have this extension installed (Pujol *et al.*, 2015; Malloy *et al.*, 2016). Since the software works by preventing the browser from requesting URLs of advertisements, we are not able to see the blocked requests in our logs of HTTP requests, therefore we are unable to estimate the use of ad-blockers on campus. Nevertheless, we were still able detect 99,658 video-ad exhibitions in our local dataset.

4.2 Capturing Global Properties of Ads

We crawled the public API[9] information provided by YouTube for each unique id of video-content and video-ad present in our HTTP request dataset. Specifically, for each video-content or video-ad, we collected the following metadata: *upload time*, *duration* (in seconds), *title*, *description*, *category*, and list of *topics*. In addition, for video-contents only, we also collected the *channel*. Title and description are provided by the video uploader as a means to describe its content to the general audience. Moreover, every video is associated with a *category*, chosen by the uploader from a pre-defined set of options, including: *Autos & Vehicles*, *Pets & Animals*, *Entertainment*, *Howto & Style*, *Sports*, *Gaming*, *Education*, *Comedy*, etc. Every video is also associated (by YouTube) to one or more *topics*, extracted from Freebase[10], a collaborative semantic knowledge database that covers over 30 million topics, ranging from sports (e.g., baseball) to individuals (e.g., Muhammad Ali). Finally, every video-content uploaded on YouTube is automatically associated with a channel, which is the home page for a user account.

[9] http://developers.google.com/youtube/
[10] http://www.freebase.com

Table 1: Summary of our datasets.

	Campus Network	API	HTML Stats
# of unique video-contents	58,082	47,007	-
# of unique video-ads	5,667	5,052	3,871
# video-ad exhibitions	99,658	-	-

For each video-content/video-ad, we also crawled the public statistic data Figueiredo *et al.*, 2014 that is provided on the HTML page identified by the video id. This data includes aggregated values of the number of views and exposure time that are accounted for by YouTube. For video-ads only, we also collected the daily time series of both popularity measures. This statistic data is illustrated in Figure 4.

We note that, since each video-ad is an independent video on the system, these global statistics of video-ad popularity include all accesses to the video, regardless of whether it was paired with a video-content (used as a video-ad) or requested directly. We shall further discuss these effects in Section 6.

4.3 Overview of our Datasets

We ran the TSTAT tool to collect HTTP requests in the campus network from March 24^{th} to November 30^{th}, 2014. Our collected dataset includes 114,709 exhibition time requests, out of which 99,658 (86%) were successfully matched to video-ad requests, following the heuristic presented in Section 4.1. Out of those matches, 2,112 (2%) were conflicts, which were solved as described in the same section. In total, we identified 58,082 unique ids of video-contents with which some video-ad was paired. Such video-ads were identified by 5,667 unique ids. Table 1 (2^{nd} column) summarizes our dataset collected in the campus network[11].

We collected the API and HTML stats datasets on a single day, May 27^{th}, 2015. A summary of both datasets is shown in Table 1 (3^{rd} and 4^{th} columns). We were able to crawl the metadata associated with 47,007 video-contents and 5,052 video-ads, and we successfully retrieved the popularity time series of 3,871 unique video-ads. We were unable to crawl data for all video-contents and video-ads mostly because of either prohibitive privacy settings by the uploaders or video deletions. We note that, although our API and HTML stats datasets were collected after the campus collection was terminated, we can still study the global popularity of video-ads on YouTube during the same period covered by the campus dataset by trimming the time series data accordingly.

Before proceeding, we briefly discuss a few properties of the video-ads in our datasets. First, we analyze the distribution of their *lifetimes* in the system. The lifetime of a video-ad is defined as the number of days since its upload until our collection of global properties. Figure 5(a) shows the complementary cumulative distribution function (CCDF) of the lifetimes for all video-ads in our API dataset (90% of all identified video-ads). Note that all video-ads have been in the system for at least 6

[11] Our dataset is provided in https://github.com/marianavsarantes/video-ads-dataset.

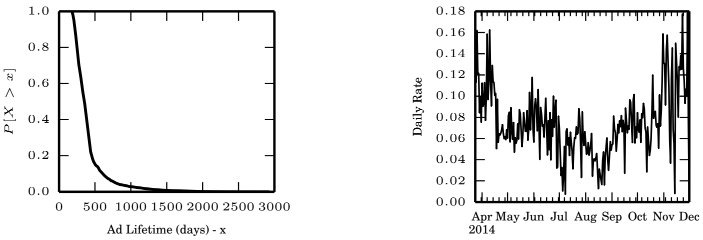

(a) Video-ad lifetimes

(b) Daily fraction of pairings

Figure 5: Overview of video-ads in our datasets.

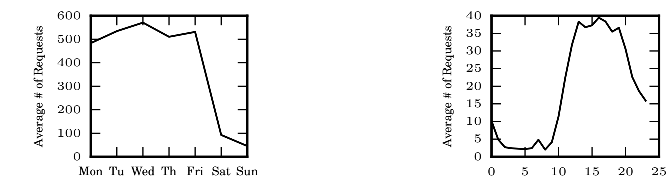

(a) Requests by day of week

(b) Requests by hour of day

Figure 6: Average number of total exhibitions per day of the week and hour of the day.

months, while around half of them have been for more than 1 year. Only a small fraction (6%) of the video-ads have lifetimes greater than 2 years, though.

Next, we look into the frequency of video-ad to video-content pairings in our campus network dataset. On Figure 5(b) we show the daily fraction of all video-content requests that are paired with any particular video-ad. We initially point out that, on average, the fraction of video-ad pairings is around 7.6%. Yet, this fraction increased significantly during the Easter period (April) and as we approached the holidays of the end of the year (starting from mid October), reaching values from 16% to 18%. Thus, in such periods, there is an increase in the expected publicity by a factor of more than 2, when compared to the overall period.

Finally, we also looked into the weekly and daily patterns of video-ad exhibitions on our campus dataset. On Figure 6(a) we show the average number of video-ad exhibitions by days of the week and Figure 6(b) presents the average number of exhibitions by hours of the day. From these figures, we can see that the requests are highly concentrated during work hours (begins rising at 9h and decreasing at 20h) and during work days (Monday to Friday). In this sense, our campus dataset cannot capture user during different periods of their daily routine (e.g., watching movies at nights or early day shopping). While this limits some of the findings that we can achieve with this dataset, as we show in the next sections, our campus traffic can be used to understand overall skipping behavior and video pairings. Moreover, our work also explores aggregated global user behavior with time series extracted from YouTube. Because of such reasons, understanding individual users on their daily and weekly routine is out of our scope.

In the next five sections we present our main findings. The specific dataset used to support each analysis can be inferred

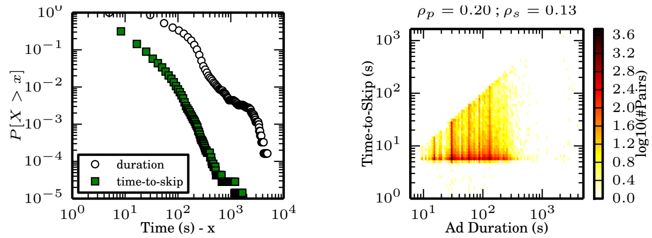

(a) Duration and time-to-skip

(b) Duration vs. time-to-skip

Figure 7: User behavior when exposed to video-ads: duration and time until user skips exhibition (time-to-skip).

based on the information exploited by it, namely, video-ad exhibitions and pairings (campus dataset), video-ad metadata (API) or video-ad popularity time series (HTML stats).

5 User Skipping Behavior

We start our study by tackling *RQ1: How do users consume video-ads?* Recall that YouTube allows users to skip a video-ad exhibition after a minimum streaming time (usually 5 seconds). Thus, we answer RQ1 by focusing on the user *skipping* behavior, as a step to analyze video-ad exhibition times.

As a basis for comparison, we first analyze video-ad durations. Figure 7(a), which presents the CCDF of video-ad durations, shows that they vary greatly across all video-ads. The mean is 107 seconds, but the median is only 60 seconds and the standard deviation is 197 seconds. Moreover, 14% of the video-ads are very short (below 30 seconds), while 35% have durations between 30 and 60 seconds, and 31% have durations above 2 minutes. We also note some rare cases of very long video-ads (over 1.5 hours) in our dataset[12].

Next, we analyze the video-ad exhibition times. The exhibition time is shorter than the duration whenever the user chooses to interrupt and *skip* video-ad exhibition. Thus, we also refer to the video-ad exhibition time as *time-to-skip*. We first note that 29,442 of the video-ad exhibitions were streamed in full. That is, in 29% of the video-ad exhibitions, users chose *not* to skip it (despite having the option to do so), watching the video-ad until completion The completion rate varies with the category of the video-ad, falling in the range of 17% (e.g., *Music*) to 49% (*Comedy*), but not exceeding 30% for most categories. We also note that the durations of the video-ads that are exhibited (at least once) in full tend to be somewhat shorter than the overall distribution, as one might expect. For example, the average duration of those video-ads is 76 seconds, and the median is only 36 seconds. Also, only 18% of them have duration above 2 minutes.

The observed fraction of video-ad exhibitions that were streamed until completion contrasts to results in (Krishnan and Sitaraman, 2013), which reports video-ad completion rates ranging from 44% to 95%. However, unlike YouTube, the applications analyzed in that work did not allow video-ad skipping. It is interesting to note that a completion rate of 29%

[12] Although rare, such ads may be exhibited to users since YouTube imposes no limit on the duration of a video-ad.

(as in our dataset) is orders of magnitude larger than the click through rates (CTR) often observed in traditional advertising (e.g., 0.01%) (Schneider *et al.*, 2009). Such higher video-ad completion rate, particularly in the presence of a skip function, might suggest a greater user engagement to this new form of online advertising. Yet, such results have to be interpreted in light of two effects. Firstly, it is impossible to skip some video-ads, a fact that increases the completion rate. Secondly, clicking on banner ads comes at a cost from the user. Streaming a video-ad is, in contrast, the *default effect*[13] provided by YouTube. There is no cost, from a user action perspective, to skip the ad. However, there is a cost related to the interest on the ad from the user. This second cost is what makes the study of the *skipping* behavior of users interesting, since it explicitly represents an action from the user of loss of interest on continuing to stream the ad. In the rest of this section, we focus on the behavior of users when they do *skip* a video-ad exhibition.

Considering only video-ad exhibitions that were skipped by the user, Figure 7(a) also shows the CCDF of the time-to-skip. Note that, in more than one third (35%) of the cases, users skip the video-ad exhibition in less than 6 seconds (one second above the minimum), whereas in only 25% of the cases users wait for more than 10 seconds before skipping the video-ad[14]. As also shown in Figure 7(a), only 1% of the video-ads have durations below 10 seconds. Thus, users often skip video-ads shortly after they are allowed to, before streaming a large fraction of their content. Indeed, we found that, on average, a user skips a video-ad after only 20% of its content has been exhibited (standard deviation of 19%). Also, in 50% of the cases, the skipping is done even earlier, after only 16% of the video-ad has been streamed.

We further analyze the skipping behavior by presenting, in Figure 7(b), a scatter plot correlating both video-ad duration and time-to-skip. Each point in the figure is a video-ad exhibition, and the colors represent the density of points. Only video-ad exhibitions that were skipped by the user before completion are included in the figure. Note that both axes are in log scale. Thus, we computed both the linear Pearson correlation (ρ_p) and the Spearman's rank correlation[15] (ρ_s) between both axes *after* taking the logarithm of all values. We found ρ_p=0.2 and ρ_s=0.13. Such low correlations are biased by the large concentration of points around a time-to-skip (y-axis) of 5 seconds. This concentration implies that many video-ad exhibitions are largely ignored by the users, who skip them as early as they are allowed to, regardless of their durations.

However, there seems to be also another (smaller) group of video-ad exhibitions that are streamed for time periods roughly proportional to their durations. To uncover this group, we focused on video-ad exhibitions that were streamed for much longer than the average, with time-to-skip above the mean (μ=12.3) plus two standard deviations (2σ=40). In those cases, which account for only 2% of all video-ad exhibitions, the correlations are indeed much higher (ρ_p=0.57 and ρ_s=0.50). Thus, those video-ad exhibition times are roughly proportional to the video-ad durations.

The results from this section may be largely impacted by users that stream a video-ad but do not necessarily watch, or pay attention to, the video-ad. That is, it is impossible to effectively say that users focused their attention to the video-ad being streamed. However, our findings on this section and the rest of the paper reflect an understanding of popularity that is based on "hits" and exhibition-times (streaming), similar to how it is accounted for at the server level (e.g., form YouTube) and exploited by video uploaders and marketeers. Thus, our findings provide a view that is perceived by analytics platforms. This factor leads the high correlations between campus views and global views that we shall study in the next section (looking into the popularity properties of video-ads).

We can summarize our main results on user skipping behavior as: (1) users often skip video-ad exhibitions as early as they are allowed to, regardless of the video-ad duration, and, on average, 20% after their beginning; (2) a small fraction of video-ad exhibitions are streamed for a time proportional to their duration; and, (3) despite this general trend towards skipping the video-ad, a considerable fraction of all video-ad exhibitions are streamed in full.

6 Video-Ad Popularity

In this section, we address *RQ2: How does video-ad popularity evolve over time?* We first analyze the overall distribution of video-ad popularity (Section 5.1). We then use the daily time series of global popularity of video-ads to analyze the dispersion of popularity temporal evolution and the amount of time until video-ads reach their daily popularity peaks (Section 5.2). Finally, we use a time series clustering algorithm to better understand the different profiles of video-ad popularity evolution (Section 5.3).

6.1 Video-Ad Popularity Distribution

We analyze the distribution of video-ad popularity using two previously used ad-efficacy metrics, namely, number of views and exposure time. The former counts the total number of times the video-ad was exhibited to a user, regardless of the time of each such exhibition, while the latter captures the total time during which users were exposed to the video-ad (i.e., total exhibition time). Our datasets provide two complementary views of each popularity measure: (1) a *local* view from the perspective of the campus network, provided by our traffic logs; (2) a *global* view from the perspective of the whole population of YouTube viewers, which is provided by the API and HTML stats pages (see Section 4). Recall that our API and HTML stat pages represent the popularity evolution of video-ads from the moment the videos were uploaded until the time we crawled YouTube (May 2015). In order to perform a fair comparison of local and global popularity of video-ads, we filtered our (global) time series data to consider only the popularity gain over the same period covered by our campus dataset (March to November 2014). We refer to this popularity view as *global filtered*.

[13] http://en.wikipedia.org/wiki/Default_effect_(psychology)

[14] The fractions are similar for all categories of video-ads.

[15] A non-parametric measure of statistical dependence between two variables that does not require linear relationships between them.

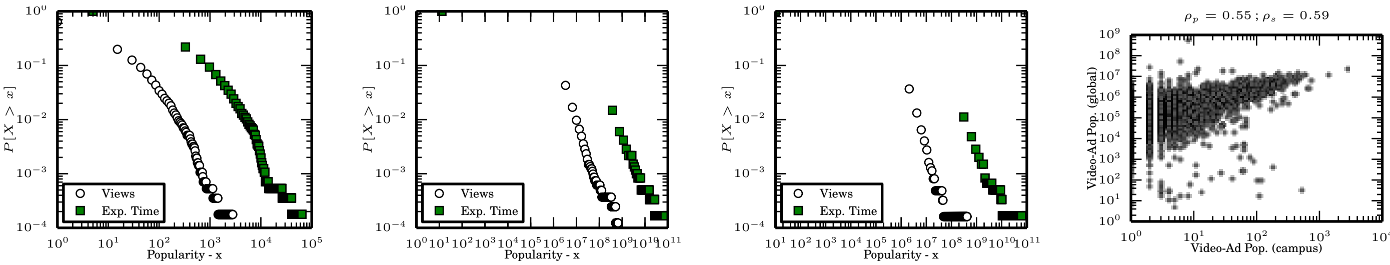

(a) Local (campus network) (b) Global (YouTube) (c) Global filtered (same period of local dataset) (d) Local vs. global filtered

Figure 8: Video-ad popularity distributions (in exposure time and number of views) according to different perspectives.

Before proceeding, we emphasize that since each video-ad is itself an independent video on YouTube, the global popularity of a video-ad accounts for *all* views of the video, regardless of whether it was paired with a video-content (exhibited as a video-ad) or accessed as an independent video. Thus, even though such global measures of popularity do not necessarily reflect, exactly, the reach of a video while promoted as a video-ad, they do capture the global interest in its content, and thus may be interpreted as the *potential efficacy* of advertisement campaigns that use the video as video-ad.

Figures 8(a-c) show the CCDFs of the two video-ad popularity measures, namely exposure time (in seconds) and number of views, for our three popularity views. As expected, the popularity measures are much higher when analyzed globally. Yet, regardless of the perspective and popularity measure, the distributions are highly skewed in nature, following a heavy tail, which is consistent with other studies of video popularity in general (Figueiredo *et al.*, 2014; Cha *et al.*, 2009). Most video-ads are exhibited only a handful of times and for very short periods, whereas a small fraction of them become very popular. For instance, only 3% of the video-ads were displayed more than 100 times on campus, while only 1.7% of them had a total (local) exposure time above 1 hour (Figure 8(a)). We also found that the most popular video-ad in our campus dataset were also very popular (within the top 0.5%) in the global and global filtered views. This particular video-ad achieved 2,812 views and was streamed for 18 hours on campus. In comparison, it received 17,859,680 views and was streamed for 389,653 hours globally during the same time period. During its whole lifetime in the system, the video-ad received 17,947,622 hits (392,239 hours of streaming).

We correlated our local popularity measures with the global filtered measures (both in log scale) to gain insights whether our local dataset reflects (to some extent) YouTube's global population in terms of video-ad popularity. This correlation is shown Figure 8(d) for popularity estimated by number of views (note the log scale on both axes). Results for exposure time are similar (omitted). We found a Spearman's rank correlation ρ_s of 0.59 (0.54 for exposure time). Such *moderate-to-strong* correlation suggests that, to a reasonable extent, our campus trace reflects the global properties of video-ad popularity on YouTube. This is an interesting result given that YouTube currently receives millions of daily viewers, whereas our local trace was collected from a campus network whose population includes only tens of thousands of users, most of whom are *not* likely to access YouTube every single day.

So far we have analyzed only the total popularity achieved by each video-ad. We are yet to discuss how this popularity evolved over time. Take the video-ad shown in Figure 4 as an example. Although it is one of the most popular video-ad in our datasets, most of its popularity is concentrated in a few weeks (based on the time series shown in the figure). Understanding how video-ad popularity evolves over time can benefit both content producers, which share a profit of the video-ad's campaign when ads are paired with their content, and content providers. For example, knowing whether the popularity of a video-ad will be concentrated on a few days or remain popular and generate revenue for longer time periods can ultimately be used to drive monetization strategies as well as caching applications (Amarie *et al.*, 2014b; Amarie *et al.*, 2014a). Thus, in the next two sections, we turn our attention to how the popularity of video-ads evolves over time.

6.2 Popularity Dispersion

To study the temporal evolution of video-ad popularity, we used the daily time series of exposure time and number of views crawled from YouTube (global view of popularity). We did not explore our campus dataset as it provides only a limited view on popularity evolution. That is, we found that no video-ad was exhibited on more than 10 days on campus. Moreover, by using the time series extracted from YouTube, we are able to analyze popularity evolution from the upload of the video-ad until the crawling time. Specifically, we address the following questions in this section: (1) How bursty is video-ad popularity evolution? (2) How much time does it take for a video-ad to reach its daily peak of popularity? We focus our discussion only on popularity in terms of number of views because very similar results were obtained for both popularity measures. Indeed, the correlations between both time series for each individual video-ad are quite strong (Pearson correlation $\rho_p = 0.99$ and Spearman correlation $\rho_s = 0.96$, on average), indicating great similarities between them (apart from scale differences).

To answer the first question, we employed a dispersion measure of inequality called Gini score (Wasserman, 2010). The Gini score can be used to measure how bursty a given time series is. Its value ranges from 0, when the total popularity acquired by a video-ad is roughly homogeneously dispersed over its lifetime, to 1, when the popularity is concentrated on a single day. According to Figure 9(a), which shows the CCDF

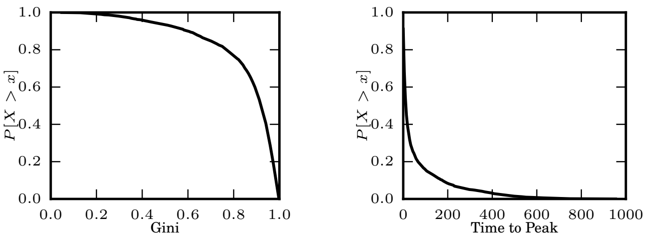

(a) Gini coefficient (dispersion)

(b) Time to peak

Figure 9: Video-ad popularity temporal evolution.

of the Gini scores computed for the video-ads in our dataset, 84% of the time series have a score higher than 0.7, and 57% have a score higher than 0.9. Thus, most video-ads have their popularity evolution concentrated on a few days. Yet, we do observe some video-ads with low Gini scores: 4% of all video-ads have scores below 0.4, suggesting that they succeeded in attracting attention for longer time periods. One might wonder whether there is a correlation between the video-ad lifetime and its Gini score (e.g., whether video-ads that have been more recently uploaded have lower Gini scores). However, we found no clear trend between video-ad lifetime and Gini score (Pearson ρ_p = -0.27 and Spearman ρ_s = -0.2).

To tackle the second question, Figure 9(b) shows the distribution of time (in days) from the video-ad upload until its daily popularity peak[16]. Typically, most video-ads (69%) reach their popularity peak within one month after upload, while for half of them the peak occurs in at most 12 days after upload. Thus, video-ads often peak in popularity very early in their lifetimes, possibly as a reflection of advertisement campaigns that are initiated shortly after the upload. However, this is not always the case. For example, for 10% of the video-ads, the popularity peak occurred only after 6 months since upload[17]. On average, the number of days until popularity peak is 56. If we normalize the time-to-peak by the video-ad lifetime, we observe that, on average, a video-ad takes only 12% of its lifetime to peak (median of 4%).

Next, we deepen our investigation of video-ad popularity by identifying common profiles (trends) of popularity temporal evolution.

6.3 Profiles of Popularity Evolution

Towards identifying profiles of popularity temporal evolution of video-ads, we made use of a time series clustering algorithm called K-Spectral Clustering (KSC) (Yang and Leskovec, 2011), which has been successfully used to study the patterns of popularity dynamics of social media content (Figueiredo *et al.*, 2014; Yang and Leskovec, 2011). KSC is a K-Means based algorithm that groups different time series into clusters based simply on the *shape* of the curves. It does so by using a distance (or similarity) metric that respects scale and time shifting invariants. That is, two video-ads that have their popularity dynamics evolving according to similar processes will be assigned to the same cluster by KSC, regardless of the popularity values. For example, two time series that are stable over time except for a peak in a day will be grouped together, regardless of when the peak occurred (time shifting invariant) and the peak value (scale invariant). By taking into account both of these invariants, we can focus on the overall *shapes*, or *trends*, that define the governing properties of popularity temporal evolution of video-ads. These trends are represented by the cluster centroids (or averaged time series) produced by the KSC algorithm.

KSC requires that all time series have the same length. Thus, we trimmed our video-ad popularity time series to include only the first 180 days. Recall that, as discussed in Section 4.3, all video-ads in our dataset have been in the system for at least 180 days. We note that such trimmed time series do include the daily popularity peaks for most video-ads: the peak occurs within the first 180 days after upload for 90% of the video-ads (see Section 6.2). Moreover, for the sake of a fair comparison between the identified profiles, we focused our analysis on video-ads that attracted at least 180,000 views (1,000 daily views *on average*). In total, we clustered 1,615 video-ads that meet this criterion.

The KSC algorithm also requires the choice of a number k of clusters. We employed various clustering quality measures suggested by previous authors (coefficient of variation, silhouette and clustering cost based scores) (Figueiredo *et al.*, 2014; Yang and Leskovec, 2011) to choose this value. We also performed a visual inspection of the cluster centroids and individual cluster members for different values of k. In a few cases, we manually merged clusters that, despite being identified as separate groups according to some clustering quality measure, did contain members with very similar popularity evolution patterns.

Based on all these heuristics, we identified $k = 6$ clusters. The cluster centroids are shown in Figure 10. Each centroid corresponds to an "average" popularity curve for the video-ads in the cluster, capturing, in general terms, the popularity dynamics of the individual members of the respective cluster. Scales on both axes are omitted to emphasize the scale and time shifting invariants. Table 2 summarizes the characteristics of each cluster by presenting the number of video-ads, as well as the average values of exposure time, number of views, ratio of exposure time to number of views, Gini score and time to peak of the members of each cluster.

It is important to notice that the popularity curves of individual video-ads may not perfectly match the corresponding centroids, however our goal is to capture the most prevalent trends of popularity evolution, respecting scale and time shifting invariants. As an illustration of this point, Figure 11 shows the popularity time series of one example of video-ad in each cluster.

Cluster C1 (Figure 10(a)) consists of video-ads that have succeeded in attracting user attention over a larger number of consecutive days. Indeed, this cluster has the smallest average Gini score (Table 2). Clusters C2 and C3 (Figures 10(b-c)), in turn, exhibit complementary popularity trends: video-ads in C2 tend to have a slow growth of popularity followed by a sharp decay, while video-ads in C3 exhibit a sharp initial growth of

[16] In case of ties – multiple days with the same popularity peak – we took the first day.

[17] Those might be videos that were first uploaded to the system and only used in video-ad campaigns much later.

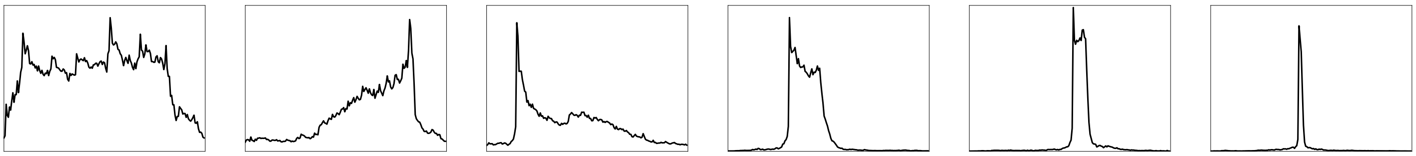

(a) C1 (b) C2 (c) C3 (d) C4 (e) C5 (f) C6

Figure 10: Trends (cluster centroids) of video-ad popularity evolution over time.

Table 2: Properties of each trend (cluster) of video-ad popularity evolution.

	C1	C2	C3	C4	C5	C6
# video-ads	69	108	109	293	467	569
Average Number of Views	1,486,175	1,869,906	4,882,094	1,789,798	1,451,894	984,175
Average Exposure Time	203,640,554	159,293,660	629,686,649	99,386,939	81,300,652	60,885,487
Average Exposure Time / Number of Views	137.02	85.19	128.98	55.53	56.0	61.86
Average Gini	0.24	0.61	0.58	0.82	0.9	0.92
Average Time to Peak	66	69	37	25	20	14

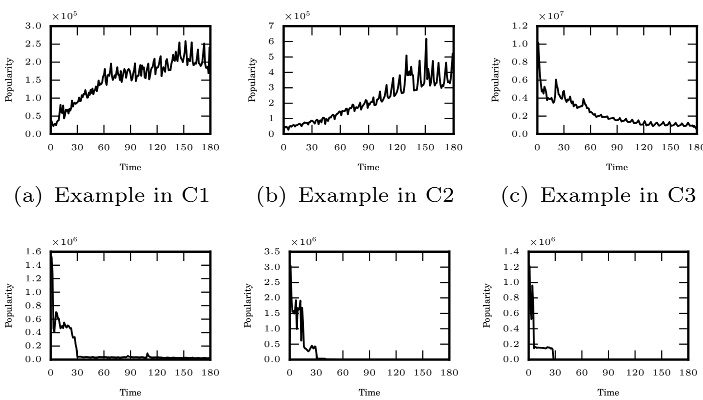

(a) Example in C1 (b) Example in C2 (c) Example in C3

(d) Example in C4 (e) Example in C5 (f) Example in C6

Figure 11: Examples of popularity time series for video-ads in each cluster.

popularity followed by a slow decay. Note that, consistently with such trends, video-ads in C2 take more than twice longer than video-ads in C3 to reach their popularity peak. These trends are interesting since similar patterns of growth and decay have previously been accounted for as viral-like propagation over social networks (Yang and Leskovec, 2011). Their average Gini scores are similar (around 0.6) but higher than that of C1. Thus, these video-ads tend to concentrate their popularity in fewer days.

The popularity trends captured by clusters C4-C6 (Figure 10) exhibit sharp increase and decrease of popularity. Also, video-ads in these clusters remain popular for much shorter time periods, compared to those in C1-C3 (note the higher Gini scores). The main distinguishing feature of C4-C6 is the time window during which the video-ad attracted user attention, which is longer in C4 and shorter in C6. These patterns may reflect advertisement campaigns having different durations. Some video-ads are publicized for a few weeks, others for only a few days. Note that the time to peak tends to decrease with the concentration of popularity, suggesting that less dispersed clusters tend to peak earlier.

Overall, video-ads in C1-C3 tend to attract more user attention in both total number of views and total exposure time, at least on average. This seems to suggest that ad-campaigns that manage to remain attractive for longer time periods will eventually become the most popular ads, which is somewhat expected. Yet, not all video-ads can remain attractive for long periods. For instance, seasonal ad campaigns, such as those related to Christmas, face the challenge of attracting a lot of attention over short time windows.

Moreover, video-ads in C1-C3 have also higher ratio of exposure time to number of views (Table 2), implying that they tend to attract more attention of *individual viewers* as well. Looking at some of the most popular video-ads in clusters C1-C2, we found two musical clips (Figures 11(a) and (b), respectively). These two videos, which were also publicized as video-ads, will likely attract viewers regardless of the ad-campaigns they are used in. Thus, one interesting direction of future work is to analyze the importance of video-ad campaigns to the ultimate growth of popularity achieved by a video.

In order to understand the nature of the video-ads in each of these clusters, we looked into their video categories (e.g., Music, Pets etc.). The clusters C1-C2 have the majority of their video-ads as members of the *Music category*. This fact can explain why the attention received by video-ads in these clusters extend for longer periods of time. Previous research has also found the effect that music videos remain attracting attention over time (Figueiredo *et al.*, 2014). The clusters C3-C6 presented *Entertainment* as the most popular category. We believe that this is a category of broad semantics (covers various topics) that is exploited by advertisers of products/goods which cannot be described by other YouTube categories. The example video-ads shown in Figure 11, for instance, promote very different products. The example video-ad for cluster C3 is about the World Cup 2014 and the video-ad in C4 is about a digital camera. The example ads for C5 and C6 are about one online service and a mobile device. Nevertheless, these ads fail to attract attention over long periods of time, and thus exhibit rise-and-fall dynamics (Figueiredo *et al.*, 2014).

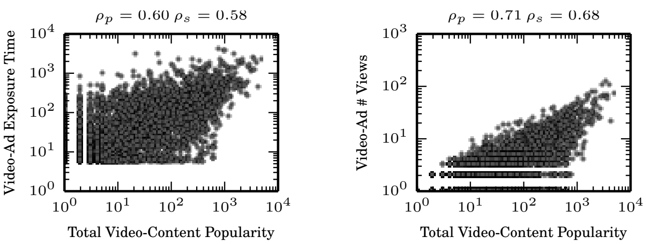

(a) Exposure Time (b) Number of views

Figure 12: Popularity of video-ad versus total popularity (in # views) of all video-contents that were paired with the video-ad (measured in the campus network).

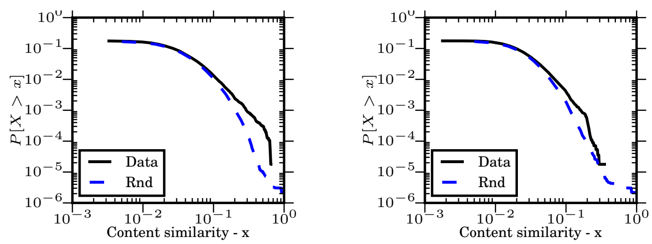

(a) TF weights (b) IDF weigths

Figure 13: Content similarity between video-ad and video-content in the real (Data) and random (Rnd) datasets.

7 Pairings

In this section, we turn to *RQ3: What are the relationships (if any) between a video-ad and the video-contents with which it is associated?* We address it by measuring the correlation between the popularity of a video-ad and the popularity of the video-contents with which it was paired (Section 7.1), and the content similarity between video-ad and video-content in each pairing (Section 7.2).

7.1 Video Popularity

We measured the correlation between video-ad popularity and video-content popularity as follows. For each video-ad in our campus dataset, we summed the total popularity, captured by the number of views, of *all* video-contents that were paired at least once with the given video-ad in the dataset. Note that this sum includes all requests to those video-contents in our dataset (even when they were not paired with the given video-ad). Figures 12(a-b) show the correlations of this value with the two previously defined video-ad popularity measures, namely, exposure time and total number of views (both axes in log scale). We focus only on popularity measures computed inside the campus, as we do not know all pairings involving a particular video-ad from the global data collected.

Figure 12 shows reasonably strong linear correlations between the popularity of the video-ad and the total popularity of all video-contents with which the ad was paired. The Pearson correlation (ρ_p) ranges from 0.6 (when correlating with exposure time) to 0.71 (when correlating with the number of views). Similarly, the Spearman's rank correlation (ρ_s) ranges from 0.58 to 0.68. Such strong correlations are intuitive: they suggest that the traffic to popular video-ads will be driven, to a large extent, by the aggregated popularity of all video-contents these ads are paired with. More popular video-contents create more opportunities for video-ads to grow in popularity. Thus, advertisement campaigns have a higher chance of being more successful when video-ads are matched to contents that are currently popular, or will grow/remain popular over time.

However, the correlations are weaker when considering individual video-contents, possibly due to the heterogeneity of the video-contents with which the same video-ad is paired. For example, when considering the average popularity of the video-contents, ρ_p ranges from 0.32 to 0.34, and ρ_s is equal to 0.36 (for both measures). Yet, the correlation between video-ad popularity and the total number of videos with which the ad was paired is quite strong (ρ_p and ρ_s exceed 0.83). Pairing a video-ad with more video-contents raises the chance of hitting a content that will be very popular, thus increasing the probability of the video-ad inheriting its audience and becoming popular as well.

These strong correlations suggest that effective content popularity prediction methods might be exploited in the design of ad-to-content pairing approaches, aiming at maximizing video-ad popularity. Indeed, content popularity prediction has recently gained a lot of attention (Radinsky *et al.*, 2013; Pinto *et al.*, 2013), often driven by the goal of designing more effective advertising services. Yet, no prior work has analyzed the correlations between video-content popularity (the target of the predictions) and video-ad popularity, thus offering quantitative results to support such goal, as we do here.

7.2 Content Similarity

We now quantify the content similarity between the video-ad and each video-content with which it was paired. To that end, we use the category, list of Freebase topics, title and description associated with each video (content and ad), crawled from the YouTube API.

Initially, we quantified the fraction of video-ad to video-content pairings in which both videos have the same YouTube category. We found that such fraction is very low (9%). Similarly, the fraction of pairings in which both videos have at least one Freebase topic in common is also very small (1%). Freebase topics are more specific than a category, and capture the semantics of a video as determined by the YouTube platform. This lack of similarity proves evidence that the videos in most pairings may have quite different semantic contents, as we further investigate next.

We then turned to the title and description features of each video to build a textual representation of the video's content. Specifically, we pre-processed the title and description features by: (1) combining the contents of both strings; (2) de-capitalizing the words; (3) removing accents, punctuation, and stop-words[18]; (4) removing words that appear only in the representation of some video-ad (but no video-content) or only in some video-content (but no video-ad).

The content of each video v was then represented as a bag of words T_v. Let $\mathcal{T}$ be the set of bags of words representing

[18] Stop-words refer to the most common words in a language. We removed stop-words in both English and Portuguese (e.g., *an*, *or*).

all videos (ads and contents) in our dataset, and $\mathcal{V}$ be the vocabulary size (i.e., total number of unique words) of $\mathcal{T}$. Each bag T_v can thus be mapped to a vector $\mathbf{t}_v$ where each entry of the vector corresponds to a word i in $\mathcal{V}$, that is: $\mathbf{t}_v =< w_{T_v,1}, w_{T_v,2}, \cdots, w_{T_v,|\mathcal{V}|} >$.

We experimented with four heuristics as weighting factors $w_{T_v,i}$. The **binary** heuristic, although simple, fails to capture the *descriptive* and *discriminative* properties of the words in T_v. We capture the descriptive strength of a word using the **Term-Frequency (TF)** heuristic. We use the **Inverse Document Frequency (IDF)** heuristic to estimate the discriminative capacity of a word and we combined both descriptive and discriminative capacities by taking the product of both metrics, using the **TF*IDF** heuristic.

Given two vectors $\mathbf{t}_a$ and $\mathbf{t}_c$ representing a video-ad and a video-content with which it was paired, we estimate the content similarity between both videos by the cosine of the corresponding vectors (using each weighting heuristic). The cosine varies from 0, when the textual representations of both videos share no common words, to 1, when they are equal.

As a baseline for comparison, we also built 500 random datasets of video-ad to video-content pairings. Each random dataset was created by taking the pairings in our real dataset and randomly shuffling the ids of the video-ad and video-content in each pair.

Figures 13(a-b) show the CCDFs of measured similarities in our real dataset (Data) and in the 500 random datasets (Rnd), for two of the four weighting heuristics. For each heuristic, we compared the two distributions by testing whether the measured similarities are greater than the similarities in the random datasets (i.e., above random chance). To that end, we applied two non-parametric statistical tests, namely one-sided Kolmogorov-Smirnov and one-sided Mann-Whitney-U. According to both tests, the similarities in our dataset are greater than the similarities in the random datasets (p-value < 0.05). Yet, as shown in Figures 13, in practice the two distributions are very similar (with differences coming up mostly in the tail). Moreover, similarity values are often very small: the median is 0 and the mean is below 0.01 in both real and random datasets, regardless of the weighting heuristic used. Similarly, the 90^{th} percentiles of the distributions do not exceed 0.04, again in both real and random datasets. These results provide evidence that most often video-ads are not paired with video-contents of similar semantic content (as captured by their title and description).

While similarities tend to be low, there is still a chance that pairings with higher similarities tend to lead to more popular video-ads. That is, users may show more interest in video-ads that are paired with similar video-contents. We investigated whether this is true in our dataset by measuring the correlation between the popularity of a video-ad and the average cosine similarity of all the pairings involving the video-ad (both in log scale). We found that these correlations are reasonably low. That is Spearman and Pearson coefficients of at most 0.33 for all four heuristics. This result indicates that popularity is not explained by similarity, as was the case when correlating video-ad and video-content popularity.

In sum, our results indicate that video-ad to video-content pairings are, in most cases, dissimilar in terms of textual content. We also found only weak evidence that more similar pairings tend to lead to more popular video-ads. One question that arises then is whether one can design novel targeted advertising techniques that, by taking the similarity between video-ads and video-contents into account when pairing them, lead to more successful (popular) video-ads. This is a subject for future work.

8 Video-Ad Monetization

We now discuss our forth research question: *RQ4: Which factors lead to video-ads being monetized?* Recall that, in order to provide a good value to advertisers, YouTube does not charge for every exhibition of video-ad on the website. When a video-ad is displayed, the reaction of the user to the advertisement (e.g., an exhibition time over 30 seconds) is taken into account to decide if the exhibition will be charged.

We start our analysis by looking at the number of exhibitions that generated revenue. These *monetized exhibitions* (see Section 3) are defined by video-ads streamed over 30 seconds. Out of the 99,658 video-ad exhibitions in our local dataset, 34,093 were monetized (34%). As we have discussed in Section 5, users will likely skip ads as soon as possible, leading to fewer monetized ads as we see here. Even though our campus trace may not necessarily reflect the global fraction of monetized views, with over 2.7 million video content streams per minute[19], we hypothesize that YouTube as a whole should have around to 1 billion monetized exhibitions daily (using our 34% estimate on the 2.7 million streams per minute). With each exhibition monetizing a few cents (Gill *et al.*, 2013), this estimate matches others that stated YouTube may generate billions of dollars yearly, translated to tens of millions of dollars daily[20]

We now look into these numbers daily. Figure 14(a) shows the volume of daily video-ad exhibitions and the volume of them that were successful and Figure 14(b) shows the fraction of exhibitions that were monetized per day. The number of video-ad exhibitions per day vary greatly, with an average of 395 and a standard deviation of 354. Looking at the fraction of monetized exhibitions, first, we can notice an increase in September, reaching a day where 83% of the video-ad exhibitions generated revenue. In this particular day, there were 208 video-ad exhibitions of 42 unique video-ads and the successful exhibitions came from only 15 of these ads. Out of curiosity, we look into these ads and they were all ads from popular brands. Next, we also notice a day, in October, with a very low rate of monetized exhibitions (5%). This day, as we can see in Figure 14(a), was also a day with only a few number of exhibitions.

So far, we have only focused on the video-ad exhibitions, without paying attention to any video-ad and video-content individually. From the 5,667 video-ads in our local dataset, 65% of them generated revenue at least once. Based on this number, we can conclude that a considerable number of video-ads were profitable to YouTube and to content providers. Nevertheless, as we present in Figure 15, the number of exhibitions

[19] http://www.visualcapitalist.com/what-happens-internet-minute-2016/

[20] http://www.forbes.com/sites/timworstall/2013/12/12/googles-youtube-ad-revenues-may-hit-5-6-billion-in-2013

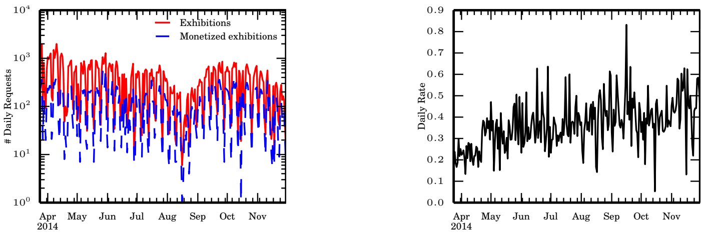

(a) Daily requests (b) Daily fraction

Figure 14: Overview of video-ad exhibitions in our dataset: volume of exhibitions and rate of exhibitions that generated revenue.

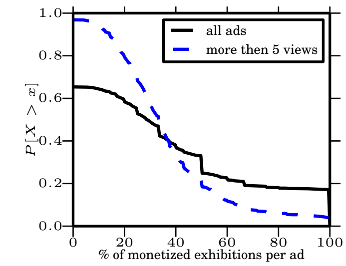

Figure 15: Distributions of the fraction of monetized exhibitions from the video-ad perspective.

that were profitable per ad is low. The figure shows the distribution of the fraction of exhibitions of each video-ad that was monetized. For comparison, we also show the distribution when considering only video-ads with more than five views (35% of all video-ads). This comparison is useful to filter out the vast majority of ads paired only a couple of times with less chance to generate revenue. We initially look at the overall ads. From the figure we can see that around 20% of ads will be monetized in over 60% of their exhibitions. This serve to show that some ads will be interesting enough to generate revenue in most of their exhibitions. This result motivate studies that aim to understanding the effect of ad quality on the generated revenue. Previous efforts looked into the effect of brands on ad interest, however social network, ad placement, ad length, and content factors (the ad itself) may also play a role (Li and Lo, 2015).

Looking at the ads with over 5 views, we can see a change in behavior. Here, most ads (over 95%) will be monetized at least once. This shows that repetition will increase the chance of monetization, as expected. Also, 18% of the video-ads had more than half of their exhibitions monetized and only 6% of them had more than 80% of their exhibitions monetized. Nevertheless, at the tail of the distribution (after 40% of exhibitions being monetized), the ads with over 5 views have a lower fraction of those views being monetized. This last effect will likely stem from those ads that are paired only a couple of times and are always monetized. Again, various factors will play a role in monetization, and repetition of ads will not necessarily lead to more monetization overall, we further explore these factors.

In order to uncover properties of video-ads that may be related to the success in attracting the attention of users, we compare two groups: (1) video-ads that generated revenue at least once, (2) and video-ads that were never successful. With these two groups we can focus on what factors lead to *at least* one monetization. We start by comparing the average duration of the advertisements. The average duration of the video-ads in the first group is 96 seconds, and 95% confidence interval ranges from 90 to 102.94, while the average duration of the video-ads in the second group is 137 seconds, with 95% confidence interval from 128.05 to 146.50. These results suggest the duration of the two groups is significantly different, indicating that shorter video-ads will have a higher chance of being monetized. In Section 5 we discussed that users will likely skip ads as soon as possible. Based on the results from that section and the results we here present, our findings show some evidence that shorter ads may attract more interest from users. Nevertheless, empirically checking these evidences require further studies.

We also aim to understand whether the categories of the video-ads have some impact on their success in generating revenue. In order to uncover the effect of categories, we conducted a chi-square test for independence for two categorical variables of our population: the categories of the video-ads and their success (the two groups defined above). Our null hypothesis states that the two categorical variables are independent. In our results, the null hypothesis was rejected, for $p = 0.05$. This result indicated evidence of dependence between categories of the video-ads and monetization. Nevertheless, we looked into categories individually to gain further insights. Recall that the video-ads in our dataset are from 15 different categories. The percentage of successful video-ads per category ranges from 42% (*Music*), to 72% (*Entertainment*). We can thus conclude that, whereas some categories appear to have a higher concentration, based on the result of chi-square test we cannot state that this effect is explained by the category itself.

Our results so far looked into the monetization of video-ads. Initially, we gave some insights on the monetization of YouTube as whole based on our campus estimate. Next, we found that a considerable number of video-ads generated revenue and the contribution of each one in particular was small, suggesting that the diversity of video-ads in the website is important. We also found that successful video-ads tend to be short and that the category of the video-ads are related to their chance of generating revenue. In the next section, we change our focus to understand monetization based on channels and video-contents.

9 Channel Perspective

YouTube allows any user to create content and earn monetary shares for video-ad exhibitions associated with their videos. Motivated by this fact, in this section, we tackle our final research question, performing a study of revenue from the perspective of the channel. That is, we turn our attention to the content creators that have their videos associated to video-ads. Whenever a user uploads a video on YouTube, the video is automatically associated with a channel. A channel is the home page for the user account and it is the place where viewers can see all the videos published by a specific user. Therefore, all videos published by the same content creator will belong to the same channel.

We start by quantifying the number of channels on our local dataset and the number of these channels that received some revenue from successful video-ad exhibitions. For each

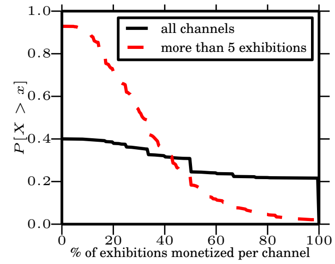

Figure 16: Distributions of the fraction of monetized exhibitions from the channel perspective

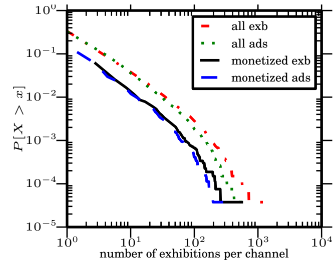

Figure 17: Distributions of exhibitions per channel

video-ad exhibition, we identified the channel of the video-content using our API dataset, explained in Section 4. Since we were not able to collect the public API information for all video-contents, we were able to find the channel information of 83% of our video-ad exhibitions and these exhibitions were related to 26,613 unique channels. Considering all video-ad exhibitions associated to each one of these channels, we found that 40% of the channels had at least one monetized exhibition. Thus, almost half of the content creators that associated their content to video-ads were able to profit from YouTube.

Similar to our analysis on video-ads, we begin by aggregating the monetized exhibitions by channel. Figure 16 presents the CCDF for the percentage of all video-ad exhibitions that were successful per channel. For comparison and to filter out tail effects (channels with few exhibitions), we again show these percentages for channels with at least 5 exhibitions. From the figure, we can see that 60% of channels are never monetized. However, 20% of channels have over 90% of their exhibitions monetized. As with Figure 15, these effects may stem from the channels with only one or two pairings. Thus, we consider channels with more than five video-ad exhibitions, 18% of them have more than half of the exhibitions monetized and only 4% have more than 80% of the exhibitions monetized. In this sense, we have evidence that monetizing most exhibitions for a single channel is rare, though it may be accomplished by a select few.

Next, we look at the popularity of exhibitions per channel. Figure 17 shows the distributions of the number of video-ad exhibitions and monetized video-ad exhibitions per channel. The distributions are very skewed, most channels are associated to video-ads just a few times, while a small fraction of channels are very popular. For instance, 7% of the channels have more than 5 video-ad exhibitions and only 0.2% of them have more than 100. The average number of exhibitions per channel is 3.11 and the standard deviation is 16. We found just one channel with more than 1,000 video-ad exhibitions. This particular channel is one of the most famous comedy channels in Brazil and it was associated, in our dataset, to 685 unique video-ads and 254 unique video-contents and, these contents were from just two categories: Comedy and Entertainment. Thus, this channel is very popular and it was able to explore a large number of distinct video-ads. When considering just the monetized exhibitions, the numbers are even lower. Just 2.5% of the channels have more than 5 monetized video-ad exhibitions and only 0.2% more than 50, hence, just a few channels were able to generate a lot of revenue. Figure 17 also shows the distributions of the number of unique video-ads and the number of unique video-ads that were monetized per channel. The similarity between the CCDF of the number of video-ad exhibitions and the number of unique video-ads present in these exhibitions shows that, in general, the same video-ad is not displayed a lot of times in the same channel. The distributions of monetized exhibitions and monetized ads are also very similar, showing that the monetized exhibitions per channel are not concentrated in just a few video-ads.

In summary, a considerable number of channels were able to profit from YouTube, although most of them were associated to video-ads just a few times, whereas only a small fraction of channels were very popular and generated revenue from several video-ad exhibitions. We also found that in general, the same video-ad is not displayed a lot of times at the same channel, the video-ads exhibited in each channel are very diverse. In the next section we conclude the paper providing discussions on the implications of our findings.

10 Discussion and Future Work

Social media applications rely heavily on their audience to generate revenue. Content providers (i.e., the application) should aim at offering an enjoyable experience to their audience, while still relying on content producers to attract users, and on online advertisers to build ad campaigns upon which all parties can profit. For example, on YouTube, advertisers usually pay the application for every 1,000 video-ad streams, while content producers receive profit for every 1,000 views of video-ads that were paired with their content. Understanding the factors behind the success of an ad campaign in a complex system is quite challenging, but it is key to the design of more effective and profitable advertising strategies.

In this paper, we took a step towards building such understanding by shedding light into how one particular type of online ad, video-ads, are currently consumed on YouTube. Driven by five research questions, we presented a thorough measurement study covering different aspects of video-ad consumption on YouTube. We now discuss some implications of our major findings.

RQ1. How users consume video-ads? Our study revealed that, even though YouTube users often skip video-ad exhibitions as early as possible, the fraction of exhibitions that are streamed until completion is reasonably high (29%). If compared to the click through rates of traditional advertising (often below 0.01%), this result might suggest a greater user engagement and thus a potentially more effective means of online advertising. Yet, this result should be taken with caution. It is important to consider that watching a video-ad in full is the *default effect* provided by YouTube. The default effect in traditional click advertising is *not to click* on the ad, which may have a role on its lower efficacy. While our work offers a first analysis of user engagement to YouTube video-ads, follow-up studies, possibly including experiments with volunteers, should be performed to compare the effectiveness of both strategies in light of default effects. Such user experiments, along with the results we present here, would provide a broader view of the user behavior, which, in turn, could offer valuable insights into the design of advertising strategies that entertain the users, while still generating profits to the other parties.

RQ2. How does video-ad popularity evolve over time? We also found that, although most video-ads have their popularity concentrated on a few days, some of them remain popular for much longer. Indeed, our study uncovered six different profiles of video-ad popularity evolution. In light of such profiles, one question that arises is: *What is the most effective means to pair video-ads and video-contents so as to increase the chance of the video-ad remaining popular for longer periods?* Content producers would be interested in attracting video-ads that remain popular for as long as possible (e.g, video-ads in clusters C1-C3) to maximize revenues. Advertisers, in turn, are interested in pairing their video-ads with contents that will lead users to their products. As we have shown, there is a trend towards video-ads that are paired with popular contents (and a larger number of video-contents) inheriting such viewers and becoming popular as well. Yet, our study also revealed that video-ad to video-content pairings are still mostly dissimilar in terms of content similarity (as captured by the textual features of both videos). This result motivates future investigations on whether contextual advertising strategies can be more effective in generating revenue for both parties.

RQ3. What are the relationships (if any) between a video-ad and the video-contents with which it is associated? On our third question, our results uncovered in this paper have focused on user behavior, popularity properties and contextual advertising. One important factor that we are currently exploring as future work is on the nature of targeted (personalized to the users demographic) ads. Targeted ads account for a large fraction of online advertising nowadays, and is the focus of studies of different ad-auction strategies. Nevertheless, the results we uncovered in this study can also be exploited by different ad-auction strategies (Gill *et al.*, 2013; Liu *et al.*, 2014). For instance, the correlations between video-content popularity and ad-popularity can be used to estimate the exhibition time of ads. Premium video-content which attracts more exhibition to ads can exploit higher prices in ad auction bids. While in contrast, the lack of correlation between the similarity of video-content and ads with popularity, indicates that this factor will likely not lead to more viewers, and thus should not affect bidding prices.

RQ4. Which factors lead to video-ads being monetized? On our forth research question we investigated the actual monetization of video-ads. Initially, we discussed the fraction of monetized exhibitions on our campus data. While this fraction may not reflect YouTube's global behavior, with it we present an educated estimate on the monetization of YouTube as whole. The lack of access to large datasets of user behavior in advertisement platforms is an issue for Web researchers nowadays. Our results here, coupled with those on RQ2, show that local campus traces may mitigate this issue. More importantly, we also discuss that shorter ads may have a higher chance of attracting user attention. Finally we also show that a small fraction of video-ads are able to monetize most of their exhibitions. Our results in this question can be explored by marketers to create more interesting ads to users. Shorter ads and some categories appear to be able to capture more attention. As stated, more entertaining advertisements to end viewers is a goal that may benefit not only advertisers, but content producers and viewers themselves.

RQ5. How successful are channels in attracting revenue? On our last question we looked at content producers. These producers gain earnings from advertisements placed with their videos. As stated, YouTube will usually pay channels after every 1,000 monetized exhibitions. Even though our campus datasets has limited information on channels, our results are able to show that some channels will monetize most of their exhibitions. Results like this one can be exploited by YouTube itself to find new partners (*The YouTube Social Network* 2012). YouTube's partners program is a worldwide initiative that aims at finding high quality channels to produce, and in consequence, monetize entertaining content for end users. Several techniques can be employed to find partners from manual inspection to machine learning algorithms (*The YouTube Social Network* 2012). Channels that are able to monetize most of their pairings can also be interpreted as a sign of possible partners. Content producers can also learn from such channels to improve their own monetization strategies, thus increasing revenues. Understanding monetization on a global level, as well as better pairing algorithms, are both interesting paths for future work. In addition, our analysis of the success of channels in attracting revenue is focused on the role of content creators, underestimating the impact of the media marketers. Therefore, another interesting path for future work is to study the role of the media marketers in the success of ad-campaigns.

With this work, we presented an understanding of video-ad consumption on YouTube. With users being more exposed to video-ads and having algorithms mediate their consumption every day, our work provides insights on how to make social media applications more beneficial to various parties. Overall, studying other applications, as well as using the knowledge we here provided to improve user experience, are both important efforts that are left for future work.

Acknowledgments

This work was partially supported by the project FAPEMIG-PRONEX-MASWeb, process number APQ-01400-14, and by individual grants from CNPq, CAPES, and Fapemig. We also thank Prof. Nazareno Andrade and the reviewers for the feedback provided on preliminary versions of this work.

References

Aaker, D. A. and D. Norris. 1982. "Characteristics of TV commercials perceived as informative." *Journal of Advertising Research*.

Abraham, I., S. Athey, M. Babaioff, and M. Grubb. 2013. "Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information". In: *Proc. EC*.

Amarie, M. C., I. Pefkianakis, and H. Lundgren. 2014a. "Mobile Video Ad Caching on Smartphones". In: *Proc. Ubicomp*.

Amarie, M. C., I. Pefkianakis, and H. Lundgren. 2014b. "Towards Understanding Mobile Video Ads". In: *Proc. HotMobile*.

Arantes, M., F. Figueiredo, and J. M. Almeida. 2016. "Understanding video-ad consumption on YouTube: a measurement study on user behavior, popularity, and content properties". In: *Proc. WebSci*.

Bachrach, Y., S. Ceppi, I. A. Kash, P. Key, and M. R. Khani. 2016. "Mechanism Design for Mixed Bidders". In: *Proc. WWW*.

Carrascosa, J. M., J. Mikians, R. Cuevas, V. Erramilli, and N. Laoutaris. 2014. "Understanding Interest-Based Behavioural Targeted Advertising". *Arxiv preprint arXiv:1411-5281*. arXiv: 1411.5281.

Carrascosa, J. M., R. González, R. Cuevas, and A. Azcorra. 2013. "Are Trending Topics Useful for Marketing?" In: *Proc. COSN*.

Cha, M., H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon. 2009. "Analyzing the Video Popularity Characteristics of Large-Scale User Generated Content Systems". *IEEE/ACM Transactions on Networking*. 17(5): 1357–1370.

Dardis, F. E., M. Schmierbach, B. Sherrick, F. Waddell, J. Aviles, S. Kumble, and E. Bailey. 2016. "Adver-Where? Comparing the Effectiveness of Banner Ads and Video Ads in Online Video Games". *Journal of Interactive Advertising*: 1–14.

Drèze, X. and F. Zufryden. 2004. "Measurement of Online Visibility and Its Impact on Internet Traffic". *Journal of Interactive Marketing*. 18(1): 20–37.

Ducoffe, R. H. 1996. "Advertising value and advertising on the web". *Journal of advertising research*. 36(5): 21–21.

Easley, D. and J. Kleinberg. 2010. *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge University Press.

Farahat, A. and M. C. Bailey. 2012. "How Effective is Targeted Advertising?" In: *Proc. WWW*.

Figueiredo, F., J. M. Almeida, M. A. Gonçalves, and F. Benevenuto. 2014. "On the Dynamics of Social Media Popularity". *ACM Transactions on Internet Technology*. 14(4): 1–23.

Finamore, A., M. Mellia, M. Meo, M. M. Munafo, and D. Rossi. 2011. "Experiences of Internet Traffic Monitoring with Tstat". *Network, IEEE*. 25(3): 8–14.

Gelb, B. D. and C. M. Pickett. 1983. "Attitude-toward-the-ad: Links to humor and to advertising effectiveness". *Journal of Advertising*. 12(2): 34–42.

Ghose, A. and S. Yang. 2010. "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets". *INFORMS: Management Science*.

Ghosh, A., M. Mahdian, R. P. McAfee, and S. Vassilvitskii. 2015. "To Match or Not to Match: Economics of Cookie Matching in Online Advertising". *ACM Transactions on Economics and Computation*. 3(2): 1–18.

Gill, P., V. Erramilli, A. Chaintreau, B. Krishnamurthy, K. Papagiannaki, and P. Rodriguez. 2013. "Follow the Money: Understanding Economics of Online Aggregation and Advertising". In: *Proc. IMC*.

Kim, E., S. Ratneshwar, and E. Thorson. 2017. "Why Narrative Ads Work: An Integrated Process Explanation". *Journal of Advertising*: 1–14.

Krishnan, S. S. and R. K. Sitaraman. 2013. "Understanding the Effectiveness of Video Ads: A Measurement Study". In: *Proc. IMC*.

Lacerda, A., M. Cristo, M. A. Gonçalves, W. Fan, N. Ziviani, and B. Ribeiro-Neto. 2006. "Learning to Advertise". In: *Proc. SIGIR*.

Li, H. and H.-Y. Lo. 2015. "Do You Recognize Its Brand? The Effectiveness of Online In-Stream Video Advertisements". *Journal of Advertising*. 44(3): 208–218.

Liu, Y., C. Kliman-Silver, R. Bell, B. Krishnamurthy, and A. Mislove. 2014. "Measurement and Analysis of OSN Ad Auctions". In: *Proc. COSN*.

Malloy, M., M. McNamara, A. Cahn, and P. Barford. 2016. "Ad Blockers: Global Prevalence and Impact". In: *Proc. IMC*. ACM. 119–125.

Manchanda, P., J.-P. Dubé, K. Y. Goh, and P. K. Chintagunta. 2006. "The Effect of Banner Advertising on Internet Purchasing". *Journal of Marketing Research*. 43(1): 98–108.

Pelsmacker, P. D. and J. Van den Bergh. 1999. "Advertising content and irritation: A study of 226 TV commercials". *Journal of international consumer marketing*. 10(4): 5–27.

Peruta, A. and J. Powers. 2017. "Look Who's Talking To Our Kids: Representations of Race and Gender In TV Commercials On Nickelodeon". *International Journal of Communication*. 11: 16.

Pinto, H., J. Almeida, and M. Gonçalves. 2013. "Using Early View Patterns to Predict the Popularity of YouTube Videos". In: *Proc. WSDM*.

Pujol, E., O. Hohlfeld, and A. Feldmann. 2015. "Annoyed users: Ads and ad-block usage in the wild". In: *Proc. IMC*. ACM. 93–106.

Radinsky, K., K. Svore, S. Dumais, J. Teevan, A. Bocharov, and E. Horvitz. 2013. "Behavioral Dynamics on the Web: Learning, Modeling, and Prediction". *ACM Transactions on Information Systems*. 32(3): 1–37.

Ribeiro, B. 2014. "Modeling and Predicting the Growth and Death of Membership-based Websites". In: *Proc. WWW*.

Rubinson, J. 2009. "Empirical evidence of TV advertising effectiveness". *Journal of Advertising Research*. 49(2): 220–226.

Schneider, F., A. Feldmann, B. Krishnamurthy, and W. Willinger. 2009. "Understanding Online Social Network Usage from a Network Perspective". In: *Proc. IMC*.

Stathopoulou, A., L. Borel, G. Christodoulides, and D. West. 2017. "Consumer branded# hashtag engagement: can creativity in TV advertising influence hashtag engagement?" *Psychology & Marketing*. 34(4): 448–462.

Sterne, J. 1997. *What makes people click: Advertising on the Web*. Que Corp.

Stone, G., D. Besser, and L. E. Lewis. 2000. "Recall, liking, and creativity in TV commercials: A new approach". *Journal of Advertising Research*. 40(3): 7–18.

The YouTube Social Network. 2012. Ed. by M. Wattenhofer, R. Wattenhofer, and Z. Zhu.

Variety. 2014. "YouTube U.S. Video-Ad Revenue Will Top $1 Billion in 2014: Report".

Variety. 2015. "Facebook's Bite into TV: Social Giant to Hit $3.8 Billion in Video-Ad Sales by 2017, Analyst Says".

Wall-Street-Journal. 2014. "YouTube Alone Brings in 20% of Digital Video Ad Revenue".

Wasserman, L. 2010. *All of Statistics*. Springer Publishing Company, Incorporated.

Wells, W. D., C. Leavitt, and M. McConville. 1971. "A reaction profile for TV commercials." *Journal of Advertising Research*.

Yang, J. and J. Leskovec. 2011. "Patterns of Temporal Variation in Online Media". In: *Proc. WSDM*.

Zhang, W., S. Yuan, and J. Wang. 2014. "Optimal real-time bidding for display advertising". In: *Proc. KDD*.