UNIVERSIDADE FEDERAL DE MINAS GERAIS
Instituto de Ciências Exatas
Programa de Pós-Graduação em Ciência da Computação

Mariana Vieira Siqueira de Arantes

Analysis of Video-Advertising Viewing Patterns on YouTube

Belo Horizonte
2018

Mariana Vieira Siqueira de Arantes

Analysis of Video-Advertising Viewing Patterns on YouTube

Final Version

Thesis presented to the Graduate Program in Computer Science of the Federal University of Minas Gerais in partial fulfillment of the requirements for the degree of Master in Computer Science.

Advisor: Jussara Marques de Almeida
Co-Advisor: Flavio Diniz de Figueiredo

Belo Horizonte
2018

Mariana Vieira Siqueira de Arantes

Análise dos Padrões de Visualização de Vídeo-Propagandas no YouTube

Versão Final

Dissertação apresentada ao Programa de Pós-Graduação em Ciência da Computação da Universidade Federal de Minas Gerais, como requisito parcial à obtenção do título de Mestre em Ciência da Computação.

Orientadora: Jussara Marques de Almeida
Coorientador: Flavio Diniz de Figueiredo

Belo Horizonte
2018

Ficha catalográfica elaborada pela Biblioteca do ICEx - UFMG

Arantes, Mariana Vieira Siqueira de.

A662a Analysis of video-advertising viewing patterns on YouTube / Mariana Vieira Siqueira de Arantes – Belo Horizonte, 2018.
xviii, 78 f. il.; 29 cm.

Dissertação (mestrado) - Universidade Federal de Minas Gerais – Departamento de Ciência da Computação.

Orientadora: Jussara Marques de Almeida Gonçalves
Coorientador: Flávio Vinícius Diniz de Figueiredo

1. Computação – Teses. 2. Redes sociais on-line. 3. YouTube (Recurso eletrônico). 4. Internet na publicidade. Orientadora. II. Coorientador. III. Título

CDU 519.6*04(043)

UNIVERSIDADE FEDERAL DE MINAS GERAIS
INSTITUTO DE CIÊNCIAS EXATAS
PROGRAMA DE PÓS-GRADUAÇÃO EM CIÊNCIA DA COMPUTAÇÃO

FOLHA DE APROVAÇÃO

ANALYSIS OF VIDEO-ADVERTISING VIEWING PATTERNS ON YOUTUBE

MARIANA VIEIRA SIQUEIRA DE ARANTES

Dissertação defendida e aprovada pela banca examinadora constituída pelos Senhores:

Jussara Marques de Almeida Gonçalves

PROFA. JUSSARA MARQUES DE ALMEIDA GONÇALVES - Orientadora
Departamento de Ciência da Computação - UFMG

Flavio V. D. de Figueiredo

PROF. FLAVIO VINICIUS DINIZ DE FIGUEIREDO - Coorientador
Departamento de Ciência da Computação - UFMG

Raquel Prates

PROFA. RAQUEL OLIVEIRA PRATES
Departamento de Ciência da Computação - UFMG

Nazareno Andrade

PROF. NAZARENO FERREIRA DE ANDRADE
Departamento de Sistemas e Computação - UFCG

Belo Horizonte, 17 de outubro de 2018.

Resumo

Publicidade online se tornou muito popular na Web. Atualmente, muitos websites, e em particular plataformas de mídia social, fornecem acesso gratuito a uma grande variedade de conteúdos e serviços para atrair a atenção de usuários. Para gerar receita, essas plataformas comercializam dados de usuários com anunciantes que exploram esses dados para promoverem suas marcas e serviços. Nesse contexto, anúncios em formato de vídeo (vídeo-propagandas) estão ganhando força e se tornando uma das principais fontes de receita na Internet. A mudança de formas mais tradicionais de anúncios, como propagandas em texto e imagens, para vídeo-propagandas trouxe novos desafios para o ecossistema de publicidade online. O formato de vídeo é um conteúdo multimídia mais rico que permite aos anunciantes promoverem seus produtos e serviços de maneira mais sofisticada e dinâmica, sendo mais invasivo para os usuários. Portanto, é importante entender este novo formato de anúncio e seu impacto no comportamento dos usuários, no que tange os seus padrões de visualização quando expostos a vídeo-propagandas.

Nessa dissertação estuda-se vídeo-propagandas, usando YouTube como estudo de caso. YouTube é atualmente a plataforma mais popular de compartilhamento de vídeos e grande parte de sua receita é gerada através do uso de vídeo-propagandas. Esta dissertação é composta por duas partes complementares. Por um lado, uma visão sobre vídeo-propagandas a partir da perspectiva do usuário é fornecida. Uma abordagem qualitativa com a aplicação de questionários e diários é utilizada para investigar as ações e experiências de usuários quando expostos a vídeo-propagandas. Em seguida, um estudo sobre monetização é apresentado, provendo uma visão de vídeo-propagandas no Youtube sob as perspectivas do criador de conteúdo e do sistema. A partir da análise de uma grande base de dados de logs de requisições HTTP da rede de um campus universitário, exibições de vídeo-propagandas são exploradas, com o intuito de mensurar o potencial delas em gerar receita para criadores de conteúdo e para o YouTube. Em resumo, esse trabalho fornece uma visão atual do ecossistema de vídeo-propagandas, apresentando resultados que motivam o desenvolvimento de estratégias mais efetivas para a criação de vídeo-propagandas potencialmente mais lucrativas.

Palavras-chave: Redes Sociais On-line, YouTube, Publicidade na Internet.

Abstract

Online advertising is ubiquitous on the Web. Nowadays, several websites, and in particular social media platforms, provide free access to content and services in exchange for user attention. In order to generate revenue, these platforms trade user data and attention with advertisers that ultimately promote their brands and content to end viewers. In this context, advertisements in the form of video (video-ads) are gaining significant traction, becoming one of the leading forms of revenue on todays' Internet. This shift from traditional text and banner ads to video-ads has brought new challenges to the ad ecosystem. The video format is a richer multimedia content that allows advertisers to promote their products and services in a more sophisticated and dynamic way, being more invasive to users. Therefore, it is important to understand this novel advertisement format and its impact on how users behave, in terms of their viewing patterns, when exposed to video-ads.

In this thesis, we study video-ads taking YouTube as our case study. YouTube is the most popular video-sharing platform nowadays and it stems most of its revenue from video-ads. Our work is composed by two main parts. We start by providing a view of video-ads from the perspective of the users. We take a qualitative approach, employing survey and diary based research to investigate the user actions and experiences when exposed to video-ads. Our aim is to bring forth the role such users play in the complex ecosystem of online video-advertisements. Next, we shift our attention to monetization. Using a large dataset of logs of HTTP requests from a university campus network, we explore video-ad exhibitions to understand their potential of generating revenue to content creators and YouTube, thus providing an overview on monetization. In sum, our work provides a timely look into the ecosystem of video advertisements, drawing insights that motivate the design of more cost-effective strategies to make online video-ads potentially more profitable.

Keywords: Online Social Networks, YouTube, Online Advertising, Video Advertising.

List of Figures

List of Tables

Contents

Chapter 1

Introduction

1.1 Motivation

Online advertising is ubiquitous on the Web. Nowadays, different websites, and in particular social media platforms, provide free access to content and services in exchange for user attention. In order to generate revenue, these platforms trade user data and attention with advertisers that ultimately promote their brands and content to end viewers. In this context, online advertising has become essential for the Web. Indeed, it has been estimated that by 2021 marketing leaders will spend nearly $119 billion on online advertising, surpassing spends with ads on both broadcast and cable television [46].

Due to the importance of online advertising, several types of ads are explored. A marketer can choose among a set of ad formats and placements. In particular, advertisement using video content (video-ad) is rising as one of the leading forms of revenue on today's Internet [47]. The video format is a much richer multimedia content that allows advertisers to promote their products and services in a more sophisticated and dynamic way. This richer content may have a higher impact on users (e.g., on whether they become more engaged with the content) and therefore has become the leading form of advertising for most of the platforms online.

YouTube is a successful example of a social media platform that stems part of its revenue from video-ads. The website has over a billion users and it is responsible alone for around 20% of the overall digital video-ad revenue [48]. Despite being very popular, the platform is also an example of a new advertising market where users not only watch a wide range of content for free, but they can also contribute uploading content and even creating their own video-ads to be exhibited in the website. Whenever a user requests some content (e.g., a YouTube video), a video-ad may be exhibited to the user (usually before the requested content is streamed [1]). In certain cases, the website allows the user to skip the ad after some seconds of streaming. If users do not skip the video-ad and stream it for a certain minimum amount of seconds, the advertiser is charged for the exhibition.

[1]There are several types of video-ads on YouTube and they can be exhibited before (pre-roll ads), in the middle (mid-roll ads) or after (post-roll ads) the streaming of the content.

The revenue from streamed video-ads is shared between the application and the owner of the content the video-ad was associated with.

This new advertising market, where any user can act as a viewer, a content producer or an advertiser, and more importantly, any user can profit from ads that are associated to their content, has attracted the attention of researchers and practitioners on the Web. Understanding the perception of users about video-ads and the factors related to the success of video-ad exhibitions can be very valuable to the ad ecosystem as a whole, helping in the design of better ad campaigns that will be potentially more profitable.

There are several previous studies on traditional online advertisements [41, 25, 50, 5, 17]. However, studies focused on video-advertisements are still rare and preliminary [1, 32, 15]. Little is known about the opinion and perception of users in regard to video-ads, as well as which factors may influence users' decision to skip or watch them.

1.2 Goals

The goal of this thesis is to broaden the understanding of the video-advertisement ecosystem, taking YouTube as our case study. Most prior efforts on video-advertisements relied on large datasets and have taken a quantitative approach in order to characterize some properties of ads. However there is a lack of research that tackles the users, i.e., those who are frequently exposed to video-ads. Also, previous studies looked at some properties of video-ads, but did not address the factors that may influence their success as well as their capacity of generating revenue to YouTube and content creators. Thus, we here aim at covering the gap on previous research by looking at video-advertisements from two complementary perspectives:

A View From the User: Our aim here is to study the user actions and experiences when exposed to YouTube video-ads, bringing forth the role such users play in the complex ecosystem of online video-advertisements. In order to perform the study, we take an exploratory approach, employing survey and diary based research.

An Overview on Monetization: Here we present a study of video-ad monetization on YouTube. We first explore a dataset of logs of HTTP requests to characterize video-ads that were exhibited on a campus network. We look at the video-ad exhibitions that were successful in generating revenue. Moreover, we also analyze the popularity of content creators and their success in profiting from ads. Then, we change our attention to the content of the video-ads. Most prior efforts have studied textual metadata information associated with video-ads (e.g., title, description, duration), here we take a step ahead and look at the multimedia properties of video-ads, analyzing to which extent those properties are related to monetization.

1.3 Contributions

The contributions of this thesis are:

- An exploratory study of the viewing behavior of users when exposed to video-advertisements. We look at the actions and experiences of users when exposed to video-ads, uncovering their reasons for deciding to skip the streaming or watch them until the end. Our work can help uncovering attributes of more engaging video-advertisements. It can also be used by providers of services in order to choose the best video-ads to be displayed to each user.

- A first study of video-ad monetization on YouTube, deepening our understanding of the video-ad market by looking at the exhibitions that were streamed long enough to generate revenue. We also provide a first study of the multimedia properties of video-ads, motivating the application of image processing to this new field of online video-advertising.

Our work has yielded the following publications:

- Towards Understanding the Consumption of Video-Ads on YouTube, featured in The Journal of Web Science 2017 [2].

- An Investigation of User Actions and Experiences when Exposed to YouTube Video Ads, featured in WebMedia 2018 [3].

1.4 Outline

The remainder of this thesis is organized as follows. Chapter 2 reviews previous work on online advertisement in general and, in special, video-advertising. Chapter 3 introduces some concepts that are used throughout this thesis and explains our data collection. In Chapter 4 we present our exploratory study of video-ads from the perspective of the users. Our overview on monetization is presented in Chapter 5. Finally, conclusions and directions for future work are presented in Chapter 6.

Chapter 2

Related Work

In this chapter we present a review of the current literature related to our work. We start by discussing previous studies of online advertisements in general and then we provide a more in-depth discussion of prior investigations related to video-advertisements.

2.1 Online Advertising

Online advertising is the economic foundation for most social media applications, services and websites. In the early days of the Web, advertisements were shown to users in banner form and the ad placement was mostly static [18, 45]. With the evolution of online advertising, today there are several formats of advertisements and smarter ad placements that can rely on the content of the ad and page, and also on user data, to dynamically choose the ads to show [7, 41, 12, 10, 17, 60, 29, 40, 38, 55].

Online auctions are currently a common practice to dynamically place ads. In these auctions the ads are selected based on the user request. Examples of requests are search engine queries, accesses to YouTube videos or simply logging in on Facebook. When a user makes a request, all ads that are competing for the placement (based on aspects such as demography, keywords, etc.) partake in an auction and the winning ad is selected to be displayed. The auctions can be performed based on different bidding strategies. Chapter 9 of the Easley and Kleinberg book gives and overview on the subject [19] and several studies explore and compare different strategies [36, 25, 5, 59].

Besides research on bidding strategies and auctions in general, researchers also focused on the development of models and algorithms to measure and improve the success of ad campaigns [44, 50, 57, 12, 35]. [44] proposed a framework for estimating the economic value of keyword advertising campaigns. The framework can be used to assess the prospect of success, as well as the expected return of investment and the possible associated risks. On the other hand, [50] presented a more general model that provides a reliable measurement of the effectiveness of an advertising campaign, using conversion rate as the proxy for success. Conversion rate is a difficult metric to estimate and [57] tackle this problem. The authors proposed a model that uses customer and product conversion patterns to estimate the rate. The model was tested using a synthetic dataset generated based on real data and the results showed that it achieved a robust prediction performance.

So far we have only discussed models focused on the success of ad campaigns from the advertiser perspective. There are also studies that proposed models for predicting and measuring the quality of advertisements from the user perspective [60, 6, 39, 30, 24]. [60] designed a framework to predict low quality native ads[1] on Yahoo News stream. Yahoo provides a feedback mechanism that allows users to hide native ads from the system when they think they are offensive. The authors used the offensive feedback as a proxy for the quality of native ads. They extracted features from the title, description and image of the ads and used these features to train a model to detect offensive ads. The authors then applied the model to filter out offensive ads, improving the user experience and quality of the system as a whole. [6] also studied native ads, but with the goal of predicting the user's dwell time on the ad landing pages. Dwell time is the amount of time users spend on the landing page of ads after clicking on them. The authors used features extracted from the native ads and its landing pages to train the predictive model. They then exploited the model to promote ads with longer dwell times on the system, thus improving the experience of users when clicking on ads.

[39] proposed a predictive model to estimate the price advertisers would pay to reach a user, based on the exposed user personal data. The motivation of the authors was to foster transparency on the Web. Using logs of HTTP requests of real mobile users and through data acquired by running their own ad-campaigns and by tapping on the Real Time Bidding protocol [58], the authors were able to develop a methodology for enabling end-users to estimate in real time their actual cost to advertisers. The methodology was then used to build a browser extension that can be installed by users that wish to be aware of their value to the ad ecosystem.

Stepping away from theoretical and predictive models, several studies focused on advertising from the user perspective, but with a broader goal of understanding the behavior of users and their preferences [34, 37, 13, 8, 4]. [34] conducted a qualitative study to explore the perception of users about contextual advertising and intrusiveness in online advertising. The author used data gathered from closed-ended questionnaires, diaries and interviews with the participants. The participants were very negative about the over-use of rich media to advertise products, as well as ads that were being forced upon them. In general they had positive experiences when the advertisements were simple, clear, short and predictable in location and form. Contextual advertising was also considered a good strategy, ads targeted by content and interests were more agreeable to participants.

[1]Native ad is a type of online advertising that replicates the look and feel of its serving platform.

In the same direction, [37] relied on data gathered from a survey and interviews to investigate user attitudes towards personalized advertising. The authors were able to uncover two distinct groups of participants. In one group, the participants had a positive attitude towards personalized ads, they saw the usefulness of those ads and were not worried about their personal data and privacy. In the other group, however, participants considered personalized ads invasive and were very concerned about their privacy. The authors also observed that, regardless of the specific group, in general, ads were deemed more relevant when they were not only related to the interests of the participants, but were also tuned to nuanced preferences of style, timing and personal taste.

As some studies discussed below have shown, in general users have bad experiences when they are exposed to annoying and intrusive advertisements. The presence of bad ads disrupt their use of the service or application and has been an incentive to the development of softwares that can detect and block advertisements. These softwares are called ad blockers and they are very popular nowadays, being the focus of recent research [33, 42]. Since most of the services offered on the Web rely on advertisements to survive, the use of ad blockers has a huge impact on the revenue of service providers and publishers, costing them billions of dollars a year [53]. Publishers sometimes try to overcome ad blocking by using softwares that can detect them, the anti-adblockers [26]. Another way to circumvent ad blocking consists on agreements with ad blocker softwares to allow advertisements that have high quality and are not intrusive or annoying. The ad blockers implement those agreements through the use of whitelists, that have also been studied recently [49].

Although the studies discussed in this section are not directly related to video-ads, they are important to the ad ecosystem as a whole. They also uncover several aspects related to advertisements that can be applied in our research. In the next section, we shift our attention to previous efforts that are more related to our work.

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 [15, 32, 1, 28, 9].

We start by discussing previous efforts that applied a quantitative approach to investigate different properties of video-ads. In [28], the authors collected and analyzed a large set of traces from professional content websites (e.g., NBC, CBS, CNN, Hulu, Fox News etc.) using Akamai's content distribution network (CDN). The aim of the work was to uncover key factors that are related to the effectiveness of video-ads, measured by their completion and abandonment rates. The results showed that the duration of video-ads can affect their effectiveness, with longer ads presenting lower completion rates. Moreover, the ad position (pre, mid or post roll) also affects completion rates, with mid-roll ads being more likely to be completed than pre-roll ads.

In the same direction, [9] relied on a small sample (458) of YouTube video-ads that were streamed in mobile devices. The aim of the study was to understand the impact of size (in bytes), display time, frequency and also the category of video-ads on the ad lifetime[2] and the number of exhibitions. The authors found that short length video-ads tend to live longer and have a larger number of occurrences. These findings were then applied to the design of a video-ad caching system for smartphones. The system was able to reduce the volume of data transference by half.

Our own previous work [1] also performed a study of some properties of video-ads. Using logs of HTTP requests originated from a university campus network, we were able to identify video-ad exhibitions on YouTube. We explored the evolution of popularity of those ads and also their success in attracting user attention. Since YouTube allows users to skip video-ads after an initial period of time (in general 5 seconds), we studied the skipping behavior of users, using it as a proxy for the effectiveness of video-ads. This work is the basis for the study presented in this thesis and is further discussed in more details in Chapter 3.

Our present work is complementary to all of these prior studies. Previous research explored metadata information about video-ads and their effectiveness. One key contribution of our present work is the analysis of monetization of video-ad exhibitions. We study the potential of video-ads in generating revenue to content creators and YouTube itself. We also look beyond metadata information, extracting multimedia features of the video-ads and understanding the impact of those features on the skipping behavior of users.

Aside from quantitative studies, the use of experiments with users and qualitative analysis have also been applied in recent research [15, 32]. [32] conducted an experiment with users to understand the impact of some properties of video-ads on brand name recognition, namely ad-length, ad-position and ad-context. Among their findings, the authors discovered that ad-length has a positive impact on brand name recognition and that ad-position can also affect the effectiveness of video-ads.

[15] also conducted user experiments in order to compare the impact of video-ads and banner ads on brand name recognition and attitude towards the brand. Two experiments were conducted. In the first experiment, video and banner ads were inserted into two types of games: non-branded games and advergames (games created with the purpose of advertising) and the impact of different settings on brand name recognition was analyzed. In the second experiment, the same two types of games were used to compare

[2]The lifetime of a video-ad was defined as the number of days since its upload until the collection of the dataset.

the impact of video-ads. The authors showed that video-ads were more effective in non-branded games and also that video-ads in mid-roll position were more influential.

Another key contribution of our present effort is an exploratory analysis that complements the aforementioned studies. We look into the actions and attitudes of users towards video-ads considering a broad environment. Previous work focused on very specific settlements and used brand name recognition as a proxy for the success of advertisements. Here, we explore the skipping behavior of users on their daily use of YouTube, uncovering their reasons for skipping video-ad exhibition.

In summary, our study provides a view of the ad ecosystem from the perspective of the users, the content creators and the service provider. Although prior efforts have focused on the user perception towards advertisements, most of these studies are very broad, studying all types of advertisements. The studies that focused on video-ads, in turn, explored very specific settings and thus, have important but limited implications. In Chapter 4, we present a study from the user perspective, covering the gap on previous research by analyzing the skipping behavior of users and their perception of video-advertisements. Moreover, in Chapter 5, we use a large dataset to study video-ads, focusing on monetization, channels and multimedia features of the ads. Previous efforts also relied on large datasets, but they characterized just metadata information of the video-ads and did not studied their role in generating revenue to YouTube and content creators.

Chapter 3

Contextualization

In this chapter, we start by presenting an overview of the YouTube ecosystem and introducing some concepts that are used throughout this thesis (Section 3.1). Next, we explain our data collection (Section 3.2) and provide details of our datasets (Section 3.3).

3.1 YouTube Ecosystem

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[1]. 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 thesis, we focus our attention to ads presented to the user in the form of a video, a currently popular format 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 of creating 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 3.1. 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

[1] www.youtube.com/yt/press/

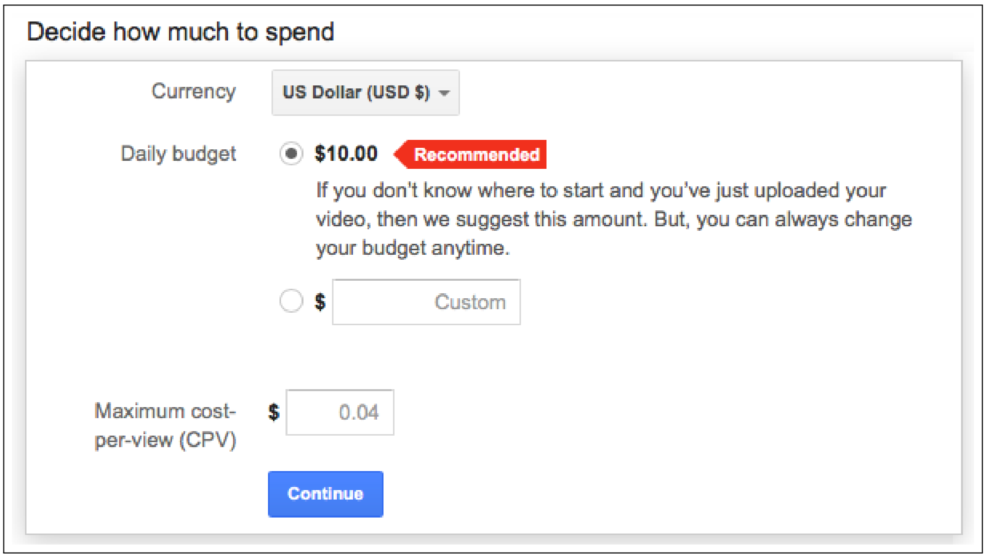

Figure 3.1: Defining the budget (screenshot taken in November 2017).

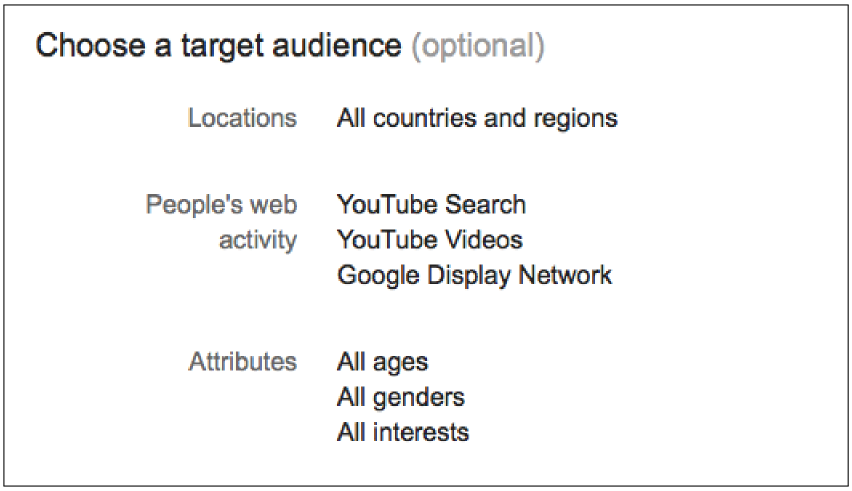

Figure 3.2: Targeting the audience (screenshot taken in November 2017).

and YouTube allows users to be target by age, gender, interests and location, as presented in Figure 3.2. 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 now introduce the notation used throughout our thesis to refer to several key concepts in the aforementioned ecosystem. 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, one video-ad may be 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. Similarly the same video-ad may be dynamically associated to different video-contents. A video-ad **exhibition** is defined as a (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).

It is also important to mention another concept we use throughout this paper, the **context** of the user. YouTube offers a huge amount of content in the website, covering a wide range of topics. For instance, we can easily find videos of music, classes, history, books, recipes, etc. Therefore, there are many different reasons that can motivate people to use the application. Sometimes users are in a moment of entertainment, other times the website is being used as a source of information for professional reasons. We use the term context to refer to the intention and actions of the user at the time when she watched a video-content (and a paired video-ad) on YouTube. In Chapter 4, we will analyze the impact of the user context on her skipping behavior.

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-content, a **channel**, to generate profit. Monetized exhibitions are defined by video-ads that are streamed for over 30 seconds or completely (whichever comes first)[2]. While this definition has changed over time (and may continue changing), we make use of the policy defined by YouTube at the time this work was developed (i.e.,the aforementioned 30 seconds or full streaming policy)[3]. Thus our findings on monetization (see Chapter 5) reflect the potential profits generated by ads *if they were exhibited at the time this thesis was developed.* Although these policies will likely change over time, our results can be adapted to newer policies if necessary. 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 of monetized exhibitions (typically 1,000). 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 video-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 video-ads on the website, while content creators receive monetary shares for video-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

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

[3] Last checking on 11/2017

revenue. Advertisers want to show ads that will attract the attention of users and viewers want video-ads that are relevant to them. Hence, these three players are important for the maintenance of the website and they can all benefit from video-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 data gathering procedure and collected datasets.

3.2 Dataset Collection

In order to study video-advertisements from the two perspectives we discussed in Chapter 1, we combined data from four different sources. First, to study video-advertisements from the perspective of users, we took an exploratory approach, gathering data through the use of a survey and a diary (dataset 1). Then, to study video-ad monetization, we initially 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 (dataset 2) 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 (dataset 3). Finally, for each unique video-ad that was displayed on campus, we also collected its audio and video, extracting multimedia features from them (dataset 4).

Datasets 2 and 3 were collected in our own previous study about video-advertisements [1]. Datasets 1 and 4 are new and their collection methodology is a contribution to this thesis. We defer the description of how we collected the new datasets to Chapters 4 and 5, focusing here on briefly describing how we collected datasets 2 and 3 and their main characteristics.

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

3.2.1 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 [21]. 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, as described in our prior work [1], 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 and (2) embedded videos on different websites[5]. These requests contain the 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 and (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
```

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

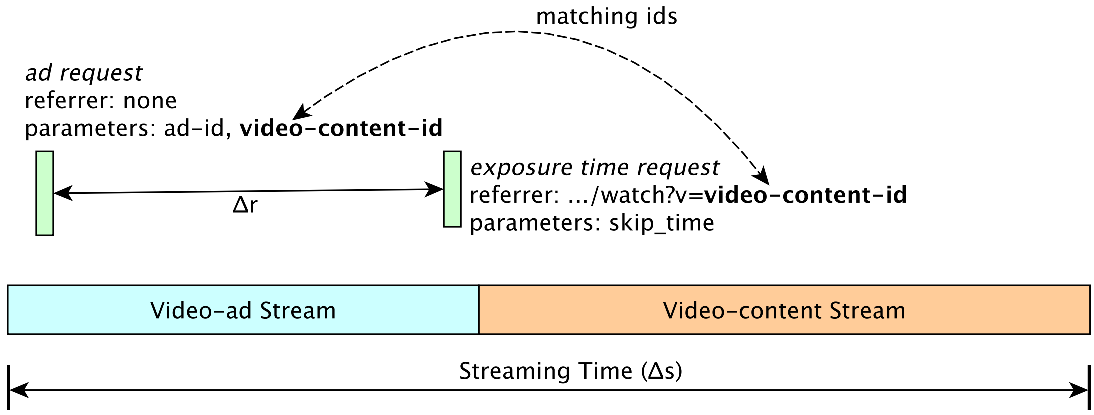

Figure 3.3: Matching video-ad ids to video-content ids to identify ad to content pairings.

```
        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 3.3.

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 and (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 considered 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 discarded the request as an unsuccessful match.

[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.

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 called this case a conflict. To deal with these conflicting matches, we initially considered as successful the match with the shortest $|\Delta_r|$ out of all matches in conflict. We then removed the matched video-ad and exhibition time requests from the HTTP trace, updating $|\Delta_r|$ for all other conflicts[8]. This was 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 was repeated for every conflict.

$|\Delta_r|$ was computed directly from the timestamps of the HTTP requests, as shown in Figure 3.3. That is, 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 (as described in the following 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 was 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 was also dealt with our conflict resolution strategy, given that we kept the match closest to when the video-content began streaming.

As in [1], we here 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 users. Thus, 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

[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 3.4: Public statistics data provided by YouTube.

advertisements exhibited online. It is raising in popularity, previous efforts estimate that around 20% of users have this extension installed [42, 33]. 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.

3.2.2 Capturing Global Properties of Ads

The second dataset used in [1] was generated by crawling 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 id*. 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

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

	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	-	-

Table 3.1: Summary of our existing datasets.

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.

For each video-content/video-ad, we also crawled the public statistic data [20] 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 number of views and exposure time. This statistic data is illustrated in Figure 3.4.

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.

3.2.3 Overview of our Existing Datasets

As presented in [1], we ran the TSTAT tool [21] 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 3.2.1. Out of those matches, 2,112 (2%) were conflicts, which were solved as described in that 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 3.1 (2^{nd} column) summarizes the dataset collected in the campus network[11].

We also collected the API and HTML stats datasets on a single day, May 27^{th}, 2015. A summary of both datasets is shown in Table 3.1 (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,

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

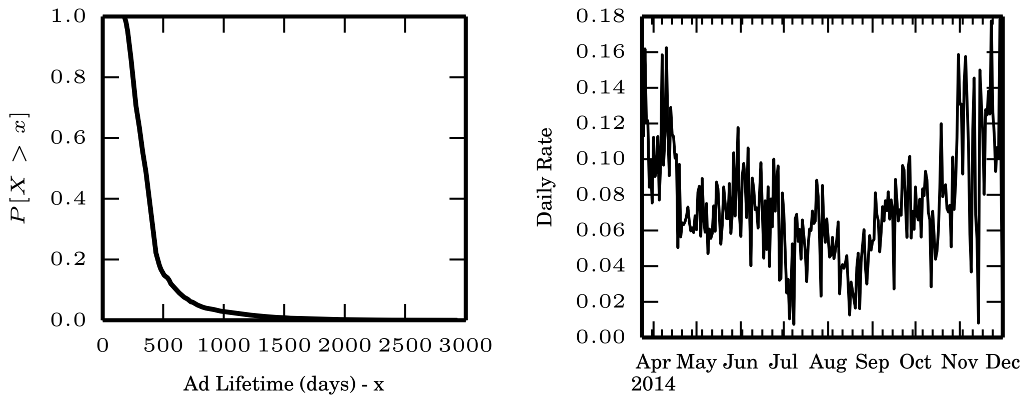

(a) *Video-ad lifetimes* (b) Daily fraction of pairings

Figure 3.5: Overview of video-ads in our datasets.

although our API and HTML stats datasets were collected after the campus collection was terminated, we were still able to 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 (see Figure 3.4).

We now summarize some few properties of the video-ads in our datasets so as to guide the reader in the analysis presented in Chapter 5. This discussion was further elaborated in [1]. We start by presenting the distribution of video-ad lifetimes, which is defined as the the number of days since the upload of the video-ad until our collection of global properties. Figure 3.5(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 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 show, in Figure 3.5(b) the daily fraction of all video-content requests that were paired with any particular video-ad. This fraction is on average only 7.6%, but it 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.

We also looked into the weekly and daily patterns of video-ad exhibitions in our campus dataset. Figures 3.6(a) and 3.6(b) show the average number of video-ad exhibitions by days of the week and hour of the day, respectively. We can see that the requests are highly concentrated during work hours (begins rising at 9AM and decreasing at 8PM) and during work days (Monday to Friday). Thus, in a sense, our campus dataset does not capture users 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, our campus traffic can be used to understand overall skipping behav-

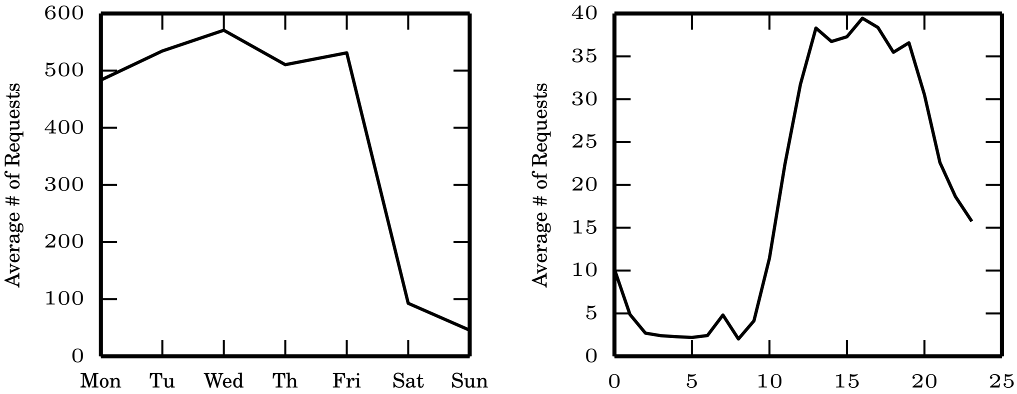

(a) *Requests by day of week* (b) Requests by hour of day

Figure 3.6: Average number of total exhibitions per day of the week and hour of the day.

ior 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 section we review the main findings of our previous analysis of our campus dataset as well as the metadata and video-ad popularity series [1]. The specific dataset used to support each analysis can be inferred based on the information exploited by it, namely, video-ad exhibitions and pairings (campus dataset), video-ad metadata (API) and video-ad popularity time series (HTML stats). In the next chapters we complement this prior work by tackling user perception of video-ad exhibitions and monetization.

3.3 Previous Results

Prior to this thesis, we explored the campus, API and HTML datasets in a preliminary work about video-advertisements [1]. We studied the popularity evolution of video-ads over time, the similarity between video-ads and video-contents that were paired on our campus dataset as well as the skipping behavior of users. In the following we briefly review the main findings of this prior study, as it has been the foundation for the work presented in this thesis.

Our study of video-ad popularity focused on two well-known metrics of ad efficacy, namely number of views and exposure time. Using the campus dataset and the global metadata, we analyzed the distributions of these two metrics, observing that they are very skewed in nature, with most video-ads being exhibited only a handful of times or with short

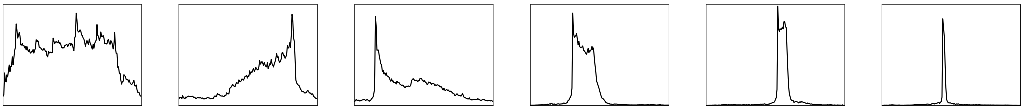

(a) *C1* (b) *C2* (c) *C3* (d) *C4* (e) *C5* (f) *C6*

Figure 3.7: Trends (cluster centroids) of video-ad popularity evolution over time [1].

exposure time. For example, only 3% of the video-ads were displayed more than 100 times on campus, while only 1.7% of them had a total exposure time on campus above 1 hour.

The distributions of popularity give us the total popularity of video-ads, but do not help to understand how the popularity has evolved over time. To that extent, we used the daily time series of number of views crawled from the HTML stats page. We started by employing a dispersion measure of inequality called Gini score [51]. The Gini score can be used to measure how bursty a given time series is. We found that most video-ads in our dataset have their popularity evolution concentrated on just a few days. Just a small fraction of ads were successful in attracting attention for longer periods. We also applied the KSC clustering algorithm [56] to find profiles of popularity evolution, for each of the two popularity metrics. We found 6 profiles, for both metrics, as illustrated in Figure 3.7. As we can see, the profiles range from ads that remained popular from long periods to ads that had just a peak in popularity and then disappeared.

In our previous work, we were also interested in uncovering the relationships (if any) between video-ads and the video-contents with which they were associated. We started by measuring the correlation between the popularity of a video-ad (using both exposure time and total number of views) and the popularity of all the video-contents with which it was paired. We found strong correlations (e.g., Pearson correlation ranged from 0.6 (when correlating with exposure time) to 0.71 (when correlating with the number of views). Similarly, the Spearman's rank correlation ranged from 0.58 to 0.68.), suggesting that video-ads that are paired with popular video-contents have higher chances of becoming popular as well. Moreover, we also studied the content similarity between the pairings. Using the title and description associated with each video (content and ad), crawled from the public API, we built a textual representation of the videos' content. We first preprocessed the title and description fields, merging them and removing stop words. Then we represented the content of each video as a bag of words. We experimented with four heuristics as weighting factors, namely, binary, term-frequency (TF), inverse document frequency (IDF) and a combination of the last two (TF*IDF). Given two vectors representing a video-ad and a video-content in a pairing, we estimated the content similarity between both videos by the cosine of the corresponding vectors. We found that video-ad to video-content pairings were, in most cases, dissimilar in terms of textual content.

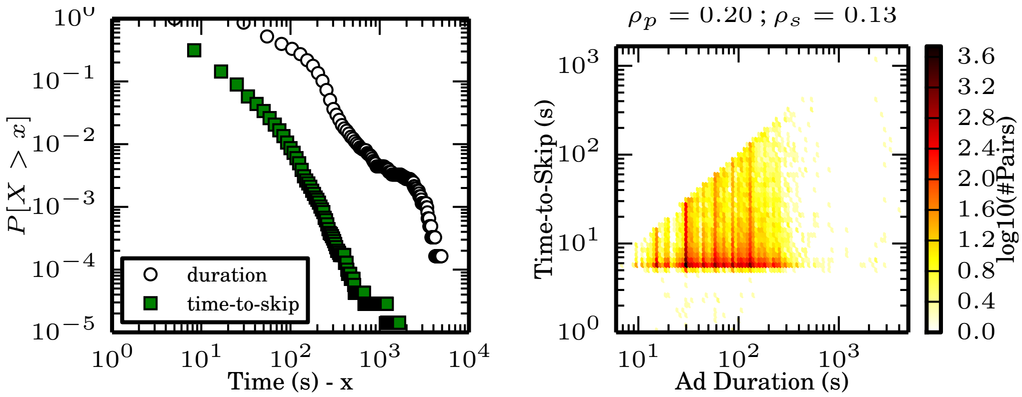

(a) *Duration and time-to-skip* (b) Duration vs. time-to-skip

Figure 3.8: User behavior when exposed to video-ads: duration and time until user skips exhibition (time-to-skip) [1].

Finally, we looked into the skipping behavior of users. Remember that, using our log of HTTP requests, we were able to identify the video-ad exhibition time for every pairing on campus. The exhibition time of a video-advertisement is shorter than the duration whenever the user decides to skip it. Therefore, we first looked at the fraction of video-ad exhibitions that were streamed in full[12]. We found this fraction to be 29%, which was surprisingly high. Next, we focused only on the exhibitions that were skipped by users. We analyzed the distribution of the time-to-skip, that is, the amount of time the video-ad was streamed before the user clicked on the skip button, as presented in Figure 3.8(a). We found that in more than one third of the cases (35%), users skipped the video-ad exhibition as soon as they were allowed (less than 6 seconds). In only 25% of the cases users waited for more than 10 seconds before skipping the video-ad. We also analyzed the correlation between time-to-skip and the duration of the video-ads. The scatterplot is presented in Figure 3.8(b) and it shows that although most exhibitions were skipped around 5 seconds, there is a small group of exhibitions that were streamed for a time proportional to their duration.

Although we shed some light on the skipping behavior of users, our previous study was merely quantitative. We were focused only on the amount of video-ad exhibitions that were streamed until completion or were skipped and the time until skip. To deepen our understanding of the skipping behavior of users, there is a need to understand how users perceive the value of video-ad exhibitions and their reasons behind the skip of video-ad exhibitions. Therefore, in Chapter 4, we present an exploratory study of skipping behavior of users, complementing the preliminary analysis summarized here.

In Chapter 5, we also complement this preliminary study, exploring different aspects of our datasets. Using the channel information crawled through the public API and

[12]We cannot guarantee that a video-ad exhibition streamed in full was actually seen by the user. The user could be doing something else or the skip button could be not available.

by further exploring the video-ad exhibition times, we provide a first study of video-ads from the perspective of the content creators, specially studying the potential of the video-ads in generating revenue. We also extend the analysis of different properties of video ads by characterizing, using our multimedia dataset, features extracted from the audio and video of the video-ads.

Chapter 4

A View from the User

In this chapter, we present a study of video-advertising from the perspective of the users. Using YouTube as our case study and employing a survey and a diary, we explore the actions and experiences of real users when exposed to individual video-ads. We start by motivating our study (Section 4.1) and presenting our methodology (Section 4.2). Then, we employ the data gathered through the survey to assess the opinions of users about video-advertisements (Section 4.3). Finally, we explore the diary data in order to shed some light into the reasons behind users skipping of video-ad exhibitions (Section 4.4) and also to study the role the multimedia content of both video-ad and video-content play on users' decisions to skip a video-ad exhibition (Section 4.5).

4.1 Motivation

Video-advertising is raising in popularity and, as discussed in Chapter 2, it has attracted the attention of researchers. However, most previous work on video-advertising has focused on algorithms and bidding strategies, looking at the ad ecosystem from the side of advertisers and service providers, such as YouTube, which rely on ads to generate revenue. There is a lack of research emphasizing the users, the ones who are constantly exposed to video-ads and advertisements in general. Thus, our aim in this chapter is to study video-advertisements providing a view from the users.

We are interested on the behavior and expressed opinions of individual users when exposed to video-ads on YouTube. Our goal in this Chapter is to bring forth the role such users play in the complex ecosystem of online video-advertisements. That is, similar to previous work, we make use of the skipping behavior as a proxy for user interest in video-ads; but unlike prior efforts we employ survey and diary [31] based research to capture the user's perspective. In particular, we focus on three research questions:

RQ1: How do users perceive the value of video-ads on YouTube?

RQ2: Why do users skip (or not) video-ad exhibitions?

RQ3: What is the role of the content of the video-ad and the video it is paired with in the user's decision to skip or not an exhibition?

Each research question captures a different perspective of YouTube video viewers. On RQ1 we focus on understanding how users perceive the value of video-ads in the system. To tackle this question we employ a survey with users. On our second and third research questions (RQ2 and RQ3) we analyze the use of the skipping feature by users to assess their perception of individual video-ads exhibited to them. We employed a structured diary consisting of a few multiple choice questions and some open questions, and asked users to add a new entry to this diary whenever a video-ad on YouTube was exhibited to them.

4.2 Methodology

In this section, we present the methodology used to gather user experience data when exposed to video-ads. Due to the exploratory nature of our study, we employed two methods of data collection: a survey and a diary [31].

Survey: The survey was comprised of a closed-ended questionnaire[1]. Our aim with this questionnaire was to collect demographic information of participants as well as their personal opinions about online advertising on YouTube. We asked participants their age and gender, the frequency at which they use YouTube, and their opinion about video-ads. The questionnaire was composed of eight questions, presented in Table 4.1. Questions S2-S5 and S8 are multiple-choice questions, while S6 and S7 are 5-point Likert scaled questions.

Diary: We used a diary to gather information about users' behavior when exposed to video-ads and the reasons behind their decisions regarding watching them or not. Unlike a lab-based experiment, where users are monitored for a short period of time (e.g., while browsing YouTube), a diary based research is conducted with no monitoring of user behavior and the data is collected by the users themselves. That is, users browse YouTube as they normally do in their own routine (e.g., in their houses and work environments), and fill out an entry in the diary every time a video-ad is displayed to them. In this sense, a diary allows us to tap into the user experience with minimum intervention and thus low impact on user experience. We developed a feedback diary[2], that is, a diary in which

[1]The complete survey in English is presented in Appendix A.
[2]The complete diary in English is presented in Appendix A.

S1	What is your name?
S2	What is your age?
S3	What is your gender?
S4	How often do you use YouTube?
S5	Have you ever subscribed to a YouTube channel?
S6	What is your opinion on the following statement: "YouTube would be better without video advertisements".
S7	What is your opinion on the following statement: "I would be willing to pay to use Youtube without advertisements".
S8	Do you use any software to block advertisements?

Table 4.1: Survey questions.

participants record events immediately or soon after they happen, based on pre-defined questions about the event [11]. The diary allowed participants to record the video-content requested by the user and the reason as to why they wanted to watch it, their behavior towards the ad (if they skipped the exhibition or watched it until the end), the reason for making that decision and whether they knew what the video-ad was about. Specifically, the diary consisted of the questions listed in Table 4.2. Questions D2, D5, D7 and D9 are multiple-choice questions and the others are open-ended questions. In question D2 the participants should inform the type of device they were using and the available answers were: “Computer”, “Smart phone”, “Tablet”, “Video game”, “Smart TV” or “Other”. In questions D5 and D7 the available answers were “Yes” or “No”. Finally, in question D9 the participants could answer “Yes”, “No” or “I don’t know”.

Participants were requested to make an entry in the diary every time they requested a YouTube video and a video-ad was exhibited to them. Since watching YouTube videos requires the participant to be online, to make it easier for them to fill out the diary, an online version of the form to be filled as an entry to the diary was developed and made available to them.

Both survey and online diary questionnaires were created using Typeform[3] and were subjected to a pilot test with six volunteers. After the pilot, minor changes in the wording of the questions were made, based on their feedback.

Recruitment of Participants. We recruited participants for our study through the Internet and also offline in Brazil and in the US. We created a web page[4] that briefly explained the study and provided instructions for volunteers to participate, and posted the invitation to participate on Facebook and Reddit (US version only). We went to classrooms in our university to recruit students, we sent email to several colleagues in our department and we also distributed the invitation to personal relations. The invitation

[3]www.typeform.com
[4]A screenshot of the web page is present in Appendix A.

D1	What is your name?
D2	Device.
D3	Describe in a few words the content (YouTube video) you were watching.
D4	Why were you watching this content?
D5	Did you skip the advertisement?
D6	Describe in a few words why you skipped or not the advertisement.
D7	Do you know what was the advertisement about?
D8	If you answered 'Yes' to the previous question, please tell us what the advertisement was about.
D9	Do you think the advertisement was related to your personal interests?

Table 4.2: Diary questions.

sent through Reddit included the links to the English versions of the survey and diary, whereas in the other places, the Portuguese version was included.

The process of recruitment occurred in 4 rounds, each one lasting around 2 weeks, spread from December 2015 to December 2016. During this period, if the person decided to participate he/she could start his/her diary. Participants were told that ideally they should try to participate for one week, but if such commitment was not possible, any feedback, even if for a single entry in the diary, would be helpful.

Before participating, we presented a term of consent to all participants explaining the research goals, data being collected and guaranteeing data confidentiality, as well as making clear that no financial compensation was being offered for their participation and that they could decide to interrupt their participation at any time without any consequences to them.

Our initial hope was that everyone who filled out the survey would also participate in the diary. However, since the diary is much more costly for participants, many of those who answered the survey decided not to participate in the diary. Yet, all the participants of the diary also filled out the survey.

Participants. In total, 117 people filled out the survey[5], 23 in English and 94 in Portuguese. Out of the 117 survey respondents, 28 also participated in the diary. Out of them, 25 filled out the Portuguese version and 3 filled out the English version. The minimum number of entries recorded by a participant was one (7 participants) and the maximum was 20 (1 participant). In total there were 135 diary entries, averaging 4.8 entries per participant.

Out of the 117 people who completed the survey, 62% were men and 38% were women, and the age of the majority of the participants ranged from 19 to 32. Most of them use YouTube at least once a day (60%), or a few times per week (32.5%) and 75%

[5]The Network Id field provided by Typeform was used to detect and filter out possible duplicates.

Content of the video (D3)	**Content of the ad (D8)**	**Reasons to watch video-content (D4)**	**Reasons to watch or skip video-ad (D6)**
Beauty	Airlines	Recommended	Already know the product
Comedy	Animals	To entertain	Don't Care About Ads
Game	Automobiles	To focus	Habit
How To	Banks	To gain information	Interesting
Lecture	Clothing and shoes	-	I was doing something else
Miscellaneous	Food and household	-	Long ad
Music	Games and toys	-	Saw the same ad before
Politics	Health and beauty	-	Short ad
Sports	Mobile operators	-	Skip not allowed
-	Movies, series and music	-	To help YouTubers
-	Online services and electronics	-	To watch content
-	Sports	-	Uninteresting

Table 4.3: Categories created through the open coding of open-ended questions in the diary.

of the participants have already subscribed to a channel. Looking at the demographics of the subgroup who also participated in the diary, we observed that the age range did not change, and 64% were men. We can notice that those who chose to participate in the diary are more frequent users than the general group: 96% of them use YouTube at least a few times per week and 82% have already subscribed to a channel.

Analysis. We applied the method of open coding [14, 31] to each open-ended question in the diary (questions D3, D4, D6 and D8 in Table 4.2). We manually assigned one code to the answer given to each question. The coding was carefully reviewed by another author. In order to validate our methodology, we also asked a volunteer to assign codes to the answers and then we measured the inter-rater agreement using the Cohen's Kappa coefficient [27]. This coefficient ranges from 0 to 1 and we achieved a value of 0.83, which indicates a strong level of agreement. During the process, we discarded 4 answers of the diary because they were not clear as to what the participant meant. In total, we created 9 categories for the content of the video (D3), 12 categories for the content of the advertisement (D8), 4 reasons for users to watch a video-content (D4) and 12 reasons for them to skip or not the advertisement (D6). The categories are presented in Table 4.3.

In the next sections, we present the results of our survey and diary, tackling the research questions we set out to investigate.

4.3 User Perception of Advertisement Value

In this section, we tackle*RQ1: How do users perceive the value of video-ads on YouTube?* In order to address this question, we used the data gathered through our survey

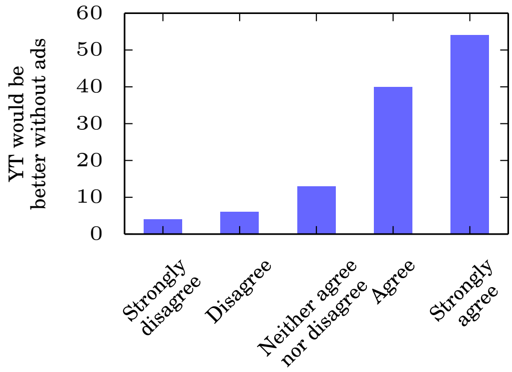

Figure 4.1: Histogram for question S6 of the survey.

and focused on three specific questions: (a) Do users believe YouTube would be better without video-ads? (b) Are users willing to pay to use YouTube without video-ads? and (c) Are users aware of software that can block advertisements? Do they use it?

Towards answering the first question, we analyzed the answers of participants to S6 of our survey, as presented in Table 4.1. In this question, participants were asked about their opinion on the statement that YouTube would be better without video-ads. The answers were scaled, ranging from “Strongly agree” to “Strongly disagree”. Figure 4.1 presents the histogram for the 117 responses. Most participants agree (or strongly agree) that the application would be better without video-ads. However, there is also a group of participants that are either indifferent to the presence of video-ads (11%) or believe their existence are important to the application (8.5%). This is an interesting result since we are usually bounded to think that users always hate advertisements, despite the quality of the ad or the moment they are exposed to it. Even in our small dataset, we were able to find users that are more open to video-advertisements[6]. This observation motivates the study of factors that lead users to like the video-ads they are exposed to as well as the development of algorithms to more effectively target users when selecting such video-ads. In fact, [34] shows in his qualitative study that, in general, personalized ads are better accepted by the users.

We continue our analysis by focusing on the responses for question S7 of our survey. In this question, we collected the opinion of the participants about paying to use YouTube without video-ads. The answers are also scaled and Figure 4.2 shows the histogram. Although the majority of the participants does not agree with the idea of paying to use the application without advertisement, 25 of our participants (21%) agree or strongly agree with that idea. This result suggests that providers should care about the needs of different types of users, offering different options to consume the service. Some popular Web applications already provide such flexibility, often offering two options: one is to use

[6] We tried to have a diverse participant pool and we recruited participants from several places. However, we still have participants that are students of computer science. We are not sure whether participants from this field of study are biased because they know to a certain extent about the importance of video-ads to the YouTube economy.

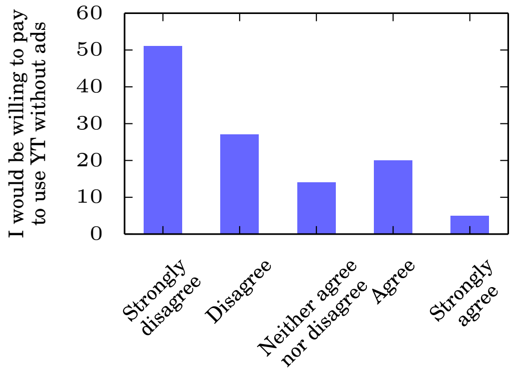

Figure 4.2: Histogram for question S7 of the survey.

the service for free with the presence of ads and the other is to pay to use it without ads. One idea is to offer more flexible options, allowing users to personalize the types of video-ads they are willing to watch according to their interests, opting in and out of specific types of advertisements, specially because there are some types of ads that are considered more intrusive than others [34].

Next, in order to deepen our understanding of different user profiles, we look into questions S6 and S7 jointly, considering the answers per participant. Figure 4.3 presents a heat map showing the number of participants who responded accordingly to each pair of responses. The map uncovers a great number of participants who thinks YouTube would be better without video-ads, but at the same time would be unwilling to pay not to watch them. In fact, 64% of the participants who answered "Agree" or "Strongly agree" to question S6 (better without ads), answered "Disagree" or "Strongly disagree" to question S7 (would be willing to pay). The contradiction between these two answers shows that video-ads are usually seen in a negative way, even though they are used by providers as a means to offer a wide range of services for free. This result raises a question of whether there are other options that would allow providers to offer free services beyond the use of advertising. Options that users may consider more enjoyable. In the very least, our observation motivates providers to try to change users' perceptions of advertising, associating to it the benefits of supporting free service.

Finally, we focus our attention on question S8 of our survey. In this question, participants were asked about their use of software to block advertisements. This type of software is used to block intrusive ads and it is usually offered for free as an extension of the browser. Three options of answer were provided (yes, no, and I don't know this software), as shown in the histogram in Figure 4.4. Surprisingly, even with the huge popularity and availability of this type of software, more than half (70) of our participants do not use them. Out of them, only 13 of the participants answered that they did not know this type of software. Thus, despite the general negative impression of video-ads, most participants do not use any software to block them, maybe because users may likely find that installing browser extensions is a bothersome task. We further correlated the

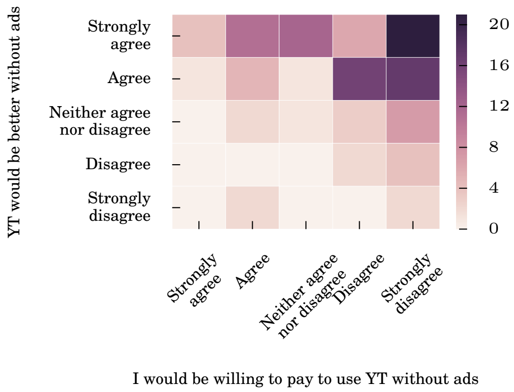

Figure 4.3: Heat map for questions S6 and S7 of the survey.

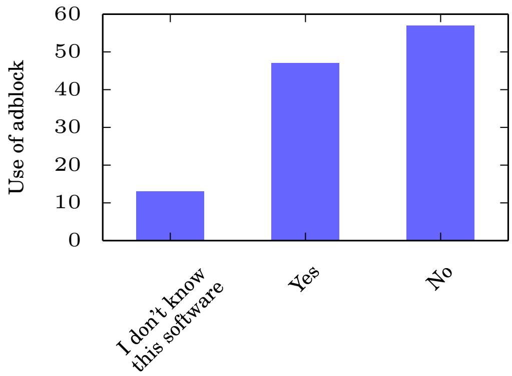

Figure 4.4: Histogram for question S8 of the survey.

answers given to questions S6 and S7 with the use of ad block but did not find any particular tendency. Thus, we have not looked further into this issue, referring to [42] and [33] for more discussions on the use of ad blocks.

In summary, users usually perceive video-ads in a negative way and would prefer to use the application without them. But we also found participants who are more open to advertisements, motivating research to improve the quality of the video-ads and their exhibitions. In the next two sections, we present our analysis of the data collected through the diary.

4.4 Reasons for Skipping Video-Ads

In this section we address *RQ2: Why do users skip (or not) video-ad exhibitions?* Previous efforts have looked into the effectiveness of video-ads [28] and the skipping behavior of users [1]. However only measurement analyses have been performed. We here complement these studies by looking into the skipping behavior from the users' perspec-

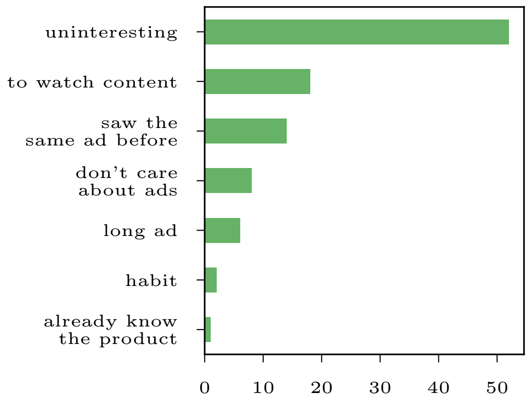

Figure 4.5: Histogram for the categories created through the open coding of question D6 (reasons to skip a video-ad).

tive, uncovering the reasons behind skips of video-ad exhibitions, and thus offering insights that can be used to build more effective video-ad campaigns.

We tackle RQ2 from two perspectives by using data gathered through the diary. We start by analyzing the overall reasons provided by users for skipping or not video-ad exhibitions (answers to questions D5 and D6 in Table 4.2). For instance, we analyze whether a skip was motivated by users not finding video-ads interesting, or due to personal reasons such as disliking video-ads in general. Next, we analyze the context of the user when the action (skip or not) was taken by using the answer to question D4, aiming at understanding the role this context plays in the user's decision to skip the video-ad exhibition.

4.4.1 Reasons to Skip

We start by focusing on answers given by participants to questions D5 and D6. In D5, participants were asked to indicate whether they had skipped the video ad-exhibition (i.e., "Yes" or "No" question). D6, in turn, is an open question where participants are requested to explain why they chose to skip or not the ad. Recall that participants were instructed to fill the diary whenever they were exposed to a video-ad on YouTube. Therefore each response in the diary corresponds to one exhibition of a video-ad. For this reason, the analysis of questions D5 and D6 allows us to capture properties related to each particular video-ad exhibition that were taken into account by the users when deciding to skip it or not.

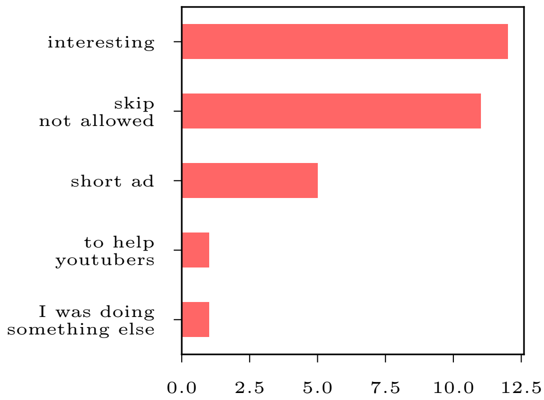

Figure 4.6: Histogram for the categories created through the open coding of question D6 (reasons to watch a video-ad).

In 30 out of the 131 diary entries analyzed (23%), participants answered "No" as to whether they had skipped the video-ad exhibition (D5). Interestingly, this percentage is similar to the one reported by our own previous study where 29% of almost 100,000 video-ad exhibitions were not skipped by the users [1]. Recall that, as explained in the methodology, we used open coding to classify the answers to question D6 into 12 categories (4th column in Table 4.3): 7 categories to skip a video-ad and 5 categories to watch it. Figure 4.5 shows the histogram of frequency of the categories for skipping a video-ad exhibition, while Figure 4.6 shows the histogram of frequency of the categories for watching it completely (not skip).

As shown in Figure 4.5, "uninteresting" was the most popular reason provided by participants for skipping a video-ad exhibition (52 responses). For example, participant P15 provided the following explanation for skipping a video-ad exhibition: *"It was not related to the content of the video"*. We infer from such answer that P15 did not find the video-ad interesting as it was unrelated to the video-content the he was first interested in. "Interesting", on the other hand, was the most popular reason given by participants for fully watching a video-ad (12 responses), as shown in Figure 4.6. Participant P14, for instance, was watching a video about movie facts when a video-ad about a particular movie was displayed. The answer provided by this participant to question D6 was *"The movie trailer was interesting and I wanted to know the name[7] "*.

It is also important to observe that most participants who fully watched the video-ad because they found it interesting think YouTube would be better without video-ads (question S6 of the survey). That is, even though these users would prefer YouTube without video-ads in general, they did enjoy watching some video-ads that they found interesting. This observation reinforces our findings that tailoring the selection of the

[7]This is the translation from the users response, we interpreted as the user being curious about the title of the movie (In Portuguese movie titles are called names).

video-ad to the (current) interests of the users is important to attract and keep user attention, even for those who do not like ads in general.

As shown in Figure 4.5, some participants also skipped video-ad exhibitions because they had seen the same ad before (14 responses). Participant P22, for example, provided the explanation: *"I already watched the same ad before"* to justify skipping the ad exhibition. In fact, this was the third most popular reason given by our participants for skipping video-ad exhibitions, showing that repetition of video ads may not help improve the effectiveness of an ad campaign, but rather it may bother the users and make them avoid (skip) the exhibition. However, this is not always the case: we found one participant (P14) who watched the same ad multiple times just for curiosity: *"I wanted to understand what the ad was about because the last time I couldn't understand it since it was fast"*. Unfortunately, we do not know if this participant was genuinely interested in the video-ad, or rather if she was influenced by our diary and watched it just to be able to answer the questions. Regardless, repetition was considered a reason for skipping video-ad exhibitions a considerable number of times by our participants, suggesting that controlling and restricting such repetition might lead to more enjoyable and thus more effective advertising campaigns.

Next, we look into the relationship between video-ad duration and the user decision to skip its exhibition or not. Figure 4.5 shows that some participants skipped video-ad exhibitions because they were considered too long (6 responses), while Figure 4.6 shows that some participants who watched the video-ad exhibitions completely did it because the video-ad was considered short (5 responses). Indeed, some participants found long video-ads annoying and irritating. As an example, participant P27 wrote *"It is long and annoying"* to justify the skipping, while P26 provided the answer *"It was super long, got on my nerves"*. In contrast, participants are more willing to watch short video-ads. P15, for example, wrote *"The ad was 15 seconds long and I had to watch the first 5 seconds, so I decided to watch the remaining of the ad"*. Therefore, the duration of the video-ad is an important factor and should be taken into account when creating video-advertisements, since short ads are more appealing to users. This result contradicts the one presented by [32]. In this study, the authors show that the duration of the video-ad has a positive impact on its effectiveness. However, the study was performed with non-skippable video-ads and brand name recognition was used to measure the effectiveness. Here, we show that when users are able to skip video-ads, they tend to do it when they consider the video-ad too long.

We also noticed the presence of some participants who do not like (or do no care about) advertisements in general and always skip video-ad exhibitions (8 responses). Participant P3, for instance, provided the explanation *"I don't like ads on YouTube"* to justify the skip, while participant P23 answered *"I am not interested, advertising bothers me"*. On the opposite direction, we also found one interesting case of a participant (P28) who watched the video-ad exhibition until the end in order to help the video content creator:

"I know that Youtubers are paid via ads. If you skip it, they get no money, so I only skip ads if they are longer than 30 seconds"[8] These examples illustrate the existence of two contrasting user profiles. On one hand, there are users who hate video-ads in general and do not watch them in any circumstance. On the other hand, there are users willing to watch video-ads to help others because they are aware of the importance of advertising in generating profit and maintaining the application. We do not know whether the users who reported disliking video-ads know about their importance for the maintenance of the application. Thus, it may be interesting to create campaigns to explain the economy behind most online applications, showing the role of advertising and generating awareness among users.

Other reasons for skipping video-ad exhibitions reported by the participants were: the participant was eager to watch the content (18 responses), it was a habit (2 responses) and the participant already knew the product being advertised (1 response). Regarding other reasons for users to watch video-ad exhibitions until the end, some participants reported that they did not skip the exhibition because the application did not allow it (11 responses) while others reported being busy doing something else and thus did not bother skipping the exhibition (1). In particular, the reasons "to watch content" and "I was doing something else" for skipping and fully watching the exhibition, respectively, indicate that there are other factors beyond the video-ad exhibition which can affect the decision of the user. Some of these factors, which are more related to the context of the user when exposed to the video-ad, are discussed in the next section.

4.4.2 Impact of User Context

We infer the context of the participant when exposed to the video-ad by the response given to question D4. In this question, participants were asked to inform the reasons to watch the *video-content*. As explained in the methodology, we coded the responses to this question creating 4 categories to represent the motivations for watching the video (3rd column in Table 4.3). Figure 4.7 shows the histogram of frequency of these categories. As shown in the figure, participants were mostly using YouTube as a means of entertainment (76 responses). However, there were also cases of participants who were watching the video-content to learn or gain information about a topic (25), to focus on their work or study (17) or because the video was recommended by a friend, an article or by YouTube itself (13).

We notice that the context defined as "to entertain" was the one with the largest presence of video-ads streamed in full (i.e., no skipping), both in absolute and relative

[8]Note once again a negative mention to long video ads.

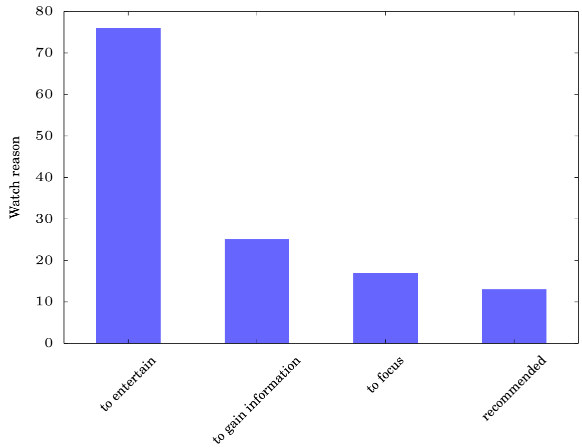

Figure 4.7: Histogram for the categories created through the open coding of question D4 (reasons to watch the video-content).

terms (14 out of 76 cases). Although these numbers can be interpreted as an indication that users using the system for personal entertainment may be more accepting of ads, the diary entries pointed out to no other particular reason for watching the video-ad. That is, these entries simply stated that the ad was interesting.

To investigate more, we then looked into users who were browsing the system "to focus". Seven different participants gave such response to question D4 (producing 17 different responses). Moreover, we found that in all 17 entries, participants used songs as a means to concentrate on their work. Again, on most of these entries users skipped video-ad exhibitions, indicating that their context (they did not want to be disturbed) was their main motivation to skip. However, when looking into these particular entries (to focus while listening to music), we found one exception of a user who streamed the ad in full. In this case, the user indicated that he/she found the ad interesting. Though this is a single example, it shows that accurate algorithms can still provide users an entertaining experience even when they are focused on other tasks. In fact, in the setting of traditional banner ads, clicks are usually observed in very small fraction [43], and effectively matching ads to content is one of the most studied problems in online advertisement. This small example shows that on video-ads, effective algorithms can still succeed in motivating full streaming of video-ads, and consequently, produce more revenue to providers and creators.

Another interesting observation based on user context is when users are browsing the system because of a recommendation. We had 13 entries of participants who were watching a recommended video (e.g., from friends or other websites). Yet, in no single case the user streamed the video-ad in full. This finding further shows that the reason that led the user to stream videos in the first place can also impact his/hers actions.

Although it is not possible to measure a clear impact of the user context on the skipping behavior through our study, the entries in the diary indicates that the context can be important to users. Therefore, understanding user context may help increasing advertising profits by exhibiting video-ads when users are more willing to watch them. In the next section we study the role of the content of the video-ad and the video-content on the skipping behavior of users.

4.5 Understanding the Role of Content Itself

As we have discussed in the previous section, the user finding or not the video-ad interesting is an important reason behind the decision to skip or not its exhibition. Moreover, as also observed, users often skip video-ad exhibitions because they are eager to watch the video-content. Motivated by these observations, we tackle RQ3 (*What is the role of the content of the video-ad and the video it is paired with in the user's decision to skip or not an exhibition?* by analyzing the role of the content of these two videos – video-ad and video-content – on the user skipping behavior.

Recall that in question D7 of the diary participants were asked to inform whether they had understood the video-ad. Complementarily, participants who answered "Yes" to this question were asked to provide information on what the video-ad was about using their own words (D8). In 106 out of the 131 entries of the diary, participants answered "Yes" to D7 and provided descriptions of the video-ads. We used these descriptions to shed some light into the role of the topic of the video-ads on user skipping behavior.

Before analyzing these responses, we first briefly discuss the 25 exhibitions for which users were unable to inform what the video-ad was about (answered "No" to question D7). Out of them, 22 exhibitions were skipped by the participant, either because the video-ad was not interesting or because the participant was eager to watch the video-content. We can observe from these exhibitions that capturing user attention in the first initial seconds is a relevant issue for advertisers. Interestingly, we observe that, for three video-ad exhibitions, despite watching the video-ad until the end, the participant was still unable to explain what it was about. Though these cases are few, they illustrate the challenges of assessing the effectiveness of online advertising. A video-ad may be fully streamed to the user without him/her actually paying enough attention to grasp what it is about. Thus, the number of video-ad streams itself may be a misleading metric. Indeed, the design of more realistic metrics is an open research topic on online advertising in general. For instance, in a recent study, [60] argued that click through rate (CTR) is not a good metric to measure the quality of native ads.

We now turn to the exhibitions for which participants reported they had understood the video-ad. As performed for the other open questions, we also used an open coding scheme to categorize the descriptions of the video-ads provided by the participants. In total, 12 categories were created (2nd column in Table 4.3), covering a wide range of

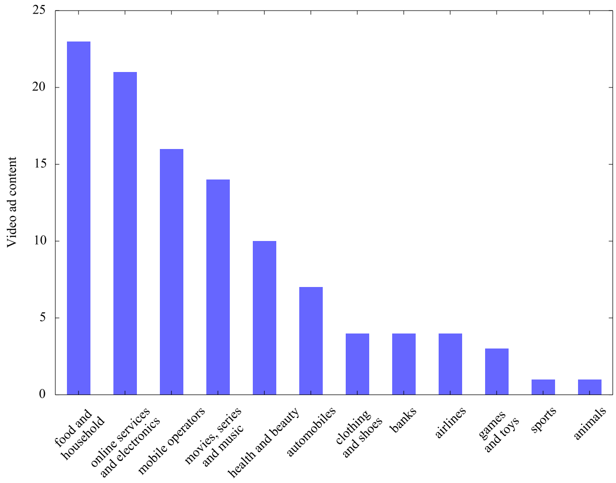

Figure 4.8: Histogram for the categories created through the open coding of question D8 (product/service being promoted).

services and products, as shown in the histogram in Figure 4.8. The three most frequent categories in our dataset cover advertisements about “food and drinks", followed by “online services and electronics", and “mobile operators". There were also a great number of video-ad exhibitions that consisted of trailers of movies, TV shows, and music. Animals and sports were the least popular categories. Moreover, 79 exhibitions were skipped and 27 were fully watched by the participants.

By cross-referencing the responses given to D5 and D6, we further focus our attention on exhibitions for which the participant reported understanding the video-ad and whose skipping decision is directly related to the participant interest (or not) in the video-ad: (1) exhibitions the participants did not skip because they found the video-ad interesting (12 exhibitions), and (2) exhibitions the participants did skip because they found them uninteresting (47).

We notice that the exhibitions the participants did not skip because they found the video-ad interesting are related to 5 different categories of ad products/services. The most frequent one is “movies, series and music” (5 exhibitions), followed by “online services and electronics” (3), “food and household” (2), “mobile operators” (1), and “games and toys” (1). By further looking at the explanations provided by the participants for not skipping these exhibitions, we notice that the content of the ad and the way it was designed were important to some participants. For example, P14 wrote: *“I found the ad interesting, the way it was designed”* whereas P9 justified fully watching the ad because: *“I was curious to know what the ad was about”*. Another participant (P16) did not skip the video-ad because of its content: *“The ad was the trailer of the second movie ‘Alice Through the Looking Glass’. It caught my attention since the beginning”*.

We also notice that, in some cases, the participants’ interest in the video-ad (and their decision not to skip the exhibition) was inferred from the similarity between the topic of the video-ad and the topic of the video-content the user originally requested to watch.

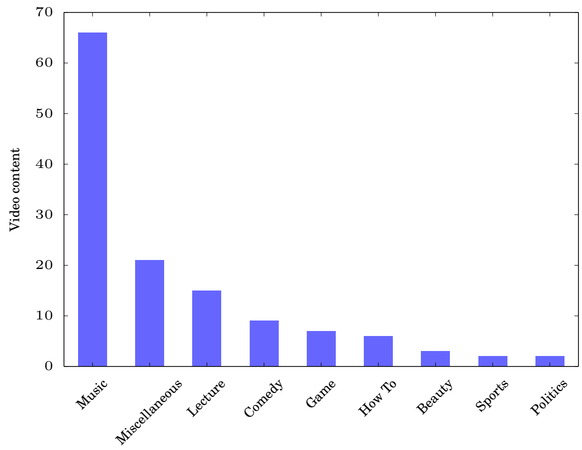

Figure 4.9: Histogram for the categories created through the open coding of question D3 (video content).

For example, participant P11 provided the following explanation for watching the video-ad until the end:*"It was similar to the content of the video"*. The response provided by P16, discussed in the last paragraph, also illustrates a similar effect: the user wanted to watch a video about movies and decided to watch an entire movie trailer. These examples suggest that the similarity between the topics of the video-ad and the video-content may be an important factor for attracting user attention towards the video-ad. This is an interesting observation, since in our own previous effort [1] we showed that, on YouTube, video-ads are often paired with video-contents that are mostly dissimilar to them (in terms of topic). The results discussed in this chapter suggest that similarity may be a factor to be explored when pairing video-ads and video-contents to enhance user experience.

Regarding the exhibitions skipped by the participants because they found the video-ads uninteresting (47 exhibitions), the 5 most frequent categories of descriptions of the video-ad provided by the participants are: "online services and electronics" (11 exhibitions), "mobile operators" (9), "food and household" (9), "health and beauty" (7) and "automobiles" (4). Interestingly, in 5 such exhibitions the participants reported that the content of the video-ad was indeed related to their interests (question D9). In these specific cases, the motivations to skip as pointed out by the users did not give us much insight further than the fact that the ads were uninteresting to them (e.g., users simply wrote that they disliked the ad). Nevertheless, in such examples, users potentially could have been attracted to the video-ad (as it is indeed related to the user interest) and, yet, they found it uninteresting and decided skipping its exhibition. These examples illustrate situations in which the users reported that the ad was on a topic of interest, but the video-ad itself was found uninteresting by the users and they were not motivated to watch it. Such examples serve to illustrate the (negative) role that the video-ad itself may have on the user's decision to watch it or not.

Finally, we look into the role of the video-content in the user skipping behavior. To that end, we analyze the responses given by participants to D3, in which they were asked to describe the YouTube content they had requested (video-content). Once again, open coding was applied to categorize the participants' responses, and the histogram of the 9 identified categories (1st column in Table 4.3) is shown in Figure 4.9. As shown, "Music" is the most popular category of video-content. Video lectures were also frequently requested, possibly due to the large number of university students in our participant pool. "Miscellaneous" is also a large category since we aggregate various videos that could not fit in a single category.

To understand the role of video-content, we looked into the specific video-ad exhibitions that were skipped by the participants because they were eager to watch the video-content ("to watch content" in D6). The most popular codings for video-contents mostly followed the same ordering as that shown in Figure 4.9. Initially, this led us to believe that the role of video-contents were hard to capture using just our diary. Nevertheless, we decided to further focus on these settings by investigating if participants understood what the video-ad was about. Surprisingly, we could not find a single case in which participants reported that they could capture the content of these video-ads (D7 in Table 4.2). This observation suggests that, in some situations, users may not pay attention to the video-ad at all, regardless of its content, simply because they are impatient to watch the requested video-content. Thus, while we could not find evidence of individual contents having some impact on user behavior, we do find that when our participants were eager (when watching any kind of content), they would always skip the ad. One important setting for future work is on capturing the eagerness of users, possibly with browsing data. By achieving such goal, providers can mitigate user fatigue with ads by simply not showing video-ads when users quickly want to watch the content.

In summary, our observations suggest that the content of both video-ads and video-contents can play a significant role in the user decision to skip a video-ad exhibition. For instance, users may skip video-ads that fail to capture their attention early on, even if these ads cover a topic of interest. Moreover, pairing video-ads with video-contents that are similar (in terms of topic) may make video-ads more appealing to users. Yet, when users are eager to watch the requested video-content, they might not pay attention to the video-ad, not even to detect whether they find it related to their own personal interests or not.

4.6 Summary of Findings

In order to provide a better understanding of user behavior when exposed to video-ads, in this chapter we presented an exploratory study to assess the perception of users about video-ads and, in particular, why users decide to view or skip ads. Our study focuses on three research questions: (1) How do users perceive the value of video-ads on YouTube? (2) Why do users skip (or not) video-ads? and (3) What is the role of the content of the video-ad and the video it is paired with in the user's decision to skip or not an exhibition?

Each research question captures a different perspective of YouTube video viewers. On RQ1 we focused on understanding how users perceive the value of video-ads in the system. To tackle this question we employed a survey and asked users if they would prefer YouTube with or without video-ads. We also inquired if users would be willing to pay not to be exposed to such ads. The results of our survey are both interesting and contradictory. While most users usually do not view video-ads as a positive feature, they are also, mostly, unwilling to pay for an ad-free service. Among other implications, this result shows that developing ads that entertain and do not detract from the users' experience is a relevant, yet challenging, issue.

On our second and third research questions (RQ2 and RQ3) we analyzed the use of the skipping feature to assess the perception of users on individual video-ads exhibited to them. We employed a structured diary consisting of a few multiple choice questions and some open questions, and asked users to add a new entry to this diary whenever a video-ad on YouTube was exhibited to them. As part of the diary, users were asked to indicate whether they had skipped the video-ad exhibition.

On RQ2 we aimed at understanding *why* users skip a video-ad exhibition. Among other findings, our results indicate that users often skip video-ads because they had seen the ad before (repeated exhibition) or the video-ad is very long or uninteresting. Past efforts that looked into offline ads indicated that repeated exhibitions and long exposures may help brands. In contrast, we find that users are actually annoyed by these factors. This result may represent a change in setting, since users on social media can explicitly skip and go on to view their content of choice. We also looked into the user context, captured by their own words on the reasons why they were streaming a particular video on YouTube. On this second setting, we find that when users are focused or studying, very rarely will they watch an ad. This finding serves as evidence that the user context at the time a video-ad is shown matters. Capturing such context to decide whether to show ads and of what type at a given moment will improve overall user experience. Strategies to capture the context of users is left as future work. Some suggestions consists of looking at the category of the video-contents as a proxy to the context of the user, as well as performing user studies to understand different profiles of users.

On RQ3 we focused on the streamed video-ads and video-content themselves. Our goal with this question was to understand the impact of the content of the videos being streamed. To achieve this, we focused on the cases where users explicitly stated that they found ads interesting. That is, we looked at the diaries and filtered entries in which users reported not having skipped and having found the ad interesting. Based on these cases, we were able to analyze the impact of the video-ad itself. Our results here are interesting since they show cases where users have a personal interest on the product being sold, but skip the ad in any case. We also find that ads that are similar to the video-content can also attract user attention. Finally, our results show that when users are eager to watch the video-content, they tend to pay very little attention to the ads regardless of their content. Based on these results we can state that video-ads and video-content do play a role in users' decisions. Their roles, however, are not trivial to delimit and assess since they are mediated by specific settings in our diary (e.g., when users are eager to watch content). Automatically inferring the role of content in user perceptions, specially when taking into account user context (e.g., when visiting a video recommended by a friend), is a challenging task and we hope our work motivates future efforts in this direction.

It is also important to mention the limitations of our work. First and foremost, our study is performed with a small participant pool. Nevertheless, it is important to point out that small samples are expected in diary based studies like ours. Moreover, our original goal was to shed some light into how users perceive video-ads. That is, their motives to skip or not exhibitions. With our diary we achieved this goal, showing how different factors may impact users' decisions. Generalizing our findings to other settings and participant pools is an important effort for future work. Secondly, several other factors such as cultural behavior may impact user experience. Understanding such factors is left as future work.

In this Chapter we discussed video-ads from the perspective of the users. In the next Chapter, we will present a complementary vision, shifting our attention to service providers and content creators. With both chapters we present an overview of video-ads on YouTube, providing insights that can be used to create better video-ads, achieving the goal of this thesis.

Chapter 5

An Overview on Video-Ad Monetization

In this chapter, we present a study of video-advertising focusing on monetization. We start by motivating our study (Section 5.1) and presenting a characterization of the video-ads and content creators in our campus dataset, focusing on their success in generating and attracting revenue (Sections 5.2 and 5.3, respectively). We then shift our attention to the content of the video-ads, analyzing the relationships between multimedia properties of video-ads and their success (Section 5.4). Finally we summarize and discuss our findings (Section 5.5).

5.1 Motivation

As discussed in Chapter 2, there are only a few prior efforts on video-advertising. In general, these studies only characterized a few properties of video-ads and did not consider their role in generating revenue to content creators and to YouTube itself. Also, prior analyses of video-ads only focused on metadata information, such as the duration of the ad in seconds or its category (e.g., Music, Entertainment). Thus, our goal in this chapter is to analyze video-ads on YouTube from the perspective of monetization, deepening our understanding of the video-ad market. In particular, we are interested in answering the following research questions:

RQ1: How successful are video-ads in generating revenue?

RQ2: How successful are channels in attracting revenue?

RQ3: How are multimedia properties of video-ads characterized? Are these properties related to monetization?

In order to answer our first research question, we analyzed the video-ad exhibitions in our campus dataset. As explained in Chapter 3, 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. Thus our goal in RQ1 is to study 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.

To address our second research question, we explored our campus dataset and also the data collected through the YouTube API. Recall that 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. Therefore, in RQ2, we analyze the popularity of channels and the success of content creators in profiting from video-ads associated with their contents.

To tackle the third research question, we used a new dataset containing video and audio features of the video-ads in our campus dataset. We crawled this dataset to provide a study of the content of the video-ad and its impact on the success of the ad in generating revenue. We characterize some classic features extracted from the audio and video of the video-ads, correlating these features with monetization. Our goal here is to bring forth the role of the multimedia content of the video-ad on its success.

With these three research questions, we complement previous efforts that have not looked into monetization and multimedia properties of video-ads. Our contribution is to shed some light on this topic, providing the first characterization of monetization on Youtube. We show our main findings related to each of the three research questions in the following sections.

5.2 Video-Ad Monetization

In this section, we tackle *RQ1: How successful are video-ads in generating revenue?* As mentioned earlier, 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 to the marketer.

First, we look at the number of exhibitions that generated revenue. These *monetized exhibitions* (see Section 3.1) are defined by video-ads streamed over 30 seconds. Out of the 99,658 video-ad exhibitions in our campus dataset, 34,093 were monetized (34%). As we have discussed in our preliminary study, which is summarized in Section 3.3, users will likely skip video-ads as soon as possible, leading to fewer monetized ads as we see here. Our campus trace may not necessarily reflect the global fraction of monetized views.

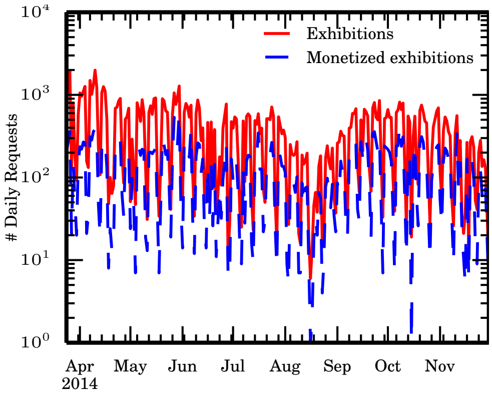

(a) *Number of exhibitions per day (total and monetized)*

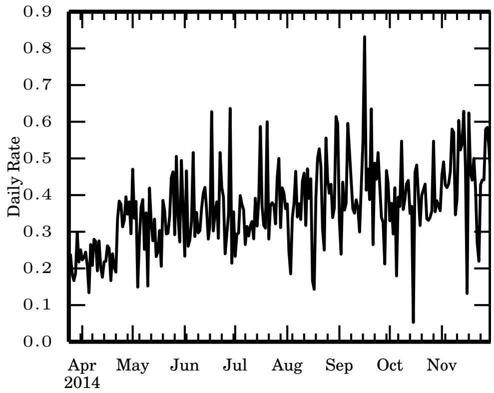

(b) Daily Fraction of monetized exhibitions

Figure 5.1: Overview of video-ad exhibitions in our dataset: volume of exhibitions and rate of exhibitions that generated revenue.

However, as the rate of video content streamed over the Internet reportedly exceeds the mark of 2.7 million streams per minute[1], 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 [22], this estimate matches others that stated YouTube may generate billions of dollars yearly, translated to tens of millions of dollars daily[2].

We now look into these numbers on a daily basis. Figure 5.1(a) shows the total number of daily video-ad exhibitions as well as the number of monetized video-ad exhibitions. Complimentary, Figure 5.1(b) shows the fraction of exhibitions that were monetized per day. The number of video-ad exhibitions per day varies 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 daily peak of 83% of the video-ad exhibitions generating revenues. In this particular day, there were 208 exhibitions of 42 unique video-ads and the monetized exhibitions came from only 15 of these ads. Out of curiosity, we looked 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 5.1(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. Out of the 5,667 unique video-ads in our campus 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 5.2, the number of exhibitions that were profitable per ad is often small. The figure shows the distribution of the fraction

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

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

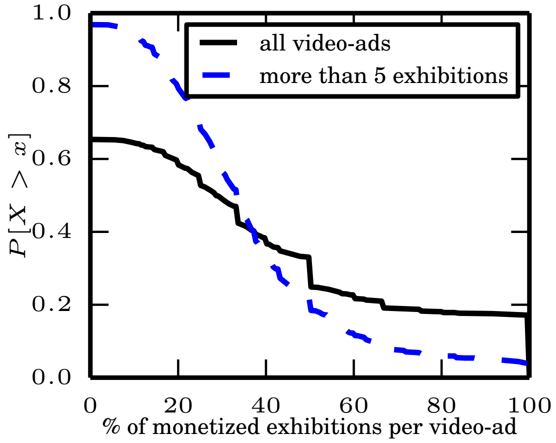

Figure 5.2: Distributions of the fraction of monetized exhibitions per video-ad (all video-ads and only video-ads with at least 5 exhibitions).

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 exhibitions (35% of all video-ads). This comparison is useful to filter out the vast majority of video-ads which were paired only a couple of times (in our dataset), and thus have a small chance of generating revenue. We initially look at the the distribution of all video-ads. From the figure we can see that around 20% of the video-ads generated revenues in over 60% of their exhibitions. In other words, some video-ads trigger enough interest in the analyzed user community to generate revenue in most of their exhibitions. This observation motivates further studies that aim to understanding the effect of video-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 [32].

Turning our attention to the video-ads that had more than 5 exhibitions, we can see a change in behavior. Here, most video-ads (over 95%) generated revenues at least once. This suggests that repetition may increase the chance of monetization. It is important to notice that exhibiting a video-ad multiple times may increase the chance of monetization as it reaches a larger audience with different users. However, exhibiting a video-ad unaware of the target user may bring discomfort and actually hurt monetization, as suggested by our study in Chapter 4. Moreover, 18% of the video-ads had more than half of their exhibitions monetized whereas 6% of them had more than 80% of their exhibitions monetized. Nevertheless, at the tail of the distribution (after 40% of exhibitions being monetized), video-ads with more exhibitions have actually a smaller fraction of those views being monetized, if compared to all video-ads. This last effect likely stems from those ads that are paired only a couple of times and are always monetized. Again, various factors may play a role in monetization.

In order to uncover properties of video-ads that may be related to the success in attracting the attention of users and generating revenues, we compare two groups: (1)

video-ads that generated revenue at least once and (2) video-ads that did not generate revenue at all (as captured by our dataset). By contrasting these two groups, we aim at shedding light into possible factors that might have led to monetization *at least once*.

The first feature that we analyze is the duration of the video-ads. The average duration of the video-ads in the first group is 96 seconds, with a 95% confidence interval ranging from 90 to 102.94. In contrast, the average duration of the video-ads in the second group is 137 seconds, ranging from 128.05 to 146.50, with 95% confidence. Therefore, with 95% confidence, the durations of the video-ads in the two groups are significantly different, indicating that shorter video-ads have a higher chance of being monetized. In Chapter 4 we discussed that users are more willing to watch video-advertisements when they are short in duration. In fact, long video-ads was a reason provided by the participants to skip the exhibitions. Based on our study presented in Chapter 4 and the results we here present, our findings show some evidence that shorter video-ads may attract more interest from users.

We also analyzed 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, that is, there is no significant difference in the categories of the video-ads in the two groups. In our results, the null hypothesis was rejected with p-value $p = 0.05$. This result offers evidence of dependence between categories of the video-ads and monetization. We further looked into the categories individually searching for possible concentration of monetized exhibitions in a few subset of the categories. The video-ads in our dataset are from 15 different categories. We found monetized video-ad exhibitions of all 15 categories. Indeed, the fraction of monetized exhibitions varied from 42% (Music category) to up 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 video-ads that are successful in leading to monetization tend to be shorter and that the category of the video-ads are to some extent related to their chance of generating revenue.

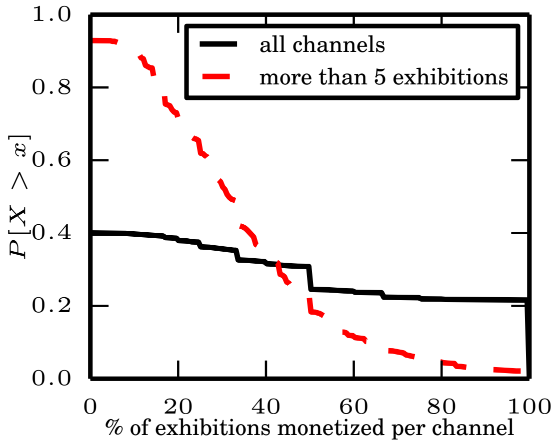

Figure 5.3: Distributions of the fraction of monetized exhibitions from the channel perspective

5.3 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, we perform a study of revenue from the perspective of the channel, tackling *RQ2: How successful are channels in attracting revenue?* That is, we turn our attention to the content creators who 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 in our campus dataset and the number of these channels that received some revenue from monetized video-ad exhibitions. For each video-ad exhibition, we identified the channel of the video-content using our API dataset, explained in Chapter 3. We were not able to collect the public API information for all video-contents: specifically, we could 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 who associated their content to video-ads were able to profit from YouTube.

As done for video-ads, we also analyzed the fraction of monetized video-ad exhibitions for each channel. To that end, we first aggregated the monetized exhibitions by channel. Figure 5.3 presents the CCDF of the percentage of all video-ad exhibitions that were monetized 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. As shown in the figure, 60% of channels are never monetized. However, 20% of channels

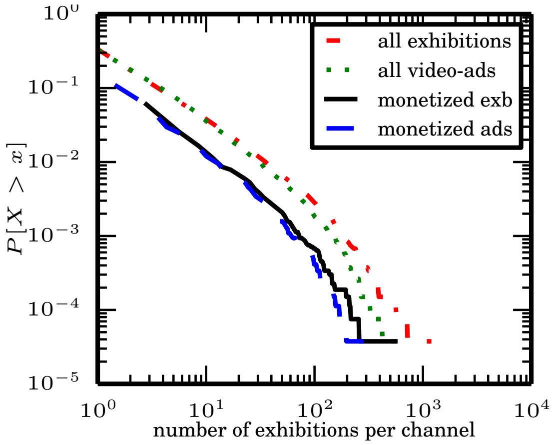

Figure 5.4: Distributions of exhibitions per channel

have over 90% of their exhibitions monetized. As with Figure 5.2, these effects may stem from the channels with only one or two pairings. Thus, we also considered only channels with more than five video-ad exhibitions, finding that 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 looked at the popularity of channels in terms of the number of video-ad exhibitions. Figure 5.4 shows the distributions of the number of video-ad exhibitions and monetized video-ad exhibitions per channel. The distributions are very skewed, that is, most channels are associated to video-ads just a few times, while a small fraction of channels are often target of video-ad pairings. For instance, 7% of the channels have more than 5 video-ad exhibitions while 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 only the monetized exhibitions, the numbers are even lower. Just 2.5% of the channels have more than 5 monetized video-ad exhibitions while 0.2% more than 50. Thus, only a few channels were able to generate a lot of revenue. Figure 5.4 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. However, most of them were associated to video-ads only 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 video-ads exhibited in each channel are often distinct.

5.4 Multimedia Properties and Monetization

In this section we tackle *RQ3: How are multimedia properties of video-ads characterized? Are these properties related to monetization?* We first present the methodology used to collect our complimentary dataset of multimedia properties (Section 5.4.1). Next, we start our analysis by exploring multimedia properties extracted from the video content of the ads (Section 5.4.2). We then shift our attention to the audio of the video-ads (Section 5.4.3).

5.4.1 Multimedia Properties Data Collection

To perform our study of the impact of multimedia features on video-ad monetization, we extended our previously collected campus dataset, described in Section 3.2.1, to collect the audio and video of each unique video-ad exhibited on our campus dataset. We used a Python library called *pytube* for downloading the YouTube videos[3]. We temporarily stored[4] the videos (including the audios) in mp4 format and each content was downloaded in 720p resolution. When the content was not provided in this resolution, we chose the highest one available.

Having the video and audio contents of each video-ad, we proceeded to extract features from the multimedia content. First, we extracted features from the content, focusing on attributes that could provide us an overview of the video-ads in terms of colors, brightness and presence of objects or humans. To that end, we first used a multimedia framework called *FFmpeg* to convert each video into a set of frames[5]. Then, for each frame we extracted the *histogram of colors*, *number of faces* and *number of blobs*. The histogram of colors and the number of faces were extracted through the use of an open

[3] `https://github.com/nficano/pytube`

[4] We stored this content only temporarily to extract the features. Original content was then deleted, to avoid storage of the video themselves. We only kept the metadata used in our study.

[5] `https://www.ffmpeg.org/about.html`

source library for computer vision called *OpenCV*[6]. The histogram of colors represents the number of pixels in an image that have colors in each of a fixed list of color ranges. We used a list of 32 color ranges and therefore, for each frame, we have 32 numbers representing the number of pixels in each one of the ranges. The number of blobs on each frame was extracted using *scikit-image*, an image processing toolbox [7]. Blobs stands for Binary Large Objects and they represent regions in an image that differ in terms of different properties (for instance brightness) compared to surrounding regions. In short, blobs can be interpreted as the number of objects in an image.

Next, we changed our focus to the audio content of the video-ads. Since we collected the videos in mp4 format, we used the framework *FFmpeg* to extract the audios from the videos, storing them in WAV format. In Section 5.2, we saw evidence of dependence between categories of the video-ads and monetization. In particular, the Music category was the one with the smallest fraction of monetized exhibitions, suggesting that the use of music in video-advertisements may not help to attract user attention. Therefore, to investigate more about the use of music in video-ads, we applied a speech versus music discriminator tool to detect whether the audio of a video-advertisement is a music or a speech. The tool is based on the open source Opus codec and it returns the probability of every 0.06 seconds of the audio being a music[8]. Thus, for each video-ad in our dataset, we have a time series of the music probabilities.

We downloaded the videos of the video-ads for two days, on May 8^{th} and 9^{th}, 2017. We were able to download the videos (and audios) of 4306 unique video-ads. We could not download some video-ads due to video-deletions, privacy settings or other problems with the video and the download.

Finally, we explore our multimedia dataset. Figure 5.5(a) shows the distribution of the size of the video-ads in megabytes and Figure 5.5(b) shows the number of frames in each video-ad. We can notice that more than 80% of the videos is larger than 1 MB, although only 20% is larger than 10 MB. The sizes vary greatly across all video-ads: the average size is 9.6 MB and the standard deviation is 18.5 MB. As expected, the number of frames also presents a high variation, with an average of 2564 frames per video-ad and a standard deviation of 2077 frames. All video-ads have more than a 100 frames and around 60% of them has more than 1,000 frames.

[6]`https://opencv.org/`
[7]`http://scikit-image.org/`
[8]`https://github.com/jzombi/opus_sm`

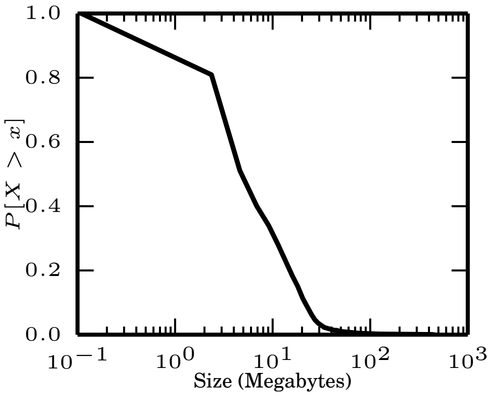

(a) *Size in MegaBytes*

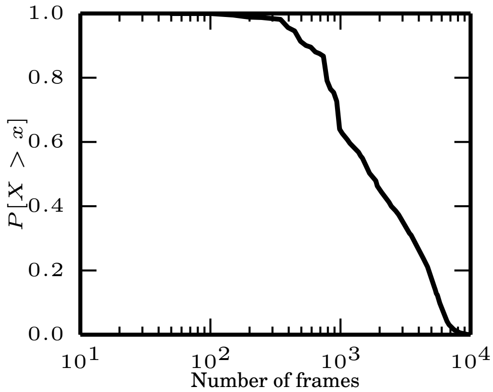

(b) Number of frames

Figure 5.5: Size and number of frames of video-ads in our campus dataset.

5.4.2 Video Properties

As mentioned in Section 5.4, we analyze the content of the video-ads using three multimedia features extracted from the videos: number of blobs, number of faces and the histogram of colors. For the latter we computed the Shannon entropy [16] as a measure of diversity of colors in each initial frame, as will be discussed below. We extracted frames for the first 5 seconds of each video-ad. Recall that YouTube usually allows users to skip video-ads after 5 seconds of exhibition. Also, the skipping behavior of users is directly tied to monetization. Therefore analyzing only the first 5 seconds of the video-ads allows us to focus on the part of the ads that may have an impact on the decision of users to skip or watch an advertisement. We generated the three aforementioned features for each extracted frame, aiming at assessing to which extent the presence of larger numbers of blobs (objects), faces or even a greater diversity of colors in these initial frames may be related to a higher success in monetization.

We start by presenting in Figure 5.6(a) the CCDF of the average number of blobs per (initial) frame for each video-ad in our multimedia dataset. The average number of blobs is moderately small (notice that both axes are in log scale): 47% of the video-ads have an average number of blobs per frame smaller than 10, and only 0.78% of the video-ads have an average number of blobs greater than 100. As explained in Section 5.4.1, the number of blobs can be seen as the number of objects in the frame. Thus, the video-ads in our dataset are mostly composed by frames without many objects. This result can be explained by the nature of video-ads itself, since most of them are promoting a product or service and in general are focused on showing the product and its details.

Next, in order to uncover possible relations between the number of blobs on a video-ad and its success in attracting the attention of users, we show in Figure 5.6(b) the

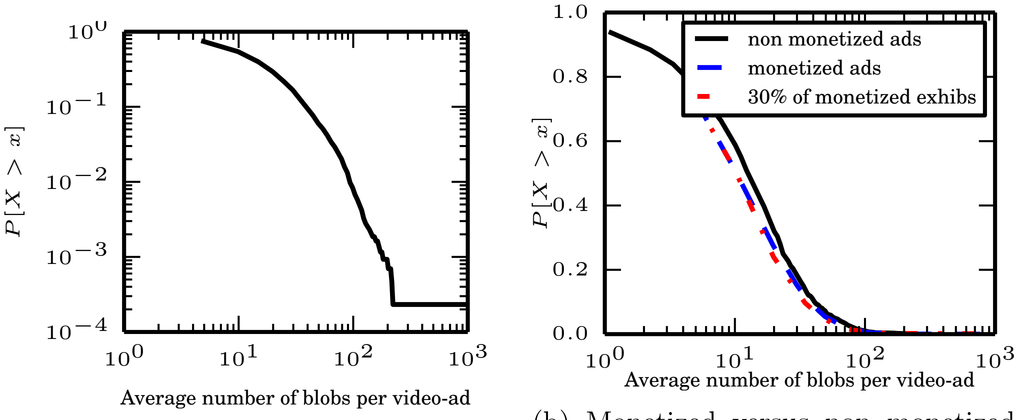

(a) All video-ads

(b) Monetized versus non monetized ads

Figure 5.6: Distributions of average number of blobs for the first 5 seconds of the video-ads.

CCDF of the average number of blobs per (initial) frame for video-ads in three different groups. Specifically, we show the distributions for: (i) video-ads that were never watched for more than 30 seconds (or completely), and thus never generated revenue, in all its exhibitions in our campus dataset (1,565 video-ads); (ii) video-ads that had at least one monetized exhibition (2,741 video-ads); and (iii) video-ads with more than 5 exhibitions and that were monetized in at least 30% of their exhibitions (793) video-ads). The latter group was aimed at reducing noise, as there is a great number of video-ads with only a few exhibitions. The distributions are quite similar, with the two groups corresponding to video-ads that were monetized at least once and at least 30% of the time presenting a slightly lower average number of blobs per frame. In order to assess whether there is a significant difference in the average number of blobs for monetized and non monetized video-ads, we applied a two sample Kolmogorov-Smirnov test [54]. This test is used to check the null hypothesis that 2 samples are drawn from the same distribution. We used as our two samples the average number of blobs for video-ads that were never monetized and the average number of blobs for video-ads that were monetized in at least 30% of their exhibitions. The null hypothesis was rejected with 95% confidence. Therefore, we can conclude that, with that level of confidence, the two samples do not belong to the same distribution.

Now we turn our attention to the number of faces detected in each frame of each video-ad in our dataset. Again, we start by presenting the CCDF of the average number of faces for *all* video-ads in our dataset. Figure 5.7(a) shows that this number is typically small: only 3.5% of the video-ads have 1 or more faces per frame, on average. The maximum number of faces detected on a frame was 13 and the minimum was 0. Although the average number of faces is very small, the presence of faces, that is, the presence of humans in the video-ads is more common. We find that 60% of the video-ads have at least

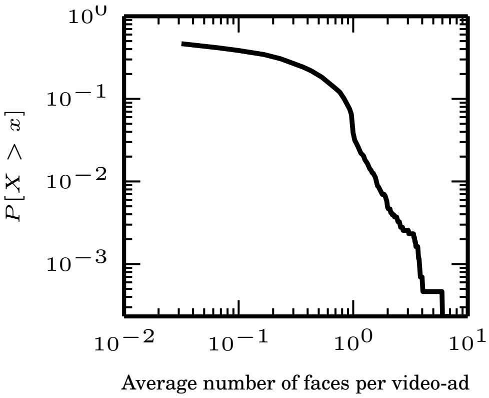

(a) All video-ads

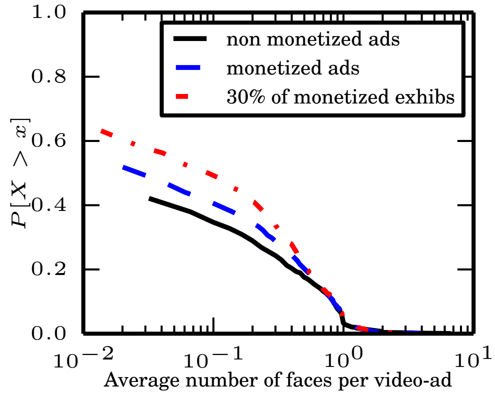

(b) Monetized versus non monetized ads

Figure 5.7: Distributions of average number of faces for the first 5 seconds of the video-ads.

one frame with a face detected, indicating that advertisers often use people to promote the products and services.

In Figure 5.7(b) we show the distribution of the average number of faces for the three aforementioned groups: video-ads that were never monetized, video-ads that were monetized at least once, and video-ads with more than 5 exhibitions and that were monetized in at least 30% of them. The distributions suggest that the more successful the video-ad is in generating revenue, the higher is its average number of faces. Once again we applied the Kolmogorov-Smirnov test to compare this feature for video-ads that were never monetized and video-ads that were monetized at least 30% of the time. The results showed that with 95% confidence, the null hypothesis that the two samples were drawn from the same distribution was rejected. Therefore, the number of faces may play a role in the success of video-ads in attracting user attention. However, this result should be interpreted with care, as it is very difficult to isolate the impact of only the number of faces on the success of advertisements. Sometimes the video-ad may have just one face per frame, but it is a face of a very famous celebrity. Other times it may show a crowd corresponding to the audience of a very famous soccer match. In both situations, the video-ad may be more appealing to users.

Finally we present our analysis of the histograms extracted from each initial frame of each video-ad in our multimedia dataset. We used histograms with 32 color ranges, therefore for each frame we have the number of pixels in each of the 32 ranges. In order to measure the diversity of colors on each frame, we calculated the Shannon entropy for the distribution of pixels [16]. The entropy is a measure of uncertainty and it is calculated by the following equation:

$$H(X) = -\sum_{i=1}^{n} P(x_i) log_b P(x_i)$$

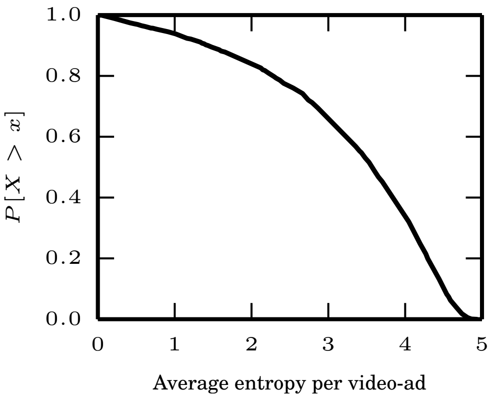

(a) All video-ads

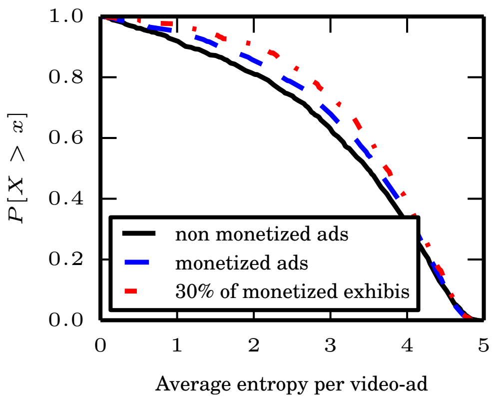

(b) Monetized versus non monetized ads

Figure 5.8: Distributions of average entropy of color histogram for the first 5 seconds of the video-ads.

where X is the discrete random variable with possible values $\{x_1, x_2, ..., x_n\}$ and b is the base of the logarithm used. Thus in our context, X is the discrete random variable that counts the number of pixels in each one of the 32 color ranges and we calculated the entropy using logarithm with base two. Thus, in the present case, the entropy values may vary between 0 and $log_2 32$. A low entropy means that the uncertainty about the random variable X is low and therefore the pixels are concentrated in just a few color ranges. A high entropy means high uncertainty suggesting that the pixels are distributed across different color ranges.

We calculated the Shannon entropy for the initial frames extracted from each video-ad and then we analyzed the average entropy of frames per video-ad. Our goal in analyzing the average entropy was to capture the dynamism of the video-ads. Videos that are very colorful and have different scenes may present a higher average entropy, while monotone videos may present lower averages. The CCDF of the average entropy per video-ad is shown in Figure 5.8(a). Once again, we also analyzed the distributions for video-ads in three selected groups, as presented in Figure 5.8(b). Overall, the video-ads are very diverse in terms of colors: 84% of the video-ads have an average entropy grater than 2 and 34% have an average entropy greater than 4. When comparing the three groups of video-ads, we can notice that the distributions are slightly different, with more successful video-ads presenting higher average entropies. Once again we applied the Kolmogorov-Smirnov test and the null hypothesis that the samples from video-ads that were never monetized and that were exhibited more than 5 times and were monetized at least 30% of the time were drawn from the same distribution was rejected (with 95% confidence). Thus the distributions are significant different, indicating that the color of the frames may be a factor that can influence the skipping behavior of users. Video-ads with more diversity of colors can be interpreted as less monotonous and thus more entertaining to users.

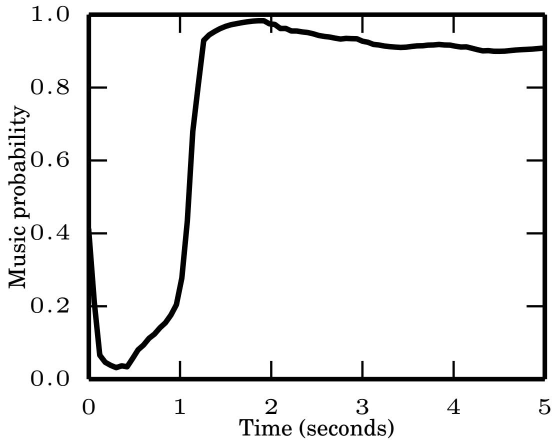

Figure 5.9: Example of a time series for the probability of the audio of a video-ad being music over time, during the first 5 seconds of streaming.

In sum, most video-ads have small numbers of blobs and faces and the colors are usually diverse, with most video-ads presenting pixels well distributed across the color ranges. Although we did find evidence that such features are correlated with the success of the video-ads in generating revenues, precisely defining the extent of such relationships is quite challenging. Here we presented a preliminary study that suggests that these features may have an impact on the user's decision to skip video-ad exhibition early on, motivating further studies to explore other multimedia features.

5.4.3 Audio Properties

As mentioned in Section 5.3.1, we downloaded the audio content of 4,306 unique video-ads exhibited in our campus dataset aiming at analyzing the relationship of audio properties and monetization. Specifically, we considered the presence of music as the main feature. Recall that in Section 5.2 we found that video-ads from the Music category were the ones with the smallest fraction of monetized exhibitions. Therefore, our aim here is to investigate the impact of music on the success of video-ads in attracting revenue. Using a speech versus music discriminator tool, we generated a time series of the probability of the audio being a music for successive 60 millisecond time windows. As done for video, we focused on the first 5 seconds of the video-ads, as this may be determinant content on the user's decision to skip the exhibition or continue streaming it. We were able to generate the time series of probabilities for 4218 out of the 4306 ads[9] and each time series is composed of 84 points (i.e., 84 non-overlapping time windows of 60 millisecond duration from 0 to 5 seconds). As an illustration, Figure 5.9 shows the time series for a video-ad in our dataset.

[9]We were not able to generate the time series for some video-ads due to the quality of the audio.

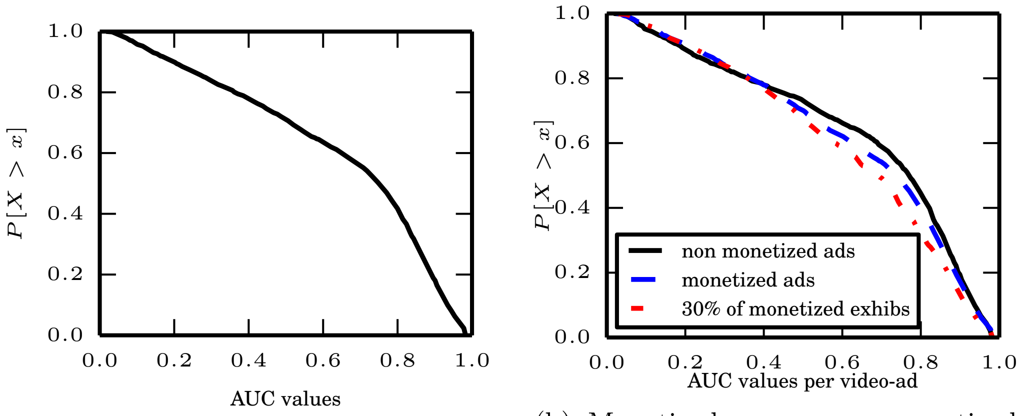

(a) All video-ads

(b) Monetized versus non monetized ads

Figure 5.10: Distributions of normalized AUC values of the probability of being music for the first 5 seconds of the video-ads.

We start by measuring the area under the curve (AUC) of each time series. The AUC is useful to give us a first cut estimate of the probability for all video-ads. A large area under the curve of probabilities indicates a high probability of the audio of the video-ad being a music through most (if not all) of its initial five seconds. A small area under the curve on the other hand indicates a low probability of the audio being a music. In order to facilitate the interpretation of the values, we normalized the areas to be between 0 and 1 (that is, we divided the calculated AUC for each video-ad by the total AUC possible). In Figure 5.10(a) we show the CCDF of the normalized AUC values. The probabilities of the video-ads being music are usually high: 71% of the video-ads have an AUC value larger than 0.5 and 42% of them have AUC larger than 0.8. However, only 18% of the video-ads have AUC greater than 0.9. Therefore there is great diversity in terms of the overall chance of being a music (estimated by AUC value) across all video-ads.

Next we looked at the distribution of the area under the curve when considering the same three groups of video-ads as in the previous section, namely: video-ads that were never monetized, video-ads that were monetized at least once and video-ads with more than 5 exhibitions and that were monetized in at least 30% of its exhibitions. By presenting the distributions for these three groups of video-ads, our aim is to uncover evidence of some impact of the audio of the video-ads on their success in attracting user attention. As we can see in Figure 5.10(b), the distributions are similar, with monetized video-ads presenting a lower area under the curve. We applied a two sample Kolmogorov-Smirnov test [54] to check whether there is a significant difference between the area under the curve for the video-ads that were never monetized and the video-ads with more than 5 exhibitions and that were monetized at least 30% of the time. We found that, with 95% confidence, the null hypothesis that the two samples were drawn from the same distribution was rejected. This result shows that the use of music has an impact on

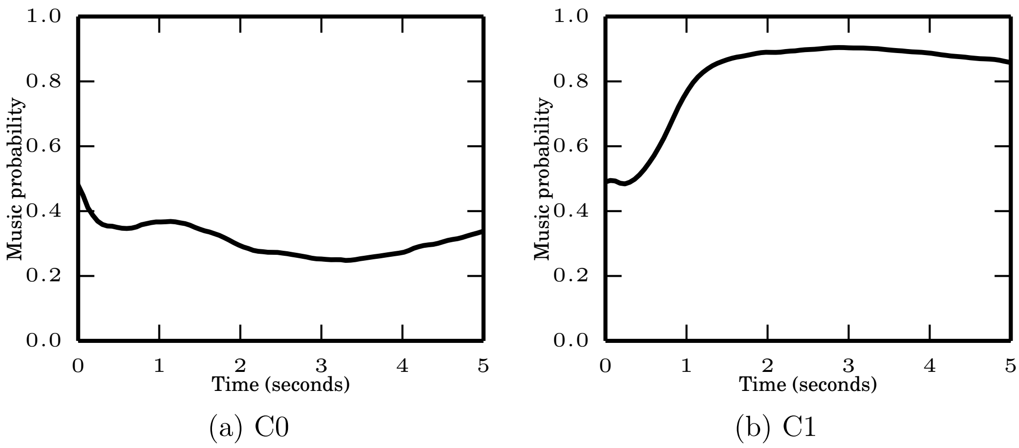

Figure 5.11: Cluster centroids of video-ad music probability over time.

monetization. Video-ads with a lower AUC (that is, video-ads with a lower probability of being music in its first 5 seconds) had a greater fraction of their exhibitions monetized. This result complements our finding in Section 5.2 which shows that the Music category was the one with the smallest fraction of exhibitions monetized, suggesting that using music in the beginning of video-ads may not help in attracting user attention[10]. This is an interesting result since YouTube is commonly used as a Jukebox and one may expect that video-ads with music may be less disturbing to users.

The AUC provides us with an aggregate measure of the probability of the video-ad being music in its first five seconds. However, this metric, by itself, does not allow us to understand how this probability evolves over time. With that aim in mind, we applied a clustering algorithm to group the time series of the music probability into patterns. The identified patterns represent profiles of the temporal evolution of the music probabilities. We made use of the K-Means algorithm [23]. This algorithm requires as an input the choice of the number k of clusters to be detected and the distance measure to be used. We used the Euclidian distance and we employed silhouette score to choose the number k of clusters. The silhouette score is a measure of how similar an instance is to its own cluster compared to other clusters. It is calculated by the following equation:

$$s(i) = \frac{b(i) - a(i)}{\max(a(i), b(i))}$$

where $a(i)$ is the average distance between instance i and all other instances within the same cluster and $b(i)$ is the smallest average distance of instance i to all instances in any other cluster, of which i is not a member. The score ranges from -1 to $+1$ and a high value indicates that the instance is well matched to its own cluster.

[10]This result may be interpreted with care, since our study is not sufficient to conclude that music may detract user attention. We can only conclude that the use of music does not help, but it may not be negative as well.

	C0	C1
Number of video-ads	1413	2805
% of monetized video-ads	66%	62%
% of monetized exhibitions	34.51%	34.27%

Table 5.1: Properties of each cluster of video-ad music probabilities over time.

We also performed a visual inspection of the cluster centroids and individual cluster members for different values of k. Based on the results, we identified k = 2 clusters. The cluster centroids are shown in Figure 5.11 and Table 5.1 summarizes some characteristics of each cluster.

Each centroid corresponds to an "average" probability curve for the video-ads in the cluster. Cluster C0 (Figure 5.11(a)) consists of video-ads with low probability of being music in the first five seconds. Cluster C1 (Figure 5.11(b)) in contrast consists of video-ads with a very high probability of being music. Out of the 4218 video-ads in our dataset, 1413 are in C0 and 2805 are in C1. Thus, video-ads with music background are more common. In order to contrast these two clusters in terms of the success of the video-ads in generating revenue, we looked into the number of video-ads that generated revenue at least once in each cluster. As presented in Table 5.1, 66% of the video-ads in C0 generated revenue at least once and 62% of the video-ads in C1 generated revenue. When looking at the video-ad exhibitions in each cluster, both clusters have also a similar percentage of monetized exhibitions. These results reinforce the discussion provided above that in general, using music in video-ads does not help to attract user attention.

In sum, most video-ads use music in its first 5 seconds. We found two profiles of video-ads, one with a low probability of having music in its first 5 seconds and one with a high probability of the audio being music. Our results also show that music does not seem to help to attract the attention of users.

5.5 Summary of Findings

Social media applications rely heavily on their audience to generate revenue. Content providers 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. 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 chapter, we took a new step towards building such understanding by providing an overview on monetization of video-ads on YouTube. In particular, we focused on three research questions: (1) How successful are video-ads in generating revenue? (2) How successful are channels in attracting revenue? and (3) How are multimedia properties of video-ads characterized? Are these properties related to monetization?

Towards tackling RQ1, we initially discussed the fraction of monetized exhibitions on our campus data. While this fraction may not reflect YouTube's global behavior, it offers 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 show that local campus traces may mitigate this issue. More importantly, we also discussed that shorter video-ads may have a higher chance of attracting user attention. Finally we also showed 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 video-ads to users. Shorter video-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.

On RQ2 we looked at content producers. These producers gain earnings from advertisements paired with their videos. As stated, YouTube will usually pay channels after every 1,000 monetized exhibitions. Even though our campus dataset has limited information on channels, our results are able to show that some channels generate revenue in most video-ad exhibitions paired with their contents. Such observation can be exploited by YouTube itself to find new partners [52]. 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 [52]. 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. One possible way to start this study is to use the channels of the video-ads as a proxy for the media marketers that publish advertisements on YouTube.

On RQ3 we studied multimedia properties of video-ads and the relation of those properties with monetization. We first characterized features extracted from the videos of the video-ads and then we explored the audio of the ads. Our results show that in general video-ads that more often generate monetization tend to have smaller numbers of blobs and faces and tend to be more diverse in terms of colors in the initial 5 seconds of content. However, the impact of these properties on the success of video-ads in generating revenue is very hard to measure. Our results also show that most video-ads have music in their first five seconds, although we found that music does not help in attracting user attention. We presented just an initial study of multimedia properties of video-ads. As future work, one may explore other features extracted from the audio and video content, correlating them with monetization.

Chapter 6

Conclusions and Future Work

Advertising is fundamental for the Web we know today. Every day billions of users access high quality content and services on the internet for free. Most of the time, those users are not even aware of the economical model that sustains the ecosystem. Also, the new generations of Internet users demand more efforts and care from service providers, since they are more worried about their data, their privacy and the type of content they are exposed to. Due to this new profile of users, it is important to understand the users, their expectations and the impact advertisements have on them, in order to keep the ecosystem healthy and economically sustainable.

In this thesis, we took a step towards such understanding of users, focusing on a specific type of advertisement that is becoming very popular on the Internet, the video-advertisements. We presented a study of video-ads on YouTube from two new perspectives. We started our study by investigating the users' perception of video-ads, bringing forth their role in the complex ecosystem of video-ads on YouTube. Then we provided an overview on monetization of video-ads, looking at the ads that were successful in generating revenue and also at the content creators that were able to profit from video-ads associated with their contents. Our work complements previous studies and provides a timely look of the YouTube ad ecosystem.

In order to provide a view of the video-ad ecosystem from the perspective of the users, we took an exploratory approach, employing survey and diary based research. The purpose of the survey was to collect demographic information of our participants, as well as their general opinion about the use of video-advertisements on YouTube. The diary was then used to collect individual experiences of participants when exposed to video-ads. Our results showed that, although most participants perceive video-ads in a negative way and would prefer to use the application without them, there were also participants more open to advertisements and that would watch video-ads they find interesting. We also found that the context of the user, as well as the content of the video-ad itself may have an impact on the users' decision to watch or skip the advertisement.

Our findings can be used by advertisers, content providers and even the users itself. Our results show that there is still room for advertisers to create interesting ad campaigns that are able to capture the attention of users. We uncovered several reasons for users to skip or watch video-ads. These reasons can serve as a guide for advertisers when creating new video-ads. The fact that users can still be interested in advertisements and the lack of knowledge from the majority of users about the underlying economic model that sustains the web can also be very enlightening to content providers. Content providers can use our results as motivation to bring more awareness to users about the need of advertisements to offer them quality content for free. Finally, our discussions about the context of the user can also benefit content providers and users. Content providers can invest in algorithms to predict the context of the users and based on that information choose the best time and the best ad to be displayed to the users. In the end, a better understanding of the whole ecosystem will benefit the users, since their needs, expectations and worries will be taken into consideration.

Next, in order to provide an overview on monetization of video-ads, we used a dataset of logs of HTTP requests originated from a university campus network. We first analyzed the fraction of video-ad exhibitions that were able to generate revenue. Then we shifted our attention to the content creators that had monetized video-ad exhibitions associated with their contents. Finally we extracted multimedia properties of video-ads, correlating those properties with the success of video-ads in attracting user attention. With our analysis we gave insights on the monetization of YouTube as a whole. We found that a considerable number of video-ads generated revenue, but the contribution of each one in particular was very small. We also found that a large number of content creators were able to profit from YouTube, although most of them had only a few video-ad exhibitions associated with their contents. When looking at the multimedia properties, we noticed that most video-ads present a small number of faces and blobs per frame and the use of music in the first five seconds is very common.

Since the Web and the users are always evolving, it is extremely important to understand the whole ecosystem of advertising in order to promote innovation, allowing services and advertisements to remain attractive to users. Our findings here can help to driven such innovation. We promoted discussions about the impact of each video-ad and content creator on YouTube as a whole and we also uncovered properties of video-ads that were successful in generating revenue. As far as we know, we were the first study to look into multimedia properties of video-ads. These properties can be applied to create new ad campaigns that will better explore the multimedia features in order to succeed. Most importantly, our study can motivate further studies of multimedia properties, that can lead to interesting and very practical results.

Our work can be extended in several directions. Our qualitative analysis can be lengthened to explore other settings and participant pools. Moreover, we only studied some factors that may impact user behavior, there are other factors to be studied, for instance, the impact of culture. Furthermore, our characterization of monetization of video-ads can also be extended. We presented a study focusing on the content creators and the video-ads, leaving behind the media marketers who also play a role in this ecosystem and should be studied. Also, our study of multimedia features was just introductory. A follow-up study would encompass an interaction with researchers in marketing in order to extract more meaningful features of the audio and video of the video-ads. To conclude, both our studies can be extended to consider other types of advertisements, as well as other platforms.

Bibliography

[1] Mariana Arantes, Flavio Figueiredo, and Jussara M Almeida. Understanding video-ad consumption on youtube: a measurement study on user behavior, popularity, and content properties. In *Proceedings of the 8th ACM Conference on Web Science*, pages 25–34. ACM, 2016.

[2] Mariana Arantes, Flavio Figueiredo, and Jussara M Almeida. Towards understanding the consumption of video-ads on youtube. *The Journal of Web Science*, 3(1), 2017.

[3] Mariana Arantes, Flavio Figueiredo, Jussara M Almeida, and Raquel O Prates. An investigation of user actions and experiences when exposed to youtube video ads. In *Proceedings of the 24th Brazilian Symposium on Multimedia and the Web*, pages 283–290. ACM, 2018.

[4] Camila Souza Araújo, Gabriel Magno, Wagner Meira, Virgilio Almeida, Pedro Hartung, and Danilo Doneda. Characterizing videos, audience and advertising in youtube channels for kids. In *International Conference on Social Informatics*, pages 341–359. Springer, 2017.

[5] Yoram Bachrach, Sofia Ceppi, Ian A Kash, Peter Key, and Mohammad Reza Khani. Mechanism design for mixed bidders. In *Proceedings of the 25th International Conference on World Wide Web*, pages 215–225. International World Wide Web Conferences Steering Committee, 2016.

[6] Nicola Barbieri, Fabrizio Silvestri, and Mounia Lalmas. Improving post-click user engagement on native ads via survival analysis. In *Proceedings of the 25th International Conference on World Wide Web*, pages 761–770. International World Wide Web Conferences Steering Committee, 2016.

[7] Paul Barford, Igor Canadi, Darja Krushevskaja, Qiang Ma, and S Muthukrishnan. Adscape: Harvesting and analyzing online display ads. In *Proceedings of the 23rd international conference on World wide web*, pages 597–608. ACM, 2014.

[8] Russell W Belk. Qualitative research in advertising. *Journal of Advertising*, 46(1):36–47, 2017.

[9] Maria Carpen Amarie, Ioannis Pefkianakis, and Henrik Lundgren. Mobile video ad caching on smartphones. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 57–61. ACM, 2014.

[10] Juan Miguel Carrascosa, Jakub Mikians, Ruben Cuevas, Vijay Erramilli, and Nikolaos Laoutaris. I always feel like somebody's watching me: measuring online behavioural advertising. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, page 13. ACM, 2015.

[11] Scott Carter and Jennifer Mankoff. When participants do the capturing: the role of media in diary studies. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 899–908. ACM, 2005.

[12] Jilin Chen, Eben M Haber, Ruogu Kang, Gary Hsieh, and Jalal Mahmud. Making use of derived personality: The case of social media ad targeting. In *Proceedings of the Ninth International AAAI Conference on Web and Social Media*, pages 51–60, 2015.

[13] Henriette Cramer. Effects of ad quality & content-relevance on perceived content quality. In *proceedings of the 33rd annual ACM conference on human factors in computing systems*, pages 2231–2234. ACM, 2015.

[14] John W Creswell. *Research design: Qualitative, quantitative, and mixed methods approaches.* Sage publications, 2013.

[15] Frank E Dardis, Mike Schmierbach, Brett Sherrick, Frank Waddell, Jose Aviles, Sushma Kumble, and Erica Bailey. Adver-where? comparing the effectiveness of banner ads and video ads in online video games. *Journal of Interactive Advertising*, 16(2):87–100, 2016.

[16] David and MacKay. *Information theory, inference and learning algorithms.* Cambridge university press, 2003.

[17] Joe DeBlasio, Saikat Guha, Geoffrey M Voelker, and Alex C Snoeren. Exploring the dynamics of search advertiser fraud. In *Proceedings of the 2017 Internet Measurement Conference*, pages 157–170. ACM, 2017.

[18] Robert H Ducoffe. Advertising value and advertising on the web. *Journal of advertising research*, 36(5):21–21, 1996.

[19] David Easley and Jon Kleinberg. *Networks, crowds, and markets: Reasoning about a highly connected world.* Cambridge University Press, 2010.

[20] Flavio Figueiredo, Jussara M. Almeida, Marcos André Gonçalves, and Fabrício Benevenuto. On the Dynamics of Social Media Popularity. *ACM Transactions on Internet Technology*, 14(4):1–23, December 2014.

[21] Alessandro Finamore, Marco Mellia, Michela Meo, Maurizio M Munafo, and Dario Rossi. Experiences of Internet Traffic Monitoring with Tstat. *Network, IEEE*, 25(3):8–14, 2011.

[22] Phillipa Gill, Vijay Erramilli, Augustin Chaintreau, Balachander Krishnamurthy, Konstantina Papagiannaki, and Pablo Rodriguez. Follow the Money: Understanding Economics of Online Aggregation and Advertising. In *Proceedings of the 2013 conference on Internet measurement conference*, 2013.

[23] Jiawei Han, Jian Pei, and Micheline Kamber. *Data mining: concepts and techniques*. Elsevier, 2011.

[24] Henning Hohnhold, Deirdre O'Brien, and Diane Tang. Focusing on the long-term: It's good for users and business. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1849–1858. ACM, 2015.

[25] Patrick Hummel and R Preston McAfee. When does improved targeting increase revenue? *ACM Transactions on Economics and Computation (TEAC)*, 5(1):4, 2016.

[26] Umar Iqbal, Zubair Shafiq, and Zhiyun Qian. The ad wars: retrospective measurement and analysis of anti-adblock filter lists. In *Proceedings of the 2017 Internet Measurement Conference*, pages 171–183. ACM, 2017.

[27] Helena C Kraemer. Kappa coefficient. *Wiley StatsRef: Statistics Reference Online*, 1982.

[28] S Shunmuga Krishnan and Ramesh K Sitaraman. Understanding the effectiveness of video ads: a measurement study. In *Proceedings of the 2013 conference on Internet measurement conference*, pages 149–162. ACM, 2013.

[29] Anísio Lacerda, Marco Cristo, Marcos André Gonçalves, Weiguo Fan, Nivio Ziviani, and Berthier Ribeiro-Neto. Learning to advertise. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 549–556. ACM, 2006.

[30] Mounia Lalmas, Janette Lehmann, Guy Shaked, Fabrizio Silvestri, and Gabriele Tolomei. Promoting positive post-click experience for in-stream yahoo gemini users. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1929–1938. ACM, 2015.

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

[32] Hao Li and Hui-Yi Lo. Do you recognize its brand? the effectiveness of online in-stream video advertisements. *Journal of Advertising*, 44(3):208–218, 2015.

[33] Matthew Malloy, Mark McNamara, Aaron Cahn, and Paul Barford. Ad blockers: Global prevalence and impact. In *Proceedings of the 2016 Internet Measurement Conference*, pages 119–125. ACM, 2016.

[34] Andrew McStay. A qualitative approach to understanding audience's perceptions of creativity in online advertising. *The Qualitative Report*, 15(1):37, 2010.

[35] Yan Mei, Weiliang Zhao, and Jian Yang. Maximizing the effectiveness of advertising campaigns on twitter. In *Big Data (BigData Congress), 2017 IEEE International Congress on*, pages 73–80. IEEE, 2017.

[36] Vahab Mirrokni and Hamid Nazerzadeh. Deals or no deals: Contract design for online advertising. In *Proceedings of the 26th International Conference on World Wide Web*, pages 7–14. International World Wide Web Conferences Steering Committee, 2017.

[37] Katie O'Donnell and Henriette Cramer. People's perceptions of personalized ads. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1293–1298. ACM, 2015.

[38] Karla Okada, Edleno S. de Moura, Marco Cristo, David Fernandes, Marcos André Gonçalves, and Klessius Berlt. Advertisement selection for online videos. In *Proceedings of the 18th Brazilian Symposium on Multimedia and the Web*, WebMedia '12, pages 367–374, New York, NY, USA, 2012. ACM.

[39] Panagiotis Papadopoulos, Pablo Rodriguez Rodriguez, Nicolas Kourtellis, and Nikolaos Laoutaris. If you are not paying for it, you are the product: How much do advertisers pay to reach you? In *Proceedings of the 2017 Internet Measurement Conference*, pages 142–156. ACM, 2017.

[40] Symeon Papadopoulos, Fotis Menemenis, Yiannis Kompatsiaris, and Ben Bratu. Lexical graphs for improved contextual ad recommendation. In *Proceedings of the 31th European Conference on IR Research on Advances in Information Retrieval*, ECIR '09, pages 216–227, Berlin, Heidelberg, 2009. Springer-Verlag.

[41] Jaimie Yejean Park, Yunkyu Sohn, and Sue Moon. Power of earned advertising on social network services: A case study of friend tagging on facebook. In *Proceedings of the Tenth International AAAI Conference on Web and Social Media*, pages 299–308, 2016.

[42] Enric Pujol, Oliver Hohlfeld, and Anja Feldmann. Annoyed users: Ads and ad-block usage in the wild. In *Proceedings of the 2015 Internet Measurement Conference*, pages 93–106. ACM, 2015.

[43] Fabian Schneider, Anja Feldmann, Balachander Krishnamurthy, and Walter Willinger. Understanding online social network usage from a network perspective. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement*, pages 35–48. ACM, 2009.

[44] Sergej Sizov and Sarah Piller. Economic value of web keyword advertising campaigns. In *Proceedings of the 8th ACM Conference on Web Science*, pages 117–126. ACM, 2016.

[45] Jim Sterne. *What makes people click: Advertising on the Web*. Que Corp., 1997.

[46] Shar VanBoskirk, Jim Nail, and Carlton Doty. Us digital marketing forecast, 2016 to 2021. *Forrester Research, Cambridge, MA*, 2017.

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

[48] Wall-Street-Journal. YouTube Alone Brings in 20% of Digital Video Ad Revenue, September 2014.

[49] Robert J Walls, Eric D Kilmer, Nathaniel Lageman, and Patrick D McDaniel. Measuring the impact and perception of acceptable advertisements. In *Proceedings of the 2015 Internet Measurement Conference*, pages 107–120. ACM, 2015.

[50] Pengyuan Wang, Yechao Liu, Marsha Meytlis, Han-Yun Tsao, Jian Yang, and Pei Huang. An efficient framework for online advertising effectiveness measurement and comparison. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 163–172. ACM, 2014.

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

[52] Mirjam Wattenhofer, Roger Wattenhofer, and Zack Zhu. The youtube social network. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*, 2012.

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

[54] R Wilcox. Kolmogorov–smirnov test. *Encyclopedia of biostatistics*, 2005.

[55] Guandong Xu, Zongda Wu, Guiling Li, and Enhong Chen. Improving contextual advertising matching by using wikipedia thesaurus knowledge. *Knowledge and Information Systems*, 43(3):599–631, 2015.

[56] Jaewon Yang and Jure Leskovec. Patterns of temporal variation in online media. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 177–186. ACM, 2011.

[57] Jinyoung Yeo, Sungchul Kim, Eunyee Koh, Seung-won Hwang, and Nedim Lipka. Predicting online purchase conversion for retargeting. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 591–600. ACM, 2017.

[58] Shuai Yuan, Jun Wang, and Xiaoxue Zhao. Real-time bidding for online advertising: Measurement and analysis. In *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*, ADKDD '13, pages 3:1–3:8, New York, NY, USA, 2013. ACM.

[59] Weinan Zhang, Shuai Yuan, and Jun Wang. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1077–1086. ACM, 2014.

[60] Ke Zhou, Miriam Redi, Andrew Haines, and Mounia Lalmas. Predicting pre-click quality for native advertisements. In *Proceedings of the 25th International Conference on World Wide Web*, pages 299–310. International World Wide Web Conferences Steering Committee, 2016.

Appendix A

An Exploratory Study

A.1 Survey Questions

1. What is your name?
2. What is your age?
    a) Under 18 years old
    b) 19-22 years old
    c) 23-27 years old
    d) 28-32 years old
    e) 33-39 years old
    f) 40-49 years old
    g) 50 years or older
3. What is your gender?
    a) Male
    b) Female
    c) Prefer not to answer
4. How often do you use YouTube?
    a) Very often (at least once a day)
    b) Often (few times per week)
    c) Ocasionally (few times per month)
    d) Rarely (few times per year)
    e) Never
5. Have you ever subscribed to a YouTube channel?

a) Yes
b) No

6. What is your opinion on the following statement: “YouTube would be better without video advertisements”.

   a) Strongly agree
   b) Agree
   c) Neither agree nor disagree
   d) Disagree
   e) Strongly disagree

7. What is your opinion on the following statement: “I would be willing to pay to use Youtube without advertisements".

   a) Strongly agree
   b) Agree
   c) Neither agree nor disagree
   d) Disagree
   e) Strongly disagree

8. Do you use any software to block advertisements?

   a) Yes
   b) No
   c) I don’t know this type of software

A.2 Diary Questions

1. What is your name?

2. Device

   a) Computer
   b) Smart phone
   c) Tablet

   d) Video Game
   e) Smart TV
   f) Other

3. Describe in a few words the content (YouTube video) you were watching.
4. Why were you watching this content?
5. Did you skip the advertisement?
   a) Yes
   b) No
6. Describe in a few words why you skipped or not the advertisement.
7. Do you know what was the advertisement about?
   a) Yes
   b) No
8. If you answered 'Yes' to the previous question, please tell us what the advertisement was about.
9. Do you think the advertisement was related to your personal interests?
   a) Yes
   b) No
   c) I don't know

A.3 Recruitment of Participants

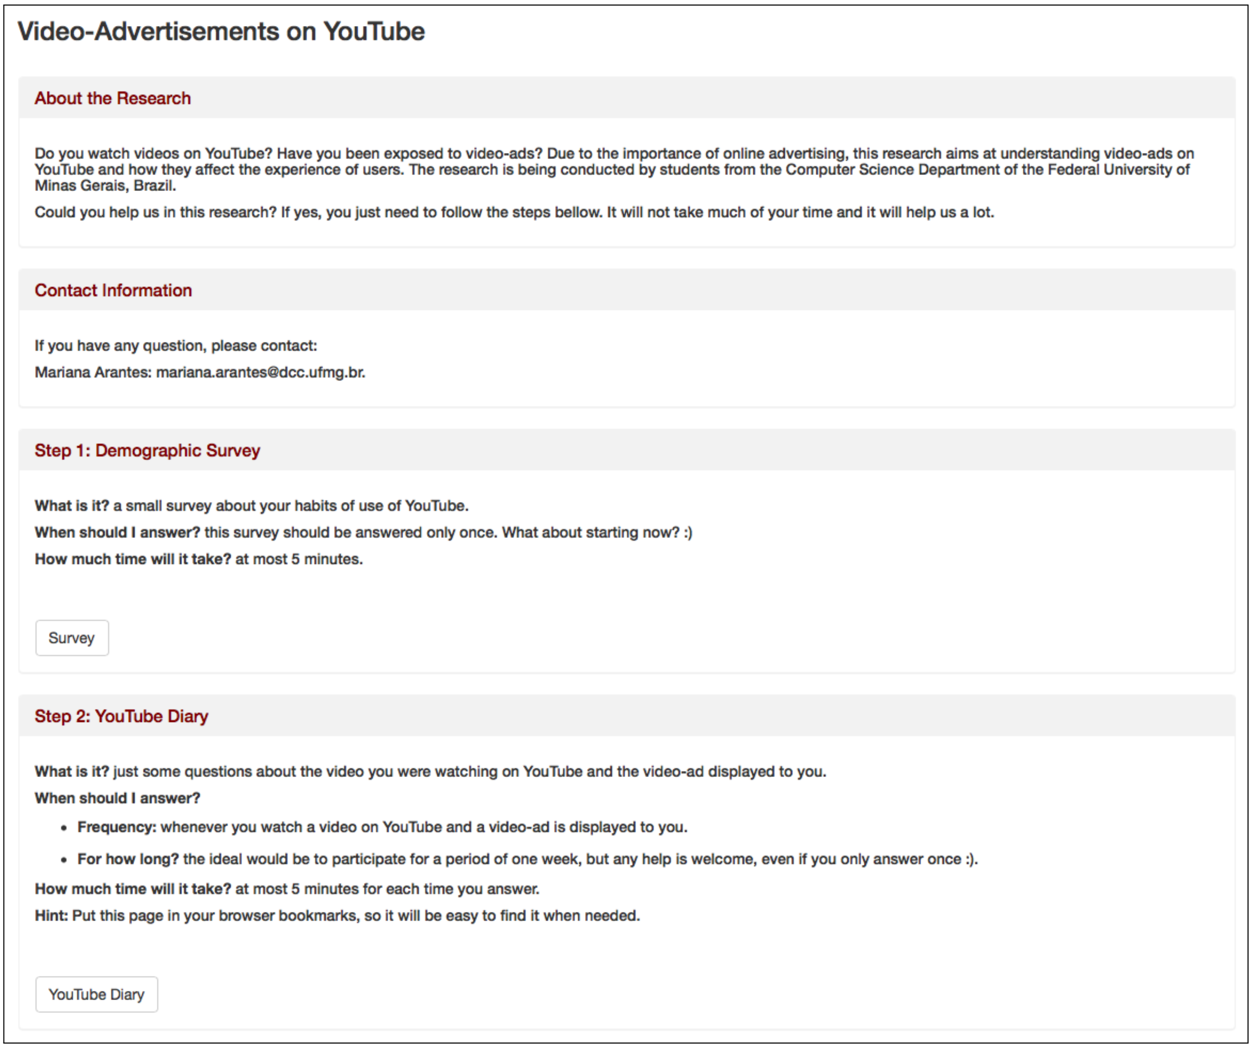

Video-Advertisements on YouTube

About the Research

Do you watch videos on YouTube? Have you been exposed to video-ads? Due to the importance of online advertising, this research aims at understanding video-ads on YouTube and how they affect the experience of users. The research is being conducted by students from the Computer Science Department of the Federal University of Minas Gerais, Brazil.

Could you help us in this research? If yes, you just need to follow the steps bellow. It will not take much of your time and it will help us a lot.

Contact Information

If you have any question, please contact:

Mariana Arantes: mariana.arantes@dcc.ufmg.br.

Step 1: Demographic Survey

What is it? a small survey about your habits of use of YouTube.

When should I answer? this survey should be answered only once. What about starting now? :)

How much time will it take? at most 5 minutes.

Survey

Step 2: YouTube Diary

What is it? just some questions about the video you were watching on YouTube and the video-ad displayed to you.

When should I answer?

- **Frequency:** whenever you watch a video on YouTube and a video-ad is displayed to you.
- **For how long?** the ideal would be to participate for a period of one week, but any help is welcome, even if you only answer once :).

How much time will it take? at most 5 minutes for each time you answer.

Hint: Put this page in your browser bookmarks, so it will be easy to find it when needed.

YouTube Diary

Figure A.1: Screenshot of web page used in the recruitment process.