



HAL
open science

Measuring Audience Retention in YouTube

Eitan Altman, Tania Jimenez

► **To cite this version:**

Eitan Altman, Tania Jimenez. Measuring Audience Retention in YouTube. VALUETOOLS 2019 - 12th EAI International Conference on Performance Evaluation Methodologies and Tools, Mar 2019, Palma, Spain. pp.79-85, 10.1145/3306309.3306322 . hal-02053300

HAL Id: hal-02053300

<https://inria.hal.science/hal-02053300v1>

Submitted on 1 Mar 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Measuring Audience Retention in YouTube

Eitan Altman
INRIA Sophia-Antipolis, Univ Cote d'Azur;
LIA, Avignon University and LINCS, Paris
France

Eitan.Altman@inria.fr

Tania Jiménez
Laboratoire Informatique d'Avignon (LIA),
Avignon University (AU)
France

tania.jimenez@univ-avignon.fr

ABSTRACT

There exist many aspects involved in a video turning viral on YouTube. These include properties of the video such as the attractiveness of its title and thumbnail, the recommendation policy of YouTube, marketing and advertising policies and the influence that the video's creator or owner has in social networks. In this work, we study audience retention measures provided by YouTube to video creators which may provide valuable information for improving the videos and for better understanding the viewers' potential interests in them. We then study the question of when is a video too long and can gain from being shortened. We examine consistency between several existing audience retention measures. We end in a proposal for a new audience retention measure and identify its advantages.

1. INTRODUCTION

In beginning of 2017, the average number YouTube videos from mobiles exceeded 1,000,000,000 views and Google's annual revenue generated from YouTube reached \$4,000,000,000.¹ This has generated a huge research interest in understanding the video propagation in the Internet and how this propagation can be accelerated.

The propagation of YouTube videos can be observed and studied at various levels:

- We have proposed in [9] an epidemic approach for the analysis of the propagation of a single video based on the viewcount process of the video. We studied the viewcount of 80,000,000 videos and showed that over 90% of them can be classified within one of the six epidemic models with a mean square error below 5%. Note that the viewcount information over a video is quite often available (it becomes unavailable if the creator does not wish it to be, or if there are copyright issues).
- The level of Internet Memes. The word Meme was coined in page 192 of [4] as a concept for describing evolutionary principles and for explaining the spread of ideas and cultural phenomena. Internet memes are a subset of this general meme concept specific to the culture and environment of the Internet. Wikipedia

¹See Facts on YouTube <https://fortunelords.com/youtube-statistics/>

cites from [12] Dawkins who characterized an *Internet meme* as being a "meme deliberately altered by human creativity, distinguished from biological genes and Dawkins' pre-Internet concept of a meme". In [2] we studied three videos at the level of both Internet meme as well as of the viewcount processes related to each video. To get information about the propagation of Internet memes we used Google trends.

- Marketing or advertisement policies play an important role in accelerating the propagation of a video. The study of this propagation aspect includes game theoretical approaches that answer the question of which content to advertize, how much and/or how long to invest in advertizing [1], when to post or to advertize a content (time of the day or of the week) [3].
- The protocols used in the recommendation system of YouTube and/or of other social networks over which videos are posted have of course an impact on the propagation of YouTube videos. In [7] we carried out an empirical analysis of some stability properties of the stochastic processes that describes the sequence of videos that are watched by a YouTube user who moves from watching one video after another according to respective recommended videos.

Analysis of audience retention has many applications. It can be used to assist with the decision of where on a video, or on what video to place advertisements so as to maximize their visibility. Such a use in the context of audience of television programs has been patented in [13].

Audience retention information can also be used in the context of caching [14] as it can help identifying what parts of a video would benefit most from caching (as we receive information on the percentage of viewers that watch the video in each point in time).

The audience retention provides a measure for the potential of a video to become epidemic. It can thus be used to select videos for uploading. Companies who produce videos further use it as a feedback on videos they create and then when possible, the video can be modified by using the insight from data on public retention, see [5].

A video version of our work can be found at <https://www.youtube.com/watch?v=H3EWUt1Ot5A>

2. EGO NETWORK APPROACH

Models for analysing popularity of videos and the virality of content have often been studied and analysed based

on statistical inference from large sets of videos. The larger the set is, the more precise are the conclusions that can be drawn on a randomly chosen "representative" of the group. Yet popularity and virality can also be studied within the framework of Ego networks [10, 11] in which the object of the study may be a given small set of videos or of video channels, which need not be representative of all videos, nor even of all videos within some communities. This small set may contain a single interesting element which we may wish to study. In the video "Why study Ego networks" [11], the author presents various examples where one is interested in the study of a small subnetwork centered around an individual called EGO, rather than in the study of the entire network. User retention is a particular property of a video that is interesting to study in the Ego network framework. Such a study enables video creators to increase the popularity of a particular video or of a set of videos. Some measurements of retention are of course needed, but in the case of YouTube they are indeed made available to the video creator by YouTube.

3. AUDIENCE RETENTION AND ENGAGEMENT GRAPHS

A key indicator for audience engagement is the average retention, this is the time spent by a viewer watching the video averaged over all viewers. The data on audience retention of a video is not publicly available and is provided by YouTube only to the video creator. We have therefore spent an effort in creating some 1000 videos since 2012 which have received more than a million of views. More precisely, we shall use in this work videos from the following YouTube channels:

1. altman2208, 129 videos, 76736 views, 102 subscribers,
2. eitan meir altman, 410 videos, 869000 views, 970 subscribers
3. eitan m. altman, 189 videos, 69534 views, 170 subscribers
4. musicandshows eitan, 274 videos, 26899 views, 30 subscribers

Note: we did not attempt to study the generic retention in a "typical" video. Neither did we restrict the study to a specific community. We used the Ego network approach since we are interested in understanding how can a video creator use retention data made available to him by YouTube in order to increase the popularity of his video channels.

Inspired by [5], we decompose the engagement graph into three disjoint parts: the beginning, the core and the end of the video. We use the corresponding terms "nose", "body" and "tail" from [5]. As an example consider Figure 1. According to [5], there is a "large drop in engagement at the very beginning of the video, a steady drop over the middle section, and then another inflection point at the end." In [5] the nose is defined as the first 2% of the video.

In a retention graph, such as in Figure 1, the horizontal axis represents the time t and the vertical axis $x(t)$ represents the percentage of persons that watched the video at time t .

Let X be the maximum fraction that watched the video at any time.

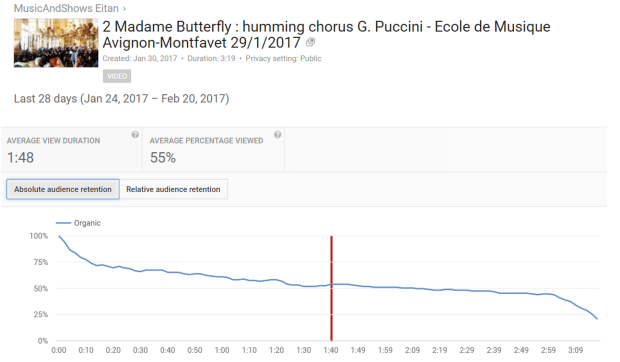


Figure 1: A typical behavior of the engagement graph

REMARK 1. YouTube allows for X to be greater than 100%. This may occur if after viewing the video viewers drag the cursor to the beginning and watches the video (or parts of it) again. This is well illustrated in Figure 2.



Figure 2: The case of X larger than 100%.

The average fraction of persons watching the video is given by

$$\bar{X} = \frac{1}{T} \int_0^T x(t) dt$$

where T is the duration of the video. But \bar{X} is also the ratio between the duration of the video watched, averaged over all those who watched the video, and the video length. It is called in YouTube the "average percentage viewed" (APV).

The average duration \bar{T} that the video is watched is given by

$$\bar{T} = \int_0^T x(t) dt = T\bar{X}. \quad (1)$$

It is called in YouTube the "Average View Duration" (AVD).

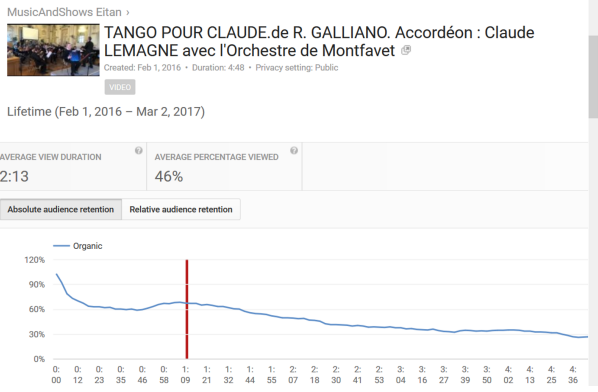


Figure 3: A non monotone nose

Figure 2 and Figure 3 are two examples of non decreasing retention curves. In the first figure, the percentage of viewers reaches 170%. YouTube explain this phenomenon by the fact that many viewers watched the video several times consecutively by dragging the mouse to the beginning. The non-monotonicity then could be due to dragging the mouse not all the way to the beginning, which results in points further away from the beginning to be viewed more than the beginning of the video. Note also that in Figure 2, the average video watching time \bar{T} exceeds the video duration.

Figure 3 is the curve corresponding to one of 14 videos that we filmed in a concert. The U form of the nose is common to all of them and is due to the fact that each of them starts with around one minute of introduction in which the conductor explains briefly the piece that is about to be played. A fraction of the viewers were not interested in the introduction but were interested in the piece. Thus the U shape of the nose. This video has an APV of 46%, its duration is 4'48" and 304 views. We filmed the concert at the end of January 2016. The retention curve in Figure 1 corresponds to one of 14 videos we took a year later with the same orchestra and at the same place. This time we did not include the introductions in the videos and the U type noses have disappeared as can be seen in the curve.

4. IMPROVING VIDEO RETENTION

We begin by discussing the shortening of videos. We shall provide below conditions under which shortening the video would result in an increase in its average fraction of viewers.

As an example, we reduced a video of 5'51" on the French song "sur le pont d'Avignon" (see Figure 7) to 2'33" (see Figure 8). As a result, we observed an increase in the averaged percentage viewed from 16% to 42%. The average viewed duration increased by 2.625 while the whole video duration decreased by 2.294. Thus not only did the average viewed time increase, it increased by a factor which is lower bounded by the amount by which we shortened the video.

We shall say that a video is causal if the corresponding $x(t)$ is causal, i.e., it is only a function of events that occurred before t . By truncating a causal video defined on $[0, T]$ to $[0, S]$ for some $S < T$, the retention curve of the shorter video is the same as the initial one restricted to $[0, S]$.

THEOREM 2. Assume that we shorten a causal video 1 to obtain a new video 2.

(i) We have

$$\bar{X}_2 = \frac{1}{T_2} \bar{T}_2 = \frac{1}{T_2} \bar{T}_1 \Delta \leq \frac{T_1}{T_2} \bar{X}_1 \quad (2)$$

where $\Delta \leq 1$ is given by

$$\Delta = \frac{\int_0^{T_2} x(t) dt}{\int_0^{T_1} x(t) dt}$$

The inequality in (2) becomes an equality and $\Delta = 1$ if $x(t) = 0$ for all $T_2 < t < T_1$.

(ii) Assume that x_1 is strictly positive and strictly decreasing. Then

$$\bar{T}_1 > \bar{T}_2$$

and

$$\bar{X}_1 < \bar{X}_2. \quad (3)$$

Proof. (i) Follows from direct computations. (ii) The first part follows from

$$\frac{d\bar{T}}{dT} = \frac{d \int_0^T x(t) dt}{dT} = x(T) > 0$$

The second statement follows from

$$\begin{aligned} \frac{d\bar{X}}{dT} &= \frac{dT^{-1} \int_0^T x(t) dt}{dT} = T^{-1}(x(T) - \bar{X}(T)) \\ &= \frac{1}{T^2} \int_0^T (x(T) - x(t)) dt < 0 \end{aligned} \quad (4)$$

The first part of the theorem identified the best possible retention curve $x(t)$ under which it is worthwhile shortening a video from T_1 to T_2 . This was done by upper bounding Δ . If $x(t) = 0$ for $T_2 < t < T_1$ then the AVD (1) is the same for the original and the shortened video while the corresponding duration was shortened by a factor of T_2/T_1 .

Similarly, the performance of the worse possible retention curve for shortening a video can be obtained by lower bounding Δ .

LEMMA 3. Assume that we shorten a causal video 1 to obtain a new video 2. Assume that x_1 is strictly positive and is decreasing. Then $\Delta \geq \frac{T_2}{T_1}$.

Proof.

$$\begin{aligned} \Delta &\geq \frac{\int_0^{T_2} x(T_2) dt}{\int_0^{T_2} x(T_2) dt + \int_{T_2}^{T_1} x(t) dt} \\ &\geq \frac{\int_0^{T_2} x(T_2) dt}{\int_0^{T_2} x(T_2) dt + \int_{T_2}^{T_1} x(T_2) dt} = \frac{T_2}{T_1} \end{aligned}$$

Note that if $x(t)$ is strictly decreasing then the last inequality is strict. This is thus an alternative way to obtain (3). We shall next identify a retention that approaches this bound.

LEMMA 4. Consider a causal video of duration T_1 which we wish to shorten to a duration T_2 . Assume that $x(0) = 1$.

Assume that the retention curve has two phases such that for some $0 \leq \epsilon < 1$: (i) it has a nose of duration $\epsilon T_1 < T_2$, and (ii) it has a flat body, i.e. $x(t) = x(T_1)$ is constant over all $t \geq \epsilon T$. Then

$$|\bar{X}_1 - \bar{X}_2| \leq \epsilon \log(T_1/T_2) \quad (5)$$

Proof. It follows from (4) that

$$\frac{d\bar{X}}{dT} = \frac{\int_0^T (x(T) - x(t))dt}{T^2} = \frac{\int_0^{\epsilon T} (x(T) - x(t))dt}{T^2}$$

We conclude that

$$\left| \frac{d\bar{X}}{dT} \right| \leq \frac{\epsilon}{T}$$

Multiplying by dT both sides and integrating between T_2 and T_1 we get (5). ■

The take away message from Theorem 2 is that shortening a video from duration T_1 to duration T_2 may improve its APV (Absolute Percentage Viewed) by a factor of up to T_1/T_2 . The smaller $x(t)$ is in the interval $[T_2, T_1]$ the more we approach this improvement. Thus shortening videos is more appealing when the percentage of those that saw the video till the end is small. On the contrary, there is little gain in shortening a video if a large part of its body has a constant retention.

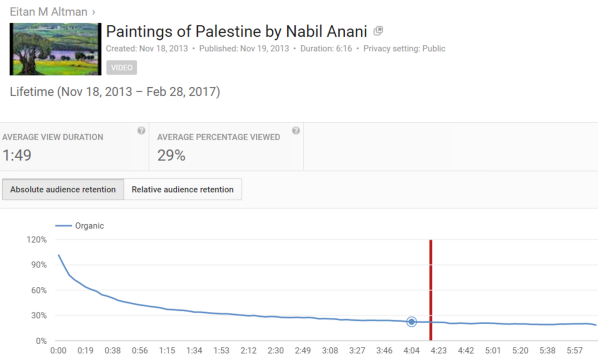


Figure 4: Relative audience retention

We used this intuition on the video of a Palestinian painter, Nabil Anani (<https://youtu.be/ZLHFWhwYqzI>), which we had prepared in the form of a slide show. Its APV is 29, its duration is 6'16" and it has 1516 views. At time 4'20", my piano music which accompanies the show sounds as if this were the end. We were therefore suggested recently to shorten the video (which is possible on YouTube). However, its retention curve 4 shows that 22% of its viewers saw the video till its end. In view of the above take away message, we decided not to do any change in the video.

Equation (3) states that under quite standard conditions, shortening a video will increase its APV index. Thus in terms of the APV, shorter videos have an advantage over longer ones. For reasons that will be stated later this is not a desired feature. In the next section we shall describe both YouTube's as well as our own proposed solution to this problem.

5. OTHER RETENTION MEASURES

To fairly compare the retention of videos of different lengths, YouTube introduced a measure called the "relative audience retention" (RAR). It consists in normalizing each point $x(t)$ of the original video by the value $Ave(t)$ obtained by averaging $x_i(t)$ over all similar YouTube videos $i \in I$ which have the same duration T as the original video. At each point in time the relative retention can take a value between 1 and 5 where the integers 1-5 correspond respectively to the categories "Low, below average, average, above average, High". These categories remind of the MOS measure for the quality of experience as standardized by the ITU.

In Figure 5 we see the relative retention curve corresponding to the absolute audience retention of Figure 4. Recall that it was suggested to end that video at time 4'20" and that we decided not to follow the suggestion since the benefit would have been small. The RAR seems definitely to agree with our decision and as we see in the Figure, it is precisely at the time that was suggested for ending the video that the rating moves steadily from 3 (average) to 4 (above average).

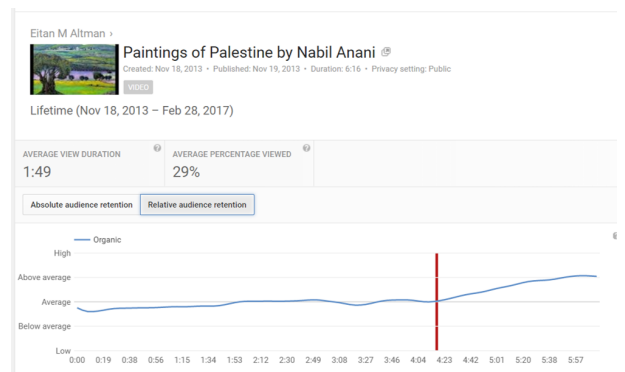


Figure 5: Relative audience retention

In Figure 6 we see the relative retention curve corresponding to the absolute audience retention of Figure 3. We see the same phenomenon again: there are above 30% of the audience who watch the video till the end, and in parallel we see that the RAR has a long increase in its score which gradually improves from 2.6 to 4 (it then decreases again during the tail part).




Figure 6: Relative audience retention

In contrast, consider the retention curve of the French song “Sur le pont d’Avignon” in Figure 7. Its duration is 5’51”, and it has an APV of 16%. We observe that a fraction of less than 5% of viewers saw it till the end. We thus decided to create a shorter version. The new short version is 2’33” long and its APV has increased to 40% or more (details are given below as well as the relative retention curve).

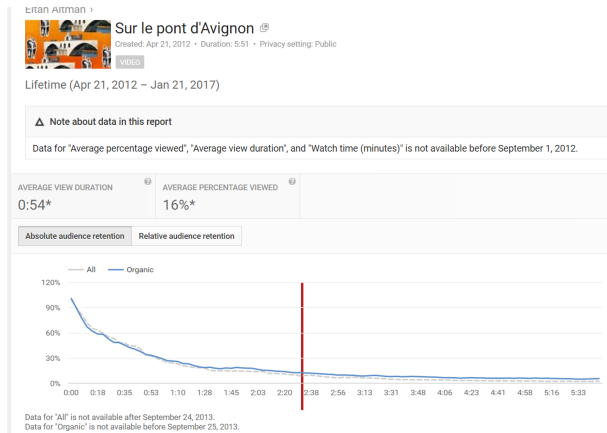


Figure 7: Refrelative audience retention

6. A NEW LOYALTY MEASURE OF RETENTION

In the previous Section we have shown that videos whose *fraction of the audience that remain till the end of the video is large* benefit little from shortening. Moreover, some evidence from our videos showed that they get better RAR ranking.

This motivates us to introduce a new audience retention measure which is the **fraction of the audience that remain till the end of the body of the video**. We call this the **loyal percentage**. Here are other advantages to judge video’s quality using this index.

1. When a person creates a video he often intends it to be viewed from the beginning to the end. Often the punchline comes in the end.
2. If a large part of the audience leaves before the end then they will probably not be interested in the next video by the same author.
3. In the context of real time broadcasts, if a person leaves before the end then she is more likely to switch to another program broadcasted at the same time and the chances to come back to watch the original channel are smaller.
4. In the context of YouTube, the video creator can embed at the 20 last seconds of the video personal recommendations to other videos or playlists. This option is called ”End Screen”. If a person leaves before the end of the video then he does not have access to these recommendations. Hence again the importance of this performance measure.



Figure 8: Absolute audience retention for the shortened video of “Sur le Pont d’Avignon” from Jan 5 till Feb 20, 2017.

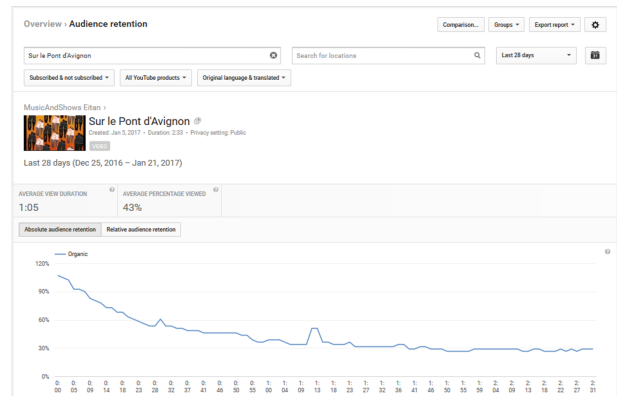


Figure 9: Absolute audience retention for the shortened video of “Sur le Pont d’Avignon” from Jan 5 till Feb 21 2017.



Figure 10: Absolute audience retention for the shortened video of “Sur le Pont d’Avignon” from Jan 5 till March 2nd 2017.

We finally present an example that shows that the APV and RAR may not be consistent. We say that two indexes are consistent if a better performance according to one of them results in a better performance with respect to the other.

Figure 8-10 correspond to the shortened video of the song “Sur le Pont d’Avignon” which we had created on January 5th, 2015. The duration of the video is 2’33”. We checked the AVP obtained by the same video in three different periods: Jan 5 till Feb 20, Jan 5 till Feb 21 and Jan 5 till March 2, 2017. The corresponding RAR are 42%, 43% and 40% (and the number of viewers increased from 61 to 70).

Figures 11-13 are the corresponding RAR ranking.

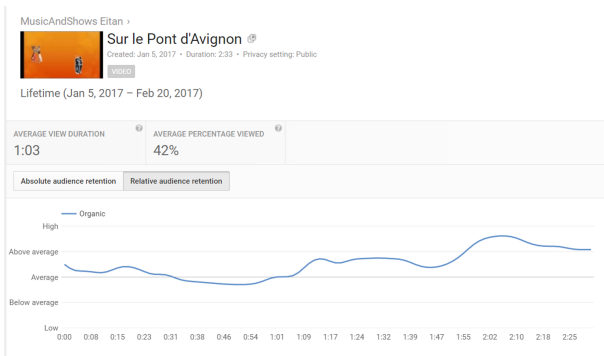


Figure 11: Relative audience retention for the shortened video of “Sur le Pont d’Avignon” from the period Jan 5 till Feb 20, 2017.

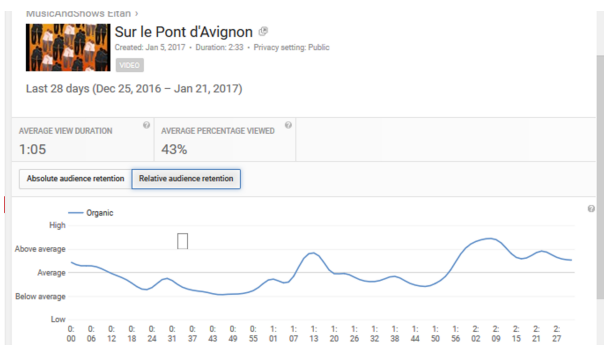


Figure 12: Relative audience retention for the shortened video of “Sur le Pont d’Avignon” from Jan 5 till Feb 21 2017.

While all three versions have almost identical APV performance (with the one of Feb 21 having the highest APV of 43% and that of March 2 the lowest at 40%), the relative retention ranking has quite a surprising behavior.

The RAR ranking of the period ending at March 2 (which has the lowest APV among the three) is almost uniformly higher in its relative ranking by 1 point with respect to that of the period ending at 21 January (which has the highest APV). The RAR ranking of the period ending at January 20 is in between the two others.

We thus have in this example a complete inversion between the rankings of the RAR and that of the APV. More-




Figure 13: Relative audience retention for the shortened video of “Sur le Pont d’Avignon” from Jan 5 till March 2nd 2017.

over, very small variations between the Absolute retention curves result in high variations between the RAR curves.

Now if we compare the APV with our loyalty measure, we see that the order of preference is the same under both rankings, and furthermore the percentage difference in the performance of the video in different periods is close under the two performance measures.

We conclude that in this example there seems to be consistency between the APV and the loyalty measure where as there is no consistency between these measures and the RAR index. Note that these conclusions were only obtained for a single example and may not hold in other ones. Further experimentation would be needed on a larger set of videos in order to understand the relation between the different performance measures.

7. PERSPECTIVES AND FUTURE WORK

In this paper, we have studied the relation between various measures of audience retention both existing ones as well as a new proposed measure, and have shown how such measures can be used in order to help deciding when a video can profit in being shortened. The measures include averaged ones in the form of a single number, such as the APV, the AVD, and the loyalty measure, as well as dynamic measures such as the retention curve or the relative audience retention.

Shortening videos

We have provided assumptions under which shortening a video seems beneficial or not. Of course there may encounter other user behavior that would provide different conclusions under different assumptions. As an example, The causality definition does not consider the effect that the knowledge of video duration has on the audience, i.e., if someone knows that there is still 1 min of video instead of 1 hour, this knowledge will most probably influence her/his interest to watch further or not. It would be interesting to study the influence of the information on the remaining video duration on the abandonment process of viewers.

Where as this work along with previous work cited here provide some understanding of when shortening of a video has a potential to improve the retention, a more systematic supervised learning may be necessary if one wishes to come up with automatic decision rules. A statistical analysis would then be required to evaluate the false positive and

false negative probabilities induced by learning rules.

In and Beyond the Ego network framework

In this work, we adopted the Ego network approach of studying a restricted small number of videos rather than a random video. This allows one to predict virality of a video at a very early stage. A high degree of audience retention would indicate that the video has a potential of growing viral even if the number of viewers of the video or the rate of views of the video is small. The relation between number of views, the rate of new views and retention measures is the subject of future work.

Even within the Ego network framework, many empirical results over retention measures have been obtained by using observations from a couple of videos to make an argument. While this approach allows us to get some insight on retention (e.g. discover phenomena such as getting non monotone retention curve), a statistical study can further provide more quantitative understanding of retention.

Other directions

There are several ways to pursue the research on audience retention. We plan to study the relation between the audience retention and network operation characteristics such as losses, throughputs and delay. These can be inserted to videos and their impact on the audience retention can then be measured using the tools that were discussed in this paper. This can be used to optimize the network resource allocation, and in particular the video resolution, caching decisions as well as radio channel characteristics (modulation and error correction).

Since the relative audience retention is obtained by a normalising procedure over a large number of videos, this opens the door for a statistical analysis that may take into account the video category (eg sport, music, travel etc...). This, as well as the analysis of the structure of a typical retention curve are properties that go beyond the Ego network framework.

Yet using this framework is of interest not only because video creators have access to the retention curve, but also because it allowed us to obtain properties that hold for any causal video, such as the impact of making the video shorter.

8. REFERENCES

- [1] Eitan Altman, "A stochastic game approach for competition over popularity in social networks", *Dynamic Games and Applications*, special Issue on stochastic games, *Dynamic Games and Applications* 3, 2 (2013) 313-323
- [2] Eitan Altman, Yonathan Portilla, "Social Networks: A Cradle of Globalized Culture in the Mediterranean Region." *International Conference on Advances in Social Networks Analysis and Mining (ICASNAM 2015)*, Jan 2015, Hammamet, Tunisia. 10.2139/ssrn.2518683 See related video (French) at https://www.youtube.com/watch?v=2H_HQwMK5pg
- [3] Eitan Altman, Nahum Shimkin, "The Ordered Timeline Game: Strategic Posting Times Over a Temporally Ordered Shared Medium", *Dynamic Games and Applications*, Springer Verlag, 2015, pp.1-25.
- [4] Dawkins, Richard (1989), "The Selfish Gene", Oxford University Press, ISBN 0-19-286092-5.
- [5] Ezra Fishman and Alyce Currier, *Understanding Audience Retention*, WISTIA, <https://wistia.com/library/understanding-audience-retention>
- [6] David Lee King, "Managing your library's social media channels", *Library Technology Reports*, Chapter 5, alatechsource.org January 2015.
- [7] Yonathan Portilla, Alexandre Reiffers, Eitan Altman, Rachid El-Azouzi, "A Study of YouTube recommendation graph based on measurements and stochastic tools", 3rd International Workshop on Big Data and Social Networking Management and Security, 2015.
- [8] Raul Rodriguez-Fernandez, Victoria Tur-Vines and Kiko Mora Contreras, "The tube on YouTube: TV series, Media strategies and user tactics in a transmedia environment", *International Journal of Communications* 10(2016), pages 1991-2013.
- [9] Cedric Richier, Eitan Altman, Rachid El-Azouzi, Tania Jimenez, Georges Linares and Yonathan Portilla, "Bio-Inspired Models for Characterizing YouTube Viewcount", 2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM).
- [10] Valerio Arnaboldi, Marco Conti, Andrea Passarella and Fabio Pezzoni; *Analysis of Ego Network Structure in Online Social Networks*, 2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing
- [11] James Cook, <https://www.youtube.com/watch?v=SM08TsCldWI>
- [12] Olivia Solon, "Richard Dawkins on the internet's hijacking of the word 'meme'". *Wired UK*. July 9, 2013. Available on <http://www.webcitation.org/6HzDGE9Go>.
- [13] M. Vinson, B Goerlich, M. Loper, M. Martin and A. Yazdani, "System and method for measuring television audience engagement", Patent US 8904419 B2
- [14] Lorenzo Maggi, Lazaros Gkatzikis, Georgios Paschos and Jeremie Leguay, "Adapting Caching to Audience Retention Rate: Which Video Chunk to Store?" [arXiv:1512.03274](https://arxiv.org/abs/1512.03274) 2017.