

The influence of visual salience on video consumption behavior: A survival analysis approach

Rafael Huber
University of Basel
Missionsstrasse 60
4055 Basel, Switzerland
+41 61 267 35 31
huber.rafael@gmail.com

Benjamin Scheibehenne
University of Basel
Missionsstrasse 60
4055 Basel, Switzerland
+41 61 267 35 31
benjamin.scheibehenne
@unibas.ch

Alexandre Chapiro
ETH Zürich
Universitätstrasse 6
8003, Zürich, Switzerland
alexandre.chapiro
@disneyresearch.com

Seth Frey
Disney Research Zürich
Stampfenbachstrasse 48
8006 Zürich, Switzerland
+41 44 632 7360
seth@disneyresearch.com

Robert W. Sumner
Disney Research Zürich
Stampfenbachstrasse 48
8006 Zürich, Switzerland
+41 44 632 7360
sumner@disneyresearch.com

ABSTRACT

In an increasingly competitive media environment, producers of online content need analytics that can predict the success of a video. In recent years the field of visual computation has produced a variety of mathematical models that quantify an image's salience, that is, its potential to capture attention. To test how a video's content might predict its success, we applied the standard saliency model of Itti, Koch, and Niebur [10] to more than 1000 video clips that were broadcast on a large video streaming website. We also obtained fine-grained data on the viewership of these clips. Based on a survival analysis, we find that that people prefer more salient videos. The results were robust towards the inclusion of other predictors such as the genre of the video, but not to video length, which remains correlated with salience even after comparing videos only within show and genre. Our analyses suggest that visual salience provides an objective and easy-to-compute supplement to previously suggested predictors of video consumption behavior.

Categories and Subject Descriptors

H.1.2 [Information Systems]: User/Machine Systems – *human factors, human information processing*. I.2.10 [Artificial Intelligence]: Vision and Scene Understanding – *intensity, color, photometry, and thresholding*.

General Terms

Human Factors.

Keywords

salience, visual salience, survival analysis, predictive analytics, video consumption behavior

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Conference '10, Month 1–2, 2010, City, State, Country.
Copyright 2010 ACM 1-58113-000-0/00/0010 ...\$15.00.

1. INTRODUCTION

Every year at the beginning of February the two most successful football teams of the year compete at the Super Bowl. But before, during, and after, another competition rages: that for the best commercial. Several of the world's most influential brands try to produce the most attractive advertisements in order to promote their company with a maximal impact. These commercials are watched by more than 100 million people in front of televisions alone. Given the huge investment in video material on both television and the internet, it is important to understand and predict what drives the success of video clips. The need to predict the success of video material also arises elsewhere, as in cinema [5]. To improve prediction accuracy, one promising approach is to examine objective properties of video content for hints about its success in front of an audience. Starting with basic psychological theories of visual perception, we test whether video consumption statistics can be predicted with quantitative measures of visual salience, as defined in terms of the levels and variability in the brightness, orientation, and color of video frames.

1.1 Predicting the success of videos

Past research has proposed several different approaches to improving prediction accuracy for the success of commercial video material [18]. When predicting the success of movies at the box office, many models rely on quantitative data that is available before or soon after the initial theatrical release. Here, typical predictors include ratings from critical reviews, genre (e.g. action movies might be more successful at the box office as compared to documentaries), stars appearing in the movie, competition for production money, the budget, whether or not a movie is a sequel, the number of screens on which a movie was shown, or the release date, to name but a few [1, 3, 5-7, 12, 17, 19]. Recently, researchers have also tried to predict movie success based on social media content [2] and Wikipedia activity [13]. Including these predictors within quantitative prediction models provides a principled and transparent way to make predictions. However, the approach relies on a limited set of variables, some of which are also somewhat subjective measures that potentially bias the results. For example, when relying on box office sales as a success




Figure 1. Schematic of video viewer interface. (A) Main panel to select a show (B) Panel to select a video clip within a show.

measure, the results strongly depend on the chosen time window (i.e., first weeks, cumulative period in theater, theater plus video sales, etc.). In addition, the predictor variables themselves can be rather subjective and thus prone to error (e.g., whether an actor is a “star,” or which critics and reviews are considered as relevant). Perhaps most importantly, available prediction models focus on full length movies and thus are of limited value when predicting the success of short video material that is becoming increasingly popular at various streaming websites on the internet such as YouTube or Vimeo.

To overcome these limitations, we propose and test an objective, quantitative, and simple measure of visual salience that is based on the raw video. Visual salience can be computed even before a video has been released, it is well suited for short video material, and it can be combined with other available predictors. While the measure requires that the video footage is readily available, knowledge about the factors that drive visual salience can also help producers to produce more attractive video material.

As a basis to test the salience measure on empirical grounds, we use data from a video streaming website. Here, we define success as how long a video was watched online before viewers navigated away from it.

1.2 Salience in Visual Computing

Visual objects differ in their salience, that is, in their capability to attract our attention. Research on visual perception indicates that this property depends on low-level perceptual features but also on features that require a higher level of cognitive processing [4, 9, 14]. From the perspective of visual computing, the salience of an element is a numerical value that quantifies the degree to which it draws attention. In this context, salience can be a property of pixels, whole images, or whole videos. This concept lends itself to an interdisciplinary approach, combining models of human behavior, knowledge of the human visual system, and low-level numerical measures that can be extracted from an image and used as predictors. Estimating the salience of an object is an active field in visual computing research, with many different models and implementations developed over the years. To estimate the salience of a two-dimensional object or scene, Itti, Koch and Niebur [10] proposed a quantitative, computationally tractable, and empirically-validated model of salience that takes objective

visual properties of the environment such as color, intensity, and orientation features as input [21]. This so-called IKN estimator is used widely, and it regularly features as a benchmark for evaluating alternative experimental measures [4]. Here, we apply their model to analyze the salience of short video clips from a streaming website. In a second step, we use statistical survival analyses to test if salience predicts how long users watch a given video.

2. METHOD

Our analysis is based on usage data from a website that allowed users to browse clips from popular television shows. The website, mocked up in Figure 1, included a video player with suggestions, channels, advertisements, and standards playback controls, similar to other streaming websites like YouTube. Because the media player was restricted to work only for U.S.-based IP addresses, all viewers are presumed to be American or, at least, English speakers familiar with American culture. On average, videos were short clips of about 3.3 minutes, though there was considerable variation in their duration ($M = 201s$, $SD = 87s$, range: 15s-462s). During a phase of several months the site aggregated users’ browsing behavior. Here, we focus on users’ probabilities of navigating away from a video before its conclusion, and how that probability is affected by the salience of the raw video that they watched.

2.1 Estimating Salience

Salience was calculated by means of the Matlab Saliency Toolbox¹, a ready-to-use package with implementations of the IKN measure. In its original form, the IKN salience assigns to each pixel a vector representing color, brightness, and orientation dimensions of visual salience. Since our aim was to obtain a single comprehensive salience value for each video, we summed the computed per-pixel saliency values across all pixels within each frame. We then computed two alternative salience measures, one based on the mean of the IKN (across all frames within each video) and the other based on its standard deviation. Both measures have a straightforward interpretation: The mean expresses the general intensity or “attention grabbing” potential that presumably demands high resources in tonic, sustained attention and the standard deviation is a proxy for the variability of the salience across a video that presumably demands high resources in phasic attention and/or alertness [8]. Knowing more about the predictive power of salience yields a better understanding of the preferences of video consumers and their underlying psychological processes.

The median of the mean salience across all videos was 31.62 with a range from 19.26 to 43.19. The median of the standard deviation-based salience was 8.41 with a range from 4.11 to 14.99.

Calculating both salience measures for each video was computationally intensive as it involved processing 1,007 videos for which raw data was available frame-by-frame, and that were watched at least once. Videos were encoded at 24–30 fps with the MPEG-4 Part 14 (MP4) codec. On an Intel® Core™ i7-3930K CPU @ 3.20GHz, computation took approximately 150 ms per frame of video, or approximately two weeks for the entire dataset.

2.2 Media Consumption and Video Data

For the final survival analyses we combined two key datasets, one consisting of the raw video data (from which we extracted saliency as well as video duration), and one consisting of users’

¹ <http://www.saliencytoolbox.net>

browsing behavior (e.g., which videos were watched, and for how long).

Our analyses did not include peripheral media content like advertisements that may have been served through the media player before or during videos. All videos were watched during the summer of 2013.

The core dataset of our analysis consisted of 81,041 data points, with behavior from 31,671 unique viewers. Most viewers watched just one video (75%) and 2% of viewers watched ten or more videos. In a few cases, a specific video was watched repeatedly by the same individual. For each viewer, we identified continuous sessions of video consumption, defined as a within-subject series of views with no more than a 30-minute gap between subsequent videos. Identifying sessions allowed us to control for large gaps in viewers' sequences of video choices, as when viewers return to a site and continue viewing at a later date.

2.3 Statistical Analysis

To test the relationship between video salience and video consumption behavior, we estimated each video's hazard rate, defined as the probability of navigating away from it before it has ended, as a function of proportion viewed, for high- and low-salience videos, separately for two measures of salience. The statistical method of survival analysis provides a feasible tool for this analysis [11]. Here, "survival" means that users have not yet clicked away from the video. Thus, the longer one watches a video, the longer this person "survives."

When conducting a survival analysis, one tries to approximate the survival function $S(t)$, that is the probability of an individual's "survival" T exceeding time t : $S(t) = P(T > t)$. Theoretically, this probability is best described by the true underlying survival function $S(t)$. Because the theoretical function is not known, one has to rely on an empirical approximation called the Kaplan-Meier (KM) method. Here, we defined "stop watching a video" as the event of interest and the percentage $p = [0,1]$ of a video watched until an event occurred as the time of survival. Because not all videos had the same duration, percentage watched can be regarded as a comparable metric across videos. We treated stopping at $p = 1$ as right-censored in our analysis because it could represent either successfully completing the video or a dropout during its last moment.

Because the Kaplan-Meier-based survival analysis requires categorical predictors, we labeled videos either "high" or "low" salience based on whether their ratings were above or below the median across all videos. Next, we estimated the survival probabilities for "high" and "low" salience values using the KM method, separately for the mean- and standard-deviation-based salience measures. For both we made an overall statistical comparison of the two empirical KM survival curves using the log-rank test ("Mantel-Haenszel test") as implemented in the `survdiff` function of the `survival` package in R [20].

3. RESULTS

To test the relationship between video consumption behavior and mean- and standard deviation-based measures of salience, we fit KM curves separately for videos with "high" and "low" values of salience, and then compared them with log-rank tests. The comparison of the KM estimates for the survival curves ("high" versus "low," see Figure 2) showed a significant difference for the mean-based measure of salience ($\chi^2 = 1,261$, $df = 1$, $p < 0.001$) and a marginally significant trend for the standard deviation-based measure ($\chi^2 = 3.7$, $df = 1$, $p = 0.0535$). Estimates for the survival probabilities at specific time points can be found in Table S3.

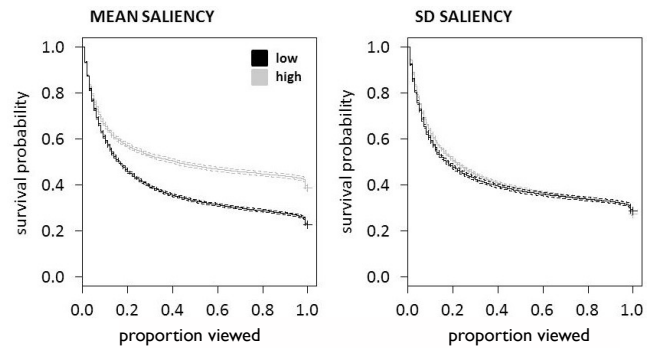


Figure 2. Kaplan-Meier estimates for the mean and standard deviation (sd) of salience, separately for high (grey) and low (black) values of the measures as defined by a median split. Videos with high and low salience show differences in the rate at which viewers drop out of videos, particularly the mean salience over the course of a video. X-axis defines dropout rates in the sense that 100% of viewers completed 0% of all videos (far right of curve) and approximately 20%–40% of viewers completed 100% of all videos (far left). Estimates for the survival probabilities at specific time points can be found in Table S3.

3.1 Robustness

To test the robustness of these results, we also conducted secondary analyses that included additional contextual predictors above and beyond salience, fully reported in the Supplementary Information, and summarized here. In particular, we repeated our analysis within the context of linear mixed-effects models that include video, the show to which it belongs, and the show's category as random effects predictors, and video salience, length in seconds, number of views, number of "like" votes, and various contextual variables as fixed effects. We repeated all of our analyses with specifications of video survival in terms of both the percentage and numbers of seconds completed. Controlling for all of these contextual factors, we identify a relationship between video length and our salience measures, such that modeling video length eliminates the linear effect of salience on video completion, whether in seconds or percentages. This effect cannot be accounted for by differences in genre. That video length is confounded with salience as a predictor of seconds watched must reflect some other unmeasured structure in our data. Of course, our dataset can offer only correlational insights, and it must be left to future work to disentangle the directionality of these influences.

4. DISCUSSION

We defined the success of a video as the amount of it that was completed (in both percentage and seconds) before a viewer navigated away, and we suggest an objective and easy to compute measure – visual salience – as a promising predictor of movie success.

Our results dovetail with past research showing that attention processes play an important role in consumer behavior [15, 16]. To measure attention, past research in consumer behavior often relied on self-reports or other "traditional" methods. Compared to this, quantitative computational methods are not susceptible to observer effects or to the vagaries of self-reports, and they are often easier to implement. IKN salience has the additional advantage that it can be readily applied to short-form free online video content, where conventional, high-level predictors for motion pictures success may be less relevant.

The ability of low-level salience measures to predict behavior in the real world may seem surprising given that individual browsing behavior presumably depends on many individual and external factors. Our results suggest that low-level psychophysical variables may play an important role in sustaining attention for the duration of the short video clips that we analyzed here. Does that imply that content providers should maximize IKN saliency to maximize video success? Maximizing this measure would involve scrambling a video so much as to make it unwatchable. Thus, presumably there is an optimal level of saliency beyond which the success of a video starts to decrease. Thus, it is an important challenge for future research to identify this ideal point and to gain a better understanding of its moderating factors. In particular, it will be crucial to better-characterize the relationship between salience and video length, a relationship that persists even after controlling for variability between different genres and different shows within a genre. Even though our bottom-up approach to salience offers encouraging results, future research may identify other, more predictive algorithms for predicting video success. Towards this goal, it would be worthwhile to apply machine learning to leverage the multidimensional output of most salience measures and find more principled ways to create summary salience measures.

In conclusion, our results on a large online behavioral dataset show that simple measures of low-level psychological salience can predict how likely viewers are to persist through online videos in a browsing environment. Our work presents raw video as a promising source of variance for predictive analytics and marketing applications.

5. ACKNOWLEDGMENTS

We thank Steven Poulakos, Aljoscha Smolic, and Maurizio Nitti for their assistance and helpful comments. Rafael Huber conducted this work with the support of Disney Research. Correspondence concerning this article should be addressed to Seth Frey.

6. REFERENCES

- [1] Ainslie, A. et al. 2005. Modeling movie life cycles and market share. *Marketing Science*. 24, 3 (2005), 508–517.
- [2] Asur, S. and Huberman, B.A. 2010. Predicting the future with social media. (2010), 492–499.
- [3] Basuroy, S. et al. 2003. How critical are critical reviews? The box office effects of film critics, star power, and budgets. *Journal of Marketing*. 67, 4 (2003), 103–117.
- [4] Borji, A. and Itti, L. 2013. State-of-the-art in visual attention modeling. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*. 35, 1 (2013), 185–207.
- [5] Chang, B.-H. and Ki, E.-J. 2005. Devising a practical model for predicting theatrical movie success: Focusing on the experience good property. *Journal of Media Economics*. 18, 4 (2005), 247–269.
- [6] Elberse, A. 2007. The power of stars: Do star actors drive the success of movies? *Journal of Marketing*. 71, 4 (2007), 102–120.
- [7] Eliashberg, J. and Shugan, S.M. 1997. Film critics: Influencers or predictors? *The Journal of Marketing*. (1997), 68–78.
- [8] Gazzaniga, M.S. 2004. *The cognitive neurosciences*. MIT press.
- [9] Itti, L. and Koch, C. 2001. Computational modelling of visual attention. *Nature reviews neuroscience*. 2, 3 (2001), 194–203.
- [10] Itti, L. et al. 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on pattern analysis and machine intelligence*. 20, 11 (1998), 1254–1259.
- [11] Kleinbaum, D.G. 1998. Survival Analysis, a Self-Learning Text. *Biometrical Journal*. 40, 1 (1998), 107–108.
- [12] Krider, R.E. and Weinberg, C.B. 1998. Competitive dynamics and the introduction of new products: The motion picture timing game. *Journal of Marketing Research*. (1998), 1–15.
- [13] Mestyán, M. et al. 2013. Early prediction of movie box office success based on Wikipedia activity big data. *PLoS one*. 8, 8 (2013), e71226.
- [14] Parkhurst, D. et al. 2002. Modeling the role of salience in the allocation of overt visual attention. *Vision research*. 42, 1 (2002), 107–123.
- [15] Pieters, R. and Warlop, L. 1999. Visual attention during brand choice: The impact of time pressure and task motivation. *Intern. J. of Research in Marketing*. 16, (1999), 1–16.
- [16] Pieters, R. et al. 1999. Visual Attention to Repeated Print Advertising: A Test of Scanpath Theory. *Journal of Marketing Research*. 36, 4 (1999), 424–438.
- [17] Ravid, S.A. 1999. Information, Blockbusters, and Stars: A Study of the Film Industry*. *The Journal of Business*. 72, 4 (1999), 463–492.
- [18] Sharda, R. and Delen, D. 2006. Predicting box-office success of motion pictures with neural networks. *Expert Systems with Applications*. 30, 2 (2006), 243–254.
- [19] Sochay, S. 1994. Predicting the performance of motion pictures. *Journal of Media Economics*. 7, 4 (1994), 1–20.
- [20] Therneau, T.M. 2014. *A package for survival analysis in S. 2014, R package version 2.37-7*. Retrieved from <http://cran.r-project.org/package=survival>.
- [21] Walther, D. and Koch, C. 2006. Modeling attention to salient proto-objects. *Neural networks*. 19, 9 (2006), 1395–1407.

Supplementary material for

The influence of visual salience on video consumption behavior: A survival analysis approach

APPENDIX

A. ROBUSTNESS CHECKS

To test the robustness of these results, we also conducted secondary analyses. The survival-based approaches do not offer very flexible support for covariates and grouping, and they force the use of categorical variables, so we repeated our analysis within the context of the following linear mixed-effects model:

$$\begin{aligned}
 video_{completion} \sim & \beta_0 + \beta_1 video_{meansalience} + \beta_2 video_{sdsalience} \\
 & + \beta_3 video_{length} + \beta_4 video_{likes} + \beta_5 \log(video_{numviews} + 1) \\
 & + \beta_6 \log(viewer_{numvideos} + 1) + \beta_7 viewer_{newcategory} \\
 & + \beta_8 viewer_{newshow} + \beta_9 viewer_{shownad} \\
 & + \beta_{10} viewer_{weekend} + \beta_{11} viewer_{daytime} \\
 & + (u_{videocategory} + u_{videoshow} + u_{video} + e)
 \end{aligned}$$

We fit two versions of the dependent, $video_{completion}$, one for amount of video watched in seconds, and the other in terms of percentage. The results of our tests are available in Table S2, Models 1.1 and 1.2 for seconds and percentage watched, respectively. The predictors are defined as follows. $video_{meansalience}$ is the median (over all videos) of the mean (over all frames of a video) values of salience. $video_{sdsalience}$ is the median (again, over all videos) of the standard deviations of frame salience within a video. Though they are listed first in the above equation, we conservatively entered them last in our model specification to allow other variables to account for as much variance as possible before fitting the salience measures. The video's length in seconds was encoded in $video_{length}$, and $video_{likes}$ gave the number of Facebook likes it elicited over the duration of the study. $video_{numviews}$ and $viewer_{numvideos}$ were, respectively, the number of views that the video received and the number of videos the viewer watched over duration of the study. We fit our model to the logarithms of these two variables because their distributions both had fat tails, with a few very popular videos and active viewers, but most videos and viewers related only once. $viewer_{newshow}$ and $viewer_{newcategory}$ were dummy variables indicating whether or not a viewer's current video is in a different show or show category (e.g. comedy, drama, reality show) than the immediately previous video that they watched in that session (If there wasn't a previous video, these dummy were set to false). The last three predictors, $viewer_{shownad}$, $viewer_{weekend}$, and $viewer_{daytime}$, were dummy variables indicating whether the video had advertising content, whether the video was viewed on a weekend, and whether it was viewed between 9AM EST and 5PM PST during a week day.

To control for differences between shows, genres, and clips, we used random effects to model the nominal variables representing each video category, each show within a category, and every video within a show (variables $u_{videocategory}$, $u_{videoshow}$, and u_{video}). The additional complexity introduced by random effects make a traditional ANOVA impossible, so we test the significance of each of these dependents with an ANOVA between the full model and a model that subtracts that specific variable. Our results show that $video_{length}$, $video_{numviews}$, $viewer_{numvideos}$, $viewer_{newcategory}$, $viewer_{newshow}$, $viewer_{shownad}$, $viewer_{weekend}$, and $viewer_{daytime}$ were significant below the $p < 0.05$ threshold for both versions of the dependent variable while $video_{likes}$, $video_{meansalience}$,

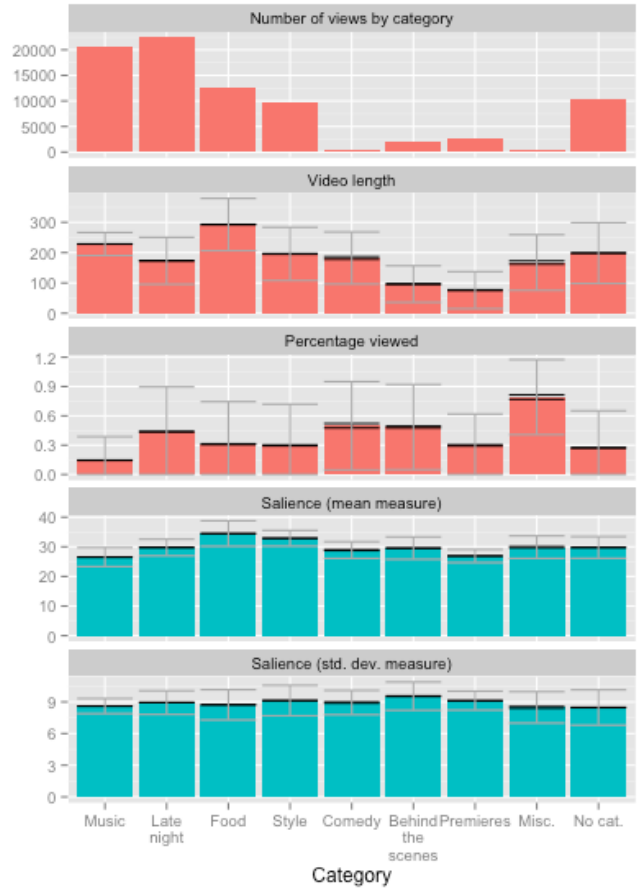


Figure S1. Descriptive statistics, and salience measures, by category. Standard errors are visualized with black error bars, standard deviations in grey.

and $video_{sdsalience}$ were insignificant (Table S2). We report covariance between these variables in Table S1. The bar charts in Figure S1 give descriptive statistics over the levels of random effect $u_{videocategory}$.

Noting the high correlation of $video_{length}$ with both salience measures and with both dependent variables, we added four additional models to better understand the colinearity between these variables. This analysis found that the deviation-based salience measure was significant for models of both dependent variables when $video_{length}$ was excluded, and that the mean-based salience measure was a significant predictor of one, the number of seconds of video watched (Table S2, models 2.1 and 2.2). Since one of the salience measures is significant for both versions of the dependent variable, the confound of $video_{length}$ with the predictive power of salience should not be tied merely to the fact that $video_{length}$ does or does not share units (seconds) with the dependent. The confound should also not be due to the fact that different shows or categories may have different lengths of video, because our analysis was performed within both show and category, modeling each as random effects. Continuing our

exploration of the relationships between video length, salience, and duration watched, we tested three interaction terms: $video_{length} * video_{meansalience}$, $video_{length} * video_{sdsalience}$, and $videomeansalience * video_{sdsalience}$. Though there were two significant effects, $video_{length}$ and its interaction with $videomeansalience$, with one of the dependents, these interactions do not do much to illuminate the nature of the relationship between their terms. That video length is confounded with salience as a

predictor of seconds watched must reflect some other unmeasured structure in our data. Of course, our dataset can offer only correlational insights, and it must be left to future work to disentangle the directionality of these influences. This work should focus on the relationship between $video_{length}$ and salience as predictors of how much a video is completed.

Table S1. Correlation matrix for fixed effects in full linear mixed effects model.

	<i>Duration</i> _{seconds}	<i>Duration</i> _{percentage}	<i>video</i> _{meansalience}	<i>video</i> _{sdsalience}	<i>video</i> _{length}	<i>video</i> _{likes}	<i>log</i> (<i>video</i> _{numviews} +1)	<i>log</i> (<i>viewer</i> _{numvideos} +1)	<i>viewer</i> _{newcategory}	<i>viewer</i> _{newshow}	<i>viewer</i> _{shownad}	<i>viewer</i> _{weekend}	<i>viewer</i> _{daytime}
<i>Duration</i> _{seconds}	1.00												
<i>Duration</i> _{percentage}	0.84	1.00											
<i>video</i> _{meansalience}	0.11	0.07	1.00										
<i>video</i> _{sdsalience}	-0.03	0.08	0.27	1.00									
<i>video</i> _{length}	0.18	-0.18	0.17	-0.29	1.00								
<i>video</i> _{likes}	-0.14	-0.18	-0.19	0.05	0.01	1.00							
<i>log</i> (<i>video</i> _{numviews} +1)	-0.22	-0.29	-0.14	0.01	0.11	0.55	1.00						
<i>log</i> (<i>viewer</i> _{numvideos} +1)	0.27	0.29	0.04	0.01	-0.07	-0.18	-0.39	1.00					
<i>viewer</i> _{newcategory}	-0.00	-0.00	0.02	0.02	-0.02	-0.02	-0.06	0.12	1.00				
<i>viewer</i> _{newshow}	0.14	0.20	-0.00	0.03	-0.11	-0.12	-0.17	0.42	0.26	1.00			
<i>viewer</i> _{shownad}	0.18	0.17	0.02	0.00	0.00	0.03	-0.14	0.08	0.01	0.04	1.00		
<i>viewer</i> _{weekend}	0.03	0.04	-0.00	-0.03	-0.03	0.00	-0.04	0.03	-0.01	0.01	-0.00	1.00	
<i>viewer</i> _{daytime}	-0.04	-0.06	-0.03	-0.01	0.03	0.02	0.04	-0.04	-0.01	-0.04	-0.03	-0.24	1.00

Notes. CI99 = 99% confidence interval. Hazard ratios are computed from survival probability S(t) according to formula $\frac{1 - S(t)_{low}}{1 - S(t)_{high}}$

Table S2. Significance of control variables in linear mixed effects models

	Model 1.1 (secs)	Model 1.2 (%)	Model 2.1 (secs)	Model 2.2 (%)	Model 3.1 (secs)	Model 3.2 (%)
<i>Intercept</i>	33.00	0.59	97.00	0.46	6.61	0.73
<i>videolength</i>	0.31 ***	0 ***			0.51 ***	0.00
<i>log(video_{numviews}+1)</i>	-4.5 ***	-0.02 ***	-3.9 ***	-0.03 ***	-4.55	-0.03
<i>log(viewer_{numvideo}_s+1)</i>	12.99 ***	0.05 ***	12.89 ***	0.05 ***	13.01	0.05
<i>viewer_{newcategory}</i>	-13.3 ***	-0.09 ***	-13.13 ***	-0.09 ***	-13.30	-0.09
<i>viewer_{newshow}</i>	6.28 ***	0.06 ***	6.36 ***	0.06 ***	6.25	0.06
<i>viewer_{shownad}</i>	45.44 ***	0.21 ***	45.49 ***	0.21 ***	45.45	0.21
<i>viewer_{weekend}</i>	2.37 ***	0.01 **	2.34 **	0.01 **	2.37	0.01
<i>viewer_{daytime}</i>	-4.78 ***	-0.02 ***	-4.75 ***	-0.02 ***	-4.78	-0.02
<i>videolikes</i>	-4.95	-0.02	-2.14	-0.02	-4.62	-0.02
<i>videomeansalienc</i>	0.26	0.00	1.01 *	0.00	1.19	-0.01
<i>videosdsalienc</i>	-1.03	0.00	-5.58 ***	0.01 *	-2.24	-0.03
<i>videolength* videomeansalienc</i>					-0.01 *	0.00
<i>videolength* videosdsalienc</i>					0.00	0.00
<i>videomeansalienc* videosdsalienc</i>					0.03	0.00

Notes. Each column represents a separate model. Rows represent the significance of each variable in that given model. Values in bold have $p < 0.05$ (*; ** for $p < 0.01$; *** for $p < 0.001$). All p -values are based on chi-square tests from an ANOVA between versions of the column's model with and without the row's variable; consequently all tests were parameterized with one degree of freedom ($\chi^2_1(N=81,041)$). Values in grey are as reported from the corresponding linear mixed-effects model but were not tested for the significance of their differences from zero.

Table S3. Survival probability for different levels of percentage viewed.

summary statistic	t	high salience		low salience		hazard ratio
		S(t)	CI _{99%}	S(t)	CI _{99%}	
	0.01	.934	[.929, .939]	.935	[.932, .939]	0.98
	0.05	.759	[.751, .767]	.729	[.723, .734]	1.12
	0.10	.652	[.643, .661]	.587	[.581, .594]	1.19
mean	0.15	.599	[.590, .609]	.511	[.504, .517]	1.22
	0.25	.545	[.536, .555]	.420	[.413, .426]	1.27
	0.50	.482	[.472, .491]	.329	[.323, .335]	1.30
	0.75	.452	[.443, .462]	.292	[.286, .298]	1.29
	1.00	.386	[.376, .395]	.227	[.221, .232]	1.26
standard deviation	0.01	.943	[.939, .946]	.924	[.919, .928]	1.33
	0.05	.749	[.743, .755]	.723	[.715, .730]	1.10
	0.10	.623	[.616, .630]	.586	[.578, .595]	1.10
	0.15	.553	[.546, .560]	.518	[.510, .527]	1.08
	0.25	.470	[.463, .477]	.445	[.436, .453]	1.05
	0.50	.382	[.375, .389]	.373	[.364, .381]	1.01
	0.75	.345	[.338, .351]	.342	[.334, .350]	1.00
	1.00	.273	[.267, .279]	.285	[.277, .292]	0.98