

The Rise of Mobile and Social Short-Form Video: An In-depth Measurement Study of Vine

Baptist Vandersmissen
Multimedia Lab, ELIS, Ghent University
iMinds, Ghent, Belgium
baptist.vandersmissen@ugent.be

Frédéric Godin
Multimedia Lab, ELIS, Ghent University
iMinds, Ghent, Belgium
frederic.godin@ugent.be

Abhineshwar Tomar
Multimedia Lab, ELIS, Ghent University
iMinds, Ghent, Belgium
abhineshwar.tomar@ugent.be

Wesley De Neve
¹Multimedia Lab, ELIS, Ghent University
iMinds, Ghent, Belgium
²Image and Video Systems Lab
KAIST, Daejeon, South Korea
wesley.deneve@ugent.be

Rik Van de Walle
Multimedia Lab, ELIS, Ghent University
iMinds, Ghent, Belgium
rik.vandewalle@ugent.be

Abstract

Thanks to the increasing popularity of mobile devices and online social networks, mobile and social video is on the rise, calling for a better understanding of its usage and future impact. In this paper, we provide an in-depth measurement study of Vine, a mobile application that is used for creating and sharing short looping videos of up to six seconds in length. Based on a dataset of 851,039 tweets containing a Vine URL, we investigate different aspects of Vine, including hashtag usage, video popularity and user attention. For the dataset used, we find that 34% of the Vine videos contain at least one hashtag, a percentage that is four times higher than the percentage of tweets that is in general annotated with at least one hashtag.

In addition, we can observe that a Vine video that is shared frequently on Twitter within hours after its creation will have more likes on Vine after one week, compared to a Vine video that is not shared frequently on Twitter during this same period of time. However, we cannot establish a clear link between the number of tweets sharing a Vine video and its resulting popularity. Finally, by analyzing the evolution of the number of likes and the number of shares received by a Vine video on Vine and Twitter, respectively, we can conclude that a Vine video receives most user attention shortly after its creation, with the amount of user attention received not stopping completely but remaining stable for days to weeks after its creation.

1 Introduction

With the increase in popularity of smartphones and online social networks, mobile and social video is on the rise. Vine, established in January 2013 and immediately acquired by Twitter, is the first well-known mobile application to focus on short-form

Copyright © by the paper's authors. Copying permitted only for private and academic purposes.

In: S. Papadopoulos, P. Cesar, D. A. Shamma, A. Kelliher, R. Jain (eds.): Proceedings of the SoMuS ICMR 2014 Workshop, Glasgow, Scotland, 01-04-2014, published at <http://ceur-ws.org>

video. In the months following its release, Vine quickly gained an active user base and was reported to be the world’s fastest growing mobile application at that time [9]. In August 2013, Vine topped 40 million users, withstanding the successful launch of Instagram Video [7].

Vine enables its users to create and distribute short looping videos of up to six seconds in length. In addition, its users can follow other users, re-broadcast videos to their followers by so-called reviving, comment on videos and embed videos on websites. Furthermore, its users also have the option of sharing videos to followers on Twitter or other online social networks.

The video length limitation of Vine resembles the message length limitation of Twitter, relying on the creativity of its users to spread essential information. Similar to Twitter, Vine is well suited for fast spreading of news, albeit on a visual level. This became clear with the Boston Marathon bombing tragedy, seeing the use of Vine as a social news platform [2]. However, the low threshold to create and share Vine videos entails a significant amount of noisy data. This, combined with the typical short video length and the limited availability of context information, makes it for instance hard to organize and browse Vine videos.

In this paper, we present an in-depth measurement study of Vine. We use Twitter as an access portal to harvest Vine videos and context information, exploiting the resulting dataset to achieve a better understanding of hashtag usage, video popularity and user attention, among other aspects. To the best of our knowledge, this is the first academic study of Vine.

We organized the rest of this paper as follows. In Section 2, we discuss related work. In Section 3, we explain the way we collected Vine videos. In Section 4, we investigate the general characteristics of Vine, subsequently focusing on creation time and origin aspects in Section 5, video popularity aspects in Section 6, and user attention aspects in Section 7. Finally, we present conclusions and directions for future research in Section 8 and in Section 9, respectively.

2 Related Work

In this section, we review a number of representative research efforts in the area of online social networks, paying particular attention to the following topics: content and audience analysis, popularity analysis and prediction, usage of online social network context, and social sensing.

2.1 Content and Audience Analysis

The authors of [6] performed a large-scale and in-depth measurement study of YouTube, discovering significant differences between YouTube videos and traditional streaming videos in terms of video lengths, access patterns and active life spans. Furthermore, they looked into growth trends and social aspects of YouTube. In [1], Twitter is used to analyze the who, what and when questions related to YouTube. Through combining the user- and sharing-centric data of Twitter with the video-centric data of YouTube, the authors are able to establish links between initial Twitter shares and the total number of views, as well as between Twitter shares and the type of content. The authors of [17] investigated the temporal, social and spatial dimensions of Flickr user behaviour. They conclude that 50% of the photo views are generated within the first two days. Furthermore, they also state that the social networking behaviour of users and photo pooling are the most important indicators of the popularity of a photo.

2.2 Popularity Analysis and Prediction

Several studies analyzed the popularity distribution of user-generated videos and images on online social networks such as YouTube, Flickr and Instagram. The authors of [4] analyzed the popularity life-cycle of user-generated content originating from YouTube in relation to the video age and level of content aliasing. The authors of [3] investigated the impact of content-agnostic factors on YouTube video popularity, finding that the current view count is the most important factor to consider when predicting the future popularity of a video, with the exception of videos that have been shortly uploaded. In the latter case, the size of the social network of the uploader is more important for future popularity prediction purposes. On Flickr, the authors of [5] analyzed how information propagates throughout the network, with the aim of gaining insight into the viral spreading of particular items. They state that information exchanged among friends is the most dominant factor leading to propagation throughout the network.

2.3 Usage of Online Social Network Context

A wide range of studies is available on the use of online social network context for designing new and improved algorithms for multimedia content analysis. In [16], a face recognition method is combined with information derived from Facebook in order to improve the accuracy of face recognition on personal photographs. Equivalent to the above, [13] used the collective knowledge in Flickr to build an image tag recommendation system.

2.4 Social Sensing

Social multimedia systems such as Vine and Twitter allow supporting studies on social behaviour. In particular, these systems can be looked upon as Participatory Sensing Systems (PSSs), making it for instance possible to study city dynamics on a large scale. In [14, 15], Instagram and Foursquare are used as PSSs, with the aim of analyzing user movement patterns, finding points of interests and observing cultural behaviour. A more general overview on the way computational analysis and visualization of PSS content can contribute to the identification of social and cultural patterns can be found [10].

3 Data Collection

In this section, we briefly describe the acquisition of Vine data. Because of the lack of an official and public Vine API, we used Twitter as a gateway to access and harvest Vine videos. Furthermore, we also used the unofficial Vine API methods to extract metadata.

We harvested tweets containing Vine URLs by tracking the keyword “vine” via the public Twitter streaming API from January 10, 2014 until January 24, 2014. This resulted in 851,039 tweets containing Vine URLs, originating from 365,188 different Twitter users. We then used HTML scraping to extract the unique Vine ID from each Vine URL. Next, we used the extracted Vine ID and the private Vine API methods to fetch information regarding the specific properties of the Vine video and its corresponding user. By making use of the aforementioned approach, we were able to collect 425,971 unique Vine videos that have been created by 193,355 unique Vine users. The key properties fetched can be found in Table 1. We note that Vine does not provide metadata regarding the view count of a Vine video. In this study, we therefore make use of an aggregation of the number of likes, revines and comments to assess the popularity of a Vine video (cf. Section 6).

The size of our dataset is not representative for the number of videos shared on Vine during the above mentioned period. However, our dataset is representative for the number of Vine shares on Twitter during this period. The strong interweaving between Vine and Twitter allows us to measure characteristics of the dataset using both Vine and Twitter metadata. The Twitter metadata consists of a tweet containing a Vine URL and the Twitter user sharing this tweet.

The above dataset is used in all of our experiments, with the exception of Section 4.3, which analyzes the popularity of the different Vine channels

Table 1: Vine metadata.

Vine	Vine User
Date Fetched	Username
Date Created	Location
Description	Followers Count
Location	Following Count
Number of Likes	Posts Count
Number of Revines	Like Count
Number of Comments	Verified
Explicit Content	

(i.e., the different Vine categories). Due to the fact that the Vine metadata do not describe to what channel a Vine video was added, we created a separate smaller dataset for assessing the distribution of Vine videos over the different Vine channels. By using the unofficial Vine API, we were able to identify the different channels and their unique IDs. We subsequently crawled each channel’s list of newly added Vine videos between December 6, 2013, and December 12, 2013 in a continuous manner. The resulting dataset contains 370,410 unique videos belonging to 16 different channels (see Section 4.3 for more details).

4 General Vine Characteristics

In this section, we investigate various characteristics of Vine, including the technical and metadata characteristics of Vine videos. We also pay attention to the characteristics of Vine channels and popular Vine content.

4.1 Technical Characteristics

Table 2: Vine video statistics

	Min	Max	Mean	Std. Dev.
Length (s)	1.4	7.6	6.1	0.73
File size (MB)	0	2.23	0.82	0.24
Bitrate (Mbps)	0.08	2.97	1.12	0.2

A Vine video can be looked upon as a visual tweet. Characterized by its limited video length of only six seconds, users are forced to be concise. Typically, a Vine video has a square frame width and height of 480 pixels. In Table 2, we summarize the size, bitrate and length properties of 5,000 videos randomly sampled from our dataset. We can observe that the average file size of a Vine video is less than 1 MB. This average file size is much smaller than the average file size of YouTube videos, which was estimated to be 8.4 MB in [6]. We can also see that the average bitrate of a Vine video is about 1.12 Mbps, thus allowing for high-quality streaming, even when the video contains a lot of motion. Finally, we can observe that Vine contains

videos with a maximum length higher than the six seconds allowed. This can most likely be attributed to a hack of the application.

4.2 Metadata Characteristics

Vine videos can be given a description that may contain hashtags or mentions of other Vine users. Hashtags and mentions assist people and algorithms in understanding the video content and the formation of communities [11]. Therefore, it is important to know to what extent users annotate videos on Vine with hashtags and mentions.

Given our dataset, we investigated the use of hashtags and mentions that have been assigned to Vine videos. Our analysis revealed that 34.0% of Vine videos contain at least one hashtag, while 9.24% of Vine videos contain one or more mentions. In this context, we would like to note that the percentage of Vine videos containing a hashtag is significantly higher than the percentage of tweets containing a hashtag (i.e., less than 8%, according to [12]). Furthermore, our analysis revealed that a Vine video contains, on average, 0.87 hashtags and 0.13 mentions.

In Figure 1, we show the distribution of the hashtag frequency on a log-log scale. The x -axis refers to the 94,716 unique hashtags, ordered by descending hashtag frequency, whereas the y -axis refers to the hashtag frequency. This distribution can be modeled accurately by a power law, with the probability of a hashtag having frequency x being proportional to $x^{-0.934}$.




Figure 1: Hashtag frequency distribution.

Similar to the hashtag frequency, we can plot the distribution of the number of hashtags per Vine video. Figure 2 shows the number of Vine videos with x hashtags.




Figure 2: Number of hashtags per Vine video.

To estimate the information content of the hashtags used, we mapped the hashtags in our dataset onto the WordNet synsets [8], finding that 11.4% of the hashtags used could be matched to the WordNet synsets. This low percentage is indicative of the use of an uncontrolled hashtag vocabulary and of the presence of a high number of noisy hashtags.

We additionally mapped the set of matched hashtags onto the WordNet categories [8]. Figure 3 shows the distribution of the hashtags matched over the different WordNet categories. We can observe that the category people or groups is tagged most frequently (20%), followed by objects or artifacts (19%), actions or events (11%), locations (7%), and emotions or cognitions (4%). The category other (39%) contains the hashtags matched that could not be mapped onto the aforementioned WordNet categories. Our results show that the hashtags used describe a wide range of concepts (i.e., people, objects, actions, events, locations, and so on), information that can be leveraged by techniques for video classification and video concept detection.

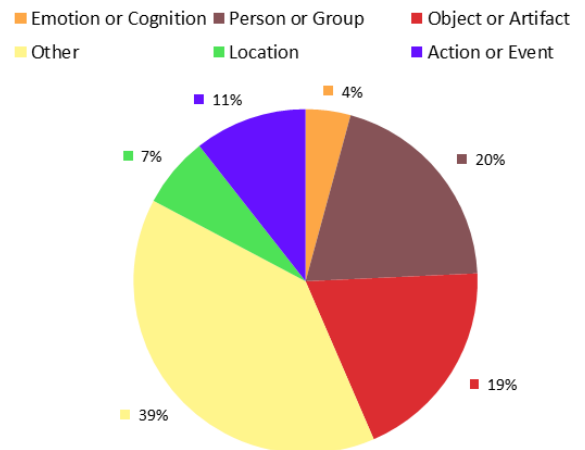


Figure 3: Most frequent WordNet categories for Vine video hashtags.

4.3 Channel Characteristics

Vine contains multiple channels (i.e., categories) to which a newly created video can optionally be published to. Table 3 gives an overview of the different channels that are currently in use. To study the popularity of these channels, we collected 370,410 Vines by following the procedure described in Section 3. Table 3 makes clear that the distribution of the number of videos over the different Vine channels is highly skewed: “Comedy” is by far the most popular channel, followed by “Music” and “Wierd”. Clearly, the focus is on entertaining and non-informative content. This is comparable to the YouTube measurement study presented in [6], who similarly reported that “Entertainment” and “Music” are the top video categories on YouTube.

Table 3: List of Vine channels

Rank	Category	Count	Pct.
1.	Comedy	225,794	60.96 %
2.	Music	33,078	8.93 %
3.	Wierd	19,513	5.27 %
4.	Dogs	16,525	4.46 %
5.	Cats	12,048	3.25 %
6.	Family	10,152	2.74 %
7.	Art & Experimental	10,141	2.74 %
8.	Sports	6,964	1.88 %
9.	Food	5,949	1.61 %
10.	Special fx	5,642	1.52 %
11.	Nature	5,458	1.47 %
12.	Urban	5,226	1.41 %
13.	Scary	4,691	1.27 %
14.	Beauty & Fashion	4,041	1.09 %
15.	News & Politics	2,774	0.75 %
16.	Health & Fitness	2,414	0.65 %

4.4 Popularity Characteristics

To gain a more detailed insight into what type of video content is popular on Vine, we collected the top 100 most popular Vine videos in our dataset, measuring popularity by multiplying the number of likes, revines and comments. Through a manual inspection, we learned that the resulting collection contains user-generated Vine videos that are not related to a particular event or brand. This is also illustrated by Figure 4, presenting an image collage of the top 6 most popular Vine videos. We can thus conclude that Vine is primarily used for producing and sharing concise and creative content among its users.



Figure 4: Image collage of the top 6 most popular Vine videos, obtained from our dataset by aggregating the number of likes, revines, and comments.

Although the majority of Vine videos can be classified as entertaining and non-informative content, we could observe that our entire dataset does contain Vine videos that are related to news or sports events. To get an impression of the nature of these Vine videos, Figure 5 and Figure 6 show snapshots of Vine videos covering a number of recent events (e.g., the Golden Globes, the Purdue shooting, the Australian Open, and so on). We retrieved these videos from our dataset by using different hashtags (e.g., #goldenglobes, #purdue, #australianopen, and so on). Note that these videos often give personal comments on events, either showing news-related images or presenting live footage of the video creator being present at the event. As such, these videos could be seen as an addition to text-based news reporting, giving different insights into a global event.



Figure 5: Image collage of Vine videos representing recent news events.



Figure 6: Image collage of Vine videos representing recent sports events.

5 Time and Place of Creation Aspects

In this section, we present findings regarding the time and place of creation of Vine videos. Recall that our dataset contains 425,971 unique Vine videos derived from 851,039 tweets containing a Vine URL, thus implying that a major part of these tweets share the same Vine video. Each Vine video possesses a creation time stamp retrieved via the private Vine API and a location field derived from the location field of the Vine user. In 27.2% of the cases, we were able to match the location of the Vine user using the Google Geocoding API.




Figure 7: Number of Vine videos created and binned per hour. General UTC time is used.

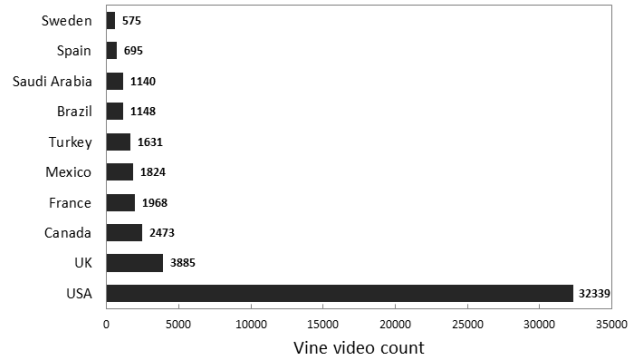


Figure 8: Number of Vine videos created per country.

In Figure 7, we show the number of Vine videos created and binned per hour, whereas Figure 8 shows which countries are creating the most videos on Vine. Clearly, Figure 7 is heavily influenced by the timezones applicable in the countries where Vine is the most popular. As such, Figure 9 also gives an overview of the number of Vine videos created in the USA, normalized per timezone. We can observe that the creation of Vine videos peaks during the afternoon and drops during night time.




Figure 9: Number of Vine videos created by USA users, normalized according to the three main timezones.

6 Video Popularity Aspects

In this section, we investigate how the popularity of Vine videos changes over time. Since there is no view count information of a Vine video, we quantify the popularity of Vine videos by three parameters: the number of likes, the number of revines, and the number of comments on Vine. We pay particular attention to the influence Twitter has on the popularity of Vine videos.

We equate Twitter attention to the number of shares S_T fetched, and where these shares have been produced by Twitter users distributing tweets that contain a link to a Vine video. We hypothesize that a higher number of shares on Twitter can be linked to a more popular Vine video (i.e., a Vine video with a relatively high number of likes, revines and comments).

We define the Twitter exposure E_T as the sum of the number of followers of the different Twitter users sharing these tweets. Similarly, we define the Vine exposure E_V as the number of followers on Vine of the original Vine video creator. We also hypothesize that sharing a Vine video with a large user base automatically results in a high popularity. Table 4 outlines the Pearson (first value) and Spearman (second value) correlation values between S_T , E_T and E_V on the one hand, and the number of likes, revines and comments on Vine on the other hand.

Table 4: Correlation between factors influencing Vine popularity and the number of likes, revines and comments on Vine.

	S_T		E_T		E_V	
Likes	0.288	0.537	0.069	0.288	0.602	0.666
Revines	0.318	0.597	0.075	0.325	0.459	0.666
Comments	0.308	0.597	0.073	0.310	0.394	0.661

We can observe that the number of shares on Twitter is only weakly correlated with each of the Vine popularity indicators used. This undermines our first hypothesis that a higher number of shares on Twitter can be linked to a more popular Vine video, showing that, despite the close relation between Vine and Twitter, each platform functions according to its own rules. Indeed, notwithstanding the fact that tweeting a Vine link automatically embeds the corresponding Vine video on Twitter, this embedding does not allow for liking, revining, or commenting. In other words, a Twitter user cannot directly add to the popularity of a Vine video, except when the user likes, revines or comments on the Vine platform itself. Furthermore, we also find no correlation between the Twitter exposure and the different popularity indicators. The second hypothesis that sharing a Vine video with a large user base automatically results in a high popularity does show to be correct, given the relatively strong correlation between the Vine exposure and the Vine popularity indicators.

Albeit, we cannot measure the actual impact of Twitter on a Vine video’s popularity due to the lack of view information and a clear correlation between the number of Twitter shares and the popularity measures, we state that a Vine video that is shared more than once on Twitter (i.e., not just by the creator of the Vine video) in an early stage (i.e., within the hour) will be an indicator for the popularity of a Vine video.




Figure 10: The number of shares of a Vine video on Twitter one hour after its creation linked to the number of likes on Vine after one week.

Figure 10 shows the number of shares on Twitter of a Vine video one hour after its creation linked to the number of likes this video has on Vine after one week. We can observe that an initial correlation exists between the number of shares on Twitter and the number of likes on Vine. In particular, this correlation is strongest in the beginning of the curve (i.e., when the number of shares on Twitter is less than 15) but less obvious when the video is shared more (i.e., when the number of shares on Twitter starts to become higher than 15). For a higher number of Twitter shares, we can similarly to the findings above state that we cannot measure the impact of Twitter. This is due to the same reasons as stated above and also related to factors such as the social impact of the Twitter users sharing the Vine video (i.e., the number of retweets per tweet) or the social network of the Vine users revining on Vine.

7 User Attention Aspects

In this section, we investigate the amount of user attention received by Vine videos. Our analysis is twofold: 1) we study the evolution of the number of likes on Vine of Vine videos and 2) we study the evolution of the number of shares on Twitter of Vine videos. Both aspects are studied in relation to the number of hours following the creation of the Vine videos. As such, we define user attention as the number of likes on Vine gained or the number of shares spread on Twitter during a certain time span. Due to the fast nature of Vine and Twitter, we hypothesize that the user attention span is short and that user attention peaks shortly after the creation of a Vine video.

First, we analyze the evolution of the number of likes given to a Vine video during the first two weeks after its creation. For this analysis, we only take into account Vine videos that have been created in the USA and that have received at least five shares one

hour after their creation, resulting in the use of 3,312 Vine videos having 32.1 Twitter shares on average.

Figure 11 shows the evolution of the average number of likes per Vine video. We can observe that the increase in the average number of likes is highest one day after the creation of a Vine video. However, we can also observe that the average number of likes keeps increasing steadily during subsequent days.

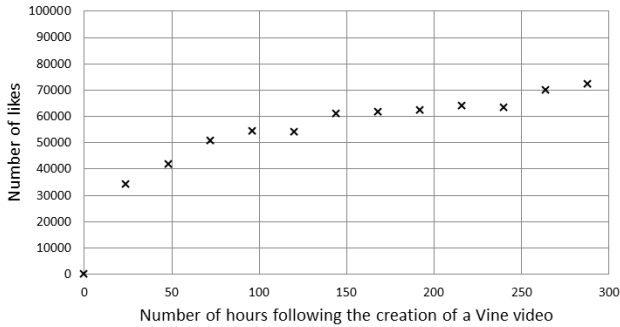


Figure 11: The average number of likes received by a Vine video as a function of the number of hours following its creation.

Second, we analyze the evolution of the number of Twitter shares given to a Vine video in relation to the number of hours following its creation. For this analysis, we only take into account Vine videos that have been shared on Twitter, both within one hour after their creation and after seven days of their creation, resulting in the use of 10,696 Vine videos. Figure 12, which uses a log-log scale, shows a trend that is comparable to the trend shown in Figure 11. We can observe that a Vine video receives most user attention on Twitter during the first hours after its creation. Note that the distribution shown in Figure 12 can be modeled by a power law-like distribution with $\alpha = 0.649$.

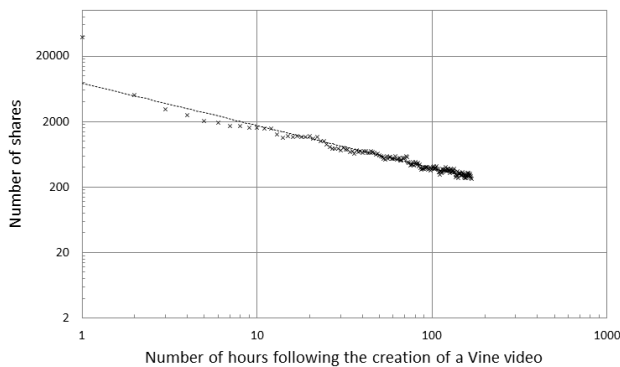


Figure 12: The average number of Twitter shares received by a Vine video as a function of the number of hours following its creation.

8 Conclusions

In this paper, we presented a large-scale measurement study of Vine, a popular mobile application for creating and posting short looping videos, paying particular attention to Vine hashtag usage, video popularity and user attention. To that end, we made use of Twitter as an access portal to Vine, harvesting 851,039 tweets containing a Vine URL. To the best of our knowledge, this is the first academic study of Vine, with the aim of achieving a better understanding of mobile and social short-form video.

In our dataset, we could observe that Vine videos have an average length of about 6.1 seconds and an average file size of 0.82 MB. We could also observe that 34% of the Vine videos in our dataset contained at least one hashtag, a percentage that is significantly higher than the 8% of tweets that is in general annotated with at least one hashtag. Furthermore, we found that 11.4% of the Vine hashtags used could be matched to the WordNet synsets. By subsequently mapping the matched Vine hashtags onto the WordNet categories, we also found that the category people or groups is tagged most frequently (20%), followed by objects or artifacts (19%), actions or events (11%), locations (7%), and emotions or cognitions (4%).

Through our study, we could learn that the content of Vine videos is typically highly personal, mostly created for entertainment purposes. However, we could also observe that Vine videos are created when notable events take place, possibly bringing Vine forward as a visual Twitter-alike social news platform in the near future.

We investigated the popularity of Vine videos by making use of both Vine and Twitter metadata, finding that Twitter cannot be used as a measure for the popularity of Vine videos. However, we did observe that Vine videos shared frequently on Twitter in an early stage after their creation are more likely to have more likes on Vine after one week, an effect that could not be observed when the number of tweets sharing the same Vine video becomes bigger. The latter can be mainly attributed to the inability to measure the amount of Twitter attention given to Vine videos. Indeed, when an embedded Vine video is viewed on Twitter, this is not notable in any Vine metadata as Twitter does not allow to directly like, revine or comment on Vine.

Finally, we also investigated the average amount of user attention given to Vine videos by studying the evolution of the number of shares on Twitter and the

number of likes on Vine. We could notice that the number of shares of Vine videos on Twitter is highest in the hours after their creation and then drops significantly, following a power law-like distribution. We expected that the user attention span of Vine videos would be short, in the order of a couple days to a week, but found that, although most user activity indeed occurs shortly after their creation, the number of likes still keeps increasing after the first week of their creation.

9 Future Research Directions

Given that the findings of our measurement study of Vine are meant to tailor future technological research in the domain of mobile and social video, we present a number of directions for future research we plan to work and collaborate on.

- To investigate what specific portion of Vine videos is news- or event-related and in what way these videos can be used to enhance news stories on the Internet.
- To create a (geo-based) hashtag recommendation and categorization system for Vine videos by making use of both content (visual) and context (textual) information.
- To leverage user mentions to detect community formation.
- To identify user segments based on categorization of the content and the context of Vine videos.
- To create personalized television channels based on hashtags and user preferences.
- To compare short-form video usage on Vine with short-form video usage on YouTube (MixBit) and Instagram Video.
- To make use of Vine as a dataset for the creation of robust video face detection and recognition algorithms.

10 Acknowledgments

The research activities described in this paper were funded by Ghent University, iMinds, the Institute for Promotion of Innovation by Science and Technology in Flanders (IWT), the FWO-Flanders, and the European Union.

References

- [1] A. Abisheva, V. R. K. Garimella, D. Garcia, and I. Weber. Who Watches (and Shares) What on YouTube? And When? Using Twitter to Understand YouTube Viewership. *CoRR*, 2013.
- [2] Adage. Boston Marathon Bombing Makes Vine a News Platform, Apr. 2013. Available at <http://adage.com/article/media/boston-marathon-bombing-makes-vine-a-news-platform/240915>.
- [3] Y. Borghol, S. Ardon, N. Carlsson, D. Eager, and A. Mahanti. The Untold Story of the Clones: Content-agnostic Factors That Impact YouTube Video Popularity. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '12, 2012.
- [4] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon. I Tube, You Tube, Everybody Tubes: Analyzing the World's Largest User Generated Content Video System. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement*, IMC '07, 2007.
- [5] M. Cha, A. Mislove, and K. P. Gummadi. A Measurement-driven Analysis of Information Propagation in the Flickr Social Network. In *Proceedings of the 18th International Conference on World Wide Web*, WWW '09, 2009.
- [6] X. Cheng, J. Liu, and C. Dale. Understanding the Characteristics of Internet Short Video Sharing: A YouTube-Based Measurement Study. *Multimedia, IEEE Transactions on*, pages 1184–1194, 2013.
- [7] CNet. Vine Grows to 40 Million Users Despite Instagram Challenge, June 2013. Available at http://news.cnet.com/8301-1023_3-57599369-93/vine-grows-to-40-million-users-despite-instagram-challenge/.
- [8] C. Fellbaum. *WordNet: An Electronic Lexical Database*. Bradford Books, 1998.
- [9] GlobalWebIndex. Fastest Growing Apps, Oct. 2013. Available at <http://blog.globalwebindex.net/mobile-app-usage>.
- [10] N. Hochman and L. Manovich. Zooming into an Instagram City: Reading the local through social media. *First Monday*, 2013.
- [11] B. Huberman, D. Romero, and F. Wu. Social networks that matter: Twitter under the microscope. *First Monday*, 14(1), 2009.
- [12] S. M. Kywe, T.-A. Hoang, E.-P. Lim, and F. Zhu. On Recommending Hashtags in Twitter Networks. In *Proceedings of the 4th International Conference on Social Informatics*, SocInfo'12, 2012.

- [13] B. Sigurbjörnsson and R. van Zwol. Flickr Tag Recommendation Based on Collective Knowledge. In *Proceedings of the 17th International Conference on World Wide Web*, WWW '08, 2008.
- [14] T. H. Silva, P. O. S. V. d. Melo, J. M. Almeida, J. Salles, and A. A. F. Loureiro. A Picture of Instagram is Worth More Than a Thousand Words: Workload Characterization and Application. In *Proceedings of the 2013 IEEE International Conference on Distributed Computing in Sensor Systems*, DCOSS '13, 2013.
- [15] T. H. Silva, P. O. S. Vaz de Melo, J. M. Almeida, J. Salles, and A. A. F. Loureiro. A Comparison of Foursquare and Instagram to the Study of City Dynamics and Urban Social Behavior. In *Proceedings of the 2Nd ACM SIGKDD International Workshop on Urban Computing*, UrbComp '13, 2013.
- [16] Z. Stone, T. Zickler, and T. Darrell. Autotagging Facebook: Social network context improves photo annotation. In *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW '08. IEEE Computer Society Conference on*, 2008.
- [17] R. van Zwol. Flickr: Who is Looking? In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence*, WI '07, 2007.