

Anatomy of Viral Social Media Events

Essi Pöyry

Consumer Society Research Centre
University of Helsinki
essi.poyry at helsinki.fi

Salla-Maaria Laaksonen

Consumer Society Research Centre
University of Helsinki
salla.laaksonen at helsinki.fi

Arto Kekkonen

Department of Computer Science
University of Helsinki
arto.kekkonen at helsinki.fi

Juho Pääkkönen

Department of Computer Science
University of Helsinki
Aalto University
juho.paakkonen at helsinki.fi

Abstract—Discussion topics go sometimes viral in social media without a seemingly coherent pattern. Existing literature shows these discussions can reach a very high level, but, notably, they prevail to varying degrees. This paper investigates the anatomy of viral social media events using a dataset of 960 viral social media discussion topics that have been identified by an algorithm from a variety of social media sources over two years' time. A negative binomial regression shows that the average daily amount and the relative change in the daily amount of social media platforms at which the event has been discussed has a positive effect on the duration of the event. Average or relative amount of posts or authors has no or very little effect on event duration. The results suggest that viral social media events last longer when people using different social media platforms get exposed to them. This finding contributes to the literature on social media events, virality, and information diffusion.

1. Introduction

In the contemporary media landscape, examples of viral events that gather a lot of attention are plentiful, and social media plays an important role both in creating and supporting these events. Examples indicate that such events also bear real-life consequences. In 2017, passengers of a United Airlines flight posted videos of a man being forcibly dragged out of a fully booked airplane, which caused an international stir mainly taking place on social media. Soon after, United Airlines stock dropped 1.4 billion dollars [53]. Another example from 2017 folds around a hashtag 'penelopegate', which refers to a scandal regarding the French presidential candidate François Fillon who allegedly had employed his

wife Penelope to a fictitious job and paid one million euros in public wages. Extensive social media coverage was followed by the news revelation and support for Fillon plummeted in the wake of accusations [57].

What is characteristic of these events is the fact that they are popular discussion topics on social media reaching a viral level. Their lifespan however varies: some catch fire extremely fast but can be soon forgotten while others can be discussed longer in a fluctuating manner. Various studies have investigated the dynamics of viral events in social media [5, 10, 39], the connections between media types in such events [29, 32], or the ways they reflect 'real world' events [1, 3, 30]. In addition, there is plenty of research on viral marketing, meaning how discussion topics can intentionally be made viral in social media [20, 22, 33].

However, due to the nature of most of these studies' topic of interest, topics discussed only in social media are often not identified. Moreover, most of the studies focus only on easily identifiable hashtags [11, 24, 52], are conducted around a single case [12, 13], or explore the unfolding of events with Twitter data only [1, 3, 4, 18, 32, 35], a medium highly characterized by certain discussion themes and types, and popular among distinct audiences [43]. Thus, there is a lack of understanding and a lack of methodological solutions to track viral social media events that may occur across various social media channels (Facebook, Twitter, Instagram, blogs, etc.). With this study, we aim to fill this gap by exploring the lifespan of viral social media events and the prediction power of various discussion features. We ask: Which factors explain the duration of a viral social media event - number of posts, number of authors, or number of social media sources the event has been discussed at?

To answer the research question, we use an existing data set on viral social media events detected by a commercial, proprietary algorithm. The data queried comes from a company that provides access to open Finnish social-media data with the aim of covering as many of the available

This work was supported by Tekes – the Finnish Funding Agency for Innovation (project grant Smarter Social Media Analytics, 4161/31/2016). The authors wish to thank Matti Nelimarkka and the three anonymous reviewers for help and comments, and people at Futusome Oy for providing access to the ViralWatch database and algorithm.

sources and domains as possible. In practice, this stands for public posts on Facebook pages, Twitter, Instagram, thousands of blogs and discussion forums, and news article comments from several Finnish news outlets. The company in question is widely known to be market leader in Finland with regards to data coverage and quality. Therefore, even though we cannot expect the data set to cover all social media discussions related to a given keyword or hashtag, we can expect a reasonable coverage. During the time between October 2014 and March 2017, the algorithm created by the same company had identified 573,858 potential viral social media events, 1335 of which were based on either a hashtag or a keyword and deemed viral enough (exceeding a pre-determined score threshold). From this data set we filtered all events lasting for less than 30 days, which resulted in 960 unique events. We use negative binomial regression to analyze the duration of these events.

This paper proceeds as follows. First, we provide an extensive literature review on virality in social media based various empirical studies. We then move on to describe our data in detail, including a description of the proprietary algorithm. Then, we explain the selections done to filter the dataset and to extract the variables, and present the results of the binomial regression analysis. Finally, we discuss our findings from the cross-media perspective by also pointing out the limitations of our research.

2. Virality and the hybrid media system

The spreading of information in social media has generally been approached using the concept of virality [31, 39, 52]. The concept has its origins in a metaphorical reference to biology and the spreading nature of viruses and microbes. Modern social research, however, has aimed to build more comprehensive and complex approaches that would take into account the complexity of social action, for example by relying on microsociology and network theory. For instance, Sampson [51] emphasizes the contagious nature of events, relations and structurations as a general force how societies come to being. Some scholars emphasize the communicative nature of virality by focusing on the transmission of meaning and affects [47, 52]. From these perspectives, the focus is not only on a particular message or content (i.e., 'a virus'), but rather on a wider perspective to the complex assemblages formed by various actors, meanings, affective attunements and contexts intersecting on various platforms. Therefore, to study virality in the modern media environment, we use the concept 'hybrid media system' [14] to review literature that discusses how various media outlets feed each other and how discussion topics spread.

2.1. Characteristics of viral events

Virality and memetic nature of online content are concepts that have become buzzwords among online content producers and marketers [34]. Viral marketing can be defined as "electronic word-of-mouth whereby some form of marketing message related to a company, brand, or product is transmitted in an exponentially growing way, often through the use of social media applications" [36]. Hence, successful content in the social media era is expected to be both sticky (attracting a large number of users for a maximum amount of time) and spreadable (widely dispersed through networks of users) [34]. Within business and marketing literature, there has been an substantive interest in the influence of social media buzz on business outcomes, as well as in different seeding strategies of generating such buzz. For example, in the case movies, it has been shown that there is a positive relationship between the volume, valence and rate of user-generated content in social media and box-office revenues [1, 17].

The assumed significance generates a need to identify possibly influential events using digital data. A stream of research has explored various ways of identifying viral events from message streams [5, 9, 24, 50] and exploring their features [10, 38]. For example, Becker et al. [4] identified real-world events from Twitter by clustering message features, while Petrovic et al. [50] designed a system to identify emerging new events from streaming Twitter data. These studies use various approaches to identify viral content or events. Most commonly, they use metrics easily available from Twitter meta data, such as number of users, timestamps, urls, or geo-location data. Some studies look more closely at the message contents, such as emoticons or sentiment [8, 21, 46], hashtags [38], or the appearance of certain words [46].

Number of unique users writing about an event is a rather direct indicator of virality. Petrovic et al. [50] show that the number of users that write about an event is a more important indicator than the amount of messages sent. Several studies emphasize the importance of online social networks in information diffusion [2, 9, 23]. A general assumption is that the so-called opinion leaders or influentials are critical in disseminating information to others. These are people who have large social networks and high influence over others [2, 16, 45, 56]. For example, Xu et al. [58] showed that well-connected users with higher issue involvement most likely influence the information flow. In viral marketing campaigns, it has been shown that seeding the content to well-connected individuals is the most successful strategy because they are typically more likely to participate in such campaigns in the first place [33]. Notably, these individuals do not exhibit higher persuasiveness in comparison to the less connected individuals – their influence is merely based on their higher reach and higher likelihood to forward the seeded content. Based on a series of computer simulations, Watts and Dodds [55] suggest that a critical mass of easily influenced individuals are more important in public opinion formation than a minority of influentials.

Another recommended seeding strategy is one that identifies those individuals who connect two otherwise unconnected networks [26, 54]. Bakshy et al. [2] remind, however, that early adopters and active users are different populations, and therefore it is not recommended to focus too heavily on those who are active and loud.

In addition, the content of the message affects which message gets virally spread. While viral marketing literature assumes that the content is related to a company or a brand [36], this is necessarily not the case when looking at what creates the most stir in social media overall. For example, majority of discussion topics on Twitter come from news, as pointed out by Kwak et al. [38]. By investigating trending hashtags, Bruns et al. [10] distinguished between media events, acute events, political events, sports events, keyword hashtags, and meme hashtags. Retweeting behavior and inclusion of an external URL were used as central categorizers. Lindgreen and Vanhamme [42] argue that the element of surprise is central in viral content, while Dobelet et al. [20] highlight the importance of fun and intrigue that capture the recipient's imagination. Eckler and Bolls [22] found that pleasant emotional tone elicit stronger intention to forward a video advertisement than if the tone was negative or coercive. In the news context, Berger and Milkman [7] found that content that evokes high-arousal emotions are more likely to go viral than content that evokes low-arousal or deactivating emotions. These findings challenge the perhaps prevailing assumption that negative messages would more easily go viral than positive messages. Instead, content that arouses feelings, either positive or negative, is more likely to get spread.

From marketing perspective, the longer a (positive) buzz lasts, the better. Twitter has entitled the most discussed hashtag topics on any given day as trending topics, and publishes real-time trending lists globally and regionally based on their full data. The availability of such lists has made the study of trending topics particularly popular [5, 10, 28]. However, viral events vary in their duration. Very peaky and short events can be regarded as bursting events, while long-lasting events are often regarded as long-term trends [24]. According to Kwak and colleagues' study, only 7% of Twitter trends lasted more than 10 days, and one third of the topics lasted one day [39]. Some viral events, hence, are very quickly over, but others remain as a discussion topic for several days.

In general, it is argued that discussion topics of societal relevance follow a specific logic. These "issues" (controversies that have implications for multiple stakeholders) progress through an evolutionary path comprised of a limited number of predictable stages, typically three or four. An issue progresses from a period of insignificance, to a period of heightened awareness, and possibly to a period where new standards and operating procedures to deal with the issue become institutionalized [59]. The focus of this paper is on the "heightened awareness" stage as discussion topics start trending when they create the most interest and discussion gets heated. Social media arguably plays a large role in creating heightened awareness. As said, many of the

discussion topics originate from news but that is not always the case. For example, the United Airlines case originated from a tweet that then spread to other social media platforms and traditional media [53]. Penelopegate, on the other hand, originated from a news article that then spread to various social media platforms [57].

Besides these two examples, however, there are also heated and trending discussion topics that can originate from a social media platform but that do not spread beyond that specific platform. This usually refers to a discussion topic that is either platform-specific or relevant to a defined crowd only. Moreover, there are also discussion topics that initially get discussed on multiple platforms but that do not spread to legacy media (i.e., 'traditional' media). Even though the two aforementioned types of viral social media events are often more well-known than the two latter types, it is still possible that the latter ones fulfill the definitions of virality [36, 39, 51]. Indeed, virality does not mean that *everyone* in a society knows of the viral event, and social media is known to create discussion network heterogeneity, meaning that people get exposed to different kinds of contents and information [37]. In this paper, our interest lies particularly in the significance of cross-platform diffusion within social media from the perspective of event duration.

2.2. Towards hybrid approaches

Within media and communication studies, the interdependence of different media forms has been recently theorized using the concept hybrid media system [14, 15]. In the hybrid media system, older and newer forms of media are intertwined in complex relationships and multiple media logics affect the circulation of content in the mediasphere. Numerous studies have investigated this interdependence, most extensively explored by looking at the diffusion of content between legacy media (e.g., TV, newspapers, radio) and social media (e.g., Facebook, Twitter, blogs). Studies show that legacy media and social media influence and feed one another [28, 48, 49]. In 2010, Kwak et al. [39] found that on Twitter, the vast majority (over 85%) of trending topics are headline news or persistent news in nature.

The influence is, however, bidirectional. Hewett et al. [32] showed that with regards to "brand buzz", the volume and valence of traditional news stories have a positive relationship with the volume and valence of Twitter posts concerning the specific brand. In other words, the more there are news stories about a brand, the more there are also social media posts, and vice versa, and these two media types amplify each other's tone of voice regarding the brand. Further, media events originating from the legacy media coverage have been found to influence the way people communicate with each other online. Using a dataset of 290 million tweets by almost 200 000 politically active Twitter users, Lin et al. [41] found that compared to baseline time periods, media events (pre-scheduled media broadcasts that encourage simultaneous use of traditional and social media) do not only generate abnormal volumes of tweets, but they are also associated with substantial declines in interpersonal

communication and more highly concentrated attention to particular users.

With regards to social media trends and virality, most studies have relied on Twitter data only – mostly due to the lack of access to other types of data in a similar scale. We argue that a hybrid, cross-media perspective, which takes into the account the fact that people are interacting on various social media platforms, and that content is being shared across those platforms, is missing from the existing literature and investigations. Hence, the focal interest of this paper is to study how different characteristic of a viral event, from the amount of users to the variety of social media platforms the event has been discussed at, relate to the duration of the event. The paper now proceeds to describe the empirical data used for the analysis and its results.

3. Data and Method

Our empirical material consists of a database of viral events detected by an algorithm created by a Finnish social media data provider Futusome. The company extensively follows Finnish-language content across various social media platforms at the approximate rate of 500 000 new messages per day and is known to provide the most comprehensive access to Finnish social-media data. Sources include publicly available, Finnish language textual content from public Facebook pages, Twitter, Instagram, YouTube, blogs, discussion forums and news article comments. Hence, their data includes only messages are posted on public arenas, excluding for example private or public Facebook posts made on individual user profile pages. The data and the analytical methods are described next.

3.1. The ViralWatch algorithm

The data is collected using a service called ViralWatch, which was coded and published in January 2015. The full dataset consists of 573,858 potential viral events spanning from October 2014 to March 2017, identified according to the criteria defined in the algorithm.

The authors were given access to the original source code of the ViralWatch algorithm to clarify the criteria used to detect the events. The algorithm was designed by using an example case of a viral social media event that occurred in January 2015. In this event, a Finnish celebrity chef tweeted a praising tweet about a tomato sauce brand, marked with the hashtag #soosi (Finnish for 'sauce'). The tweet turned viral and generated a large deal of discussion both in traditional and social media, a buzz which resulted in a situation where the brand's products were sold out from numerous supermarkets. This example refers to a viral event comparable to the United Airlines case. The trajectory of the spread and occurrence of the #soosi hashtag was used as a reference of a successful viral social media event.

The ViralWatch algorithm works by querying Futusome's database for social media messages posted within the last hour. A list of words (lemmatized to their basic form) that have occurred sufficiently many times and more often

than is expected given their usual occurrence in Finnish language texts is returned. Each term is considered a 'candidate event' and given a score. The time the event was detected is marked as the origin time. The ViralWatch algorithm separately follows each social media platform Futusome crawls and monitors for specific hashtags, keywords (nouns), or unique popular posts through their source IDs. In the case of Twitter, the algorithm only tracks hashtags.

Table 1 presents the conditions under which an event increases its scoring. In the table, let $|P(a, b, t)|$ equal the number of posts containing term t that have been posted between times a and b , where a and b are given as hours after the event's origin. Similarly, let $|U(a, b, t)|$ equal the number of unique users that have used the term; and $RT|(a, b, t)|$ equal the number of retweets (Twitter only). Let c equal a threshold; a single point is given if the scoring condition is satisfied. An event is initially given a score from 0 to 5. The maximum score an event may obtain is therefore 9.

3.2. Data filtering

As the algorithm was generated using a hashtag-based viral event as the ideal measure, we chose to focus on hashtag and keyword-based event types only in our analysis. We expect keywords to follow a similar trending logic as hashtags, while unique popular posts on Facebook, Instagram or YouTube do not. Overall, there were 86,849 hashtag and keyword events in the full database, while the rest consisted of unique popular social media posts and are thus beyond the scope of this study. Further, as the ViralWatch was originally built to start following only events that had received an initial score of 5, we chose to focus on such events only. After this filtering, 1335 events remained. In this respect, we consider the database to extensively represent viral hashtags and keywords denoting a public interest in the Finnish social media ecosystem. Based on manual checks on the detected events with high scores, this assumption seems valid. Further, as we were interested in viral events from the perspective of the hybrid media system and with a cross-platform focus, we used the identified hashtags and keywords to query the full Futusome database (consisting 660 million unique Finnish-language messages) to extract all the posts mentioning the specific keywords or hashtags across all platforms available. Prior to querying, platform-specific information was discarded. This means that if, for example, a keyword had been detected as an event on multiple platforms, these events were combined into a single query. Events were similarly combined temporally.

Next, we operationalized a bursting viral event by introducing a duration limit of 30 days in order to exclude common popular words or long-lasting popular discussions such as elections. For example, Kwak et al. [38] noticed that only a minor share of trending Twitter topics last more than ten days. Thus, the thirty-days perspective was deemed appropriate when analyzing the duration of viral social media events. The data corresponding to each query was filtered to cover a period of 30 days following the origin of the event. If a keyword or a hashtag had been detected

TABLE 1: Event scores and conditions

Scorer	Description	Scoring condition	Thresholds				
			Facebook hashtag	Facebook keyword	Twitter hashtag	Instagram hashtag	Text keyword
Term novelty	Whether the term is a novel one or has been used before a certain number of times.	$ P(-1, 0, t) < c$	2	0.4	0	5	0.4
Initial burst	The strength in the surge of popularity of the term after it has appeared.	$\frac{ P(-1, 0, t) }{ P(-168, -1, t) } < c$	0.4	-	0.4	0.2	-
Initial volume growth	The growth in post volume in a period of time after origin compared to the hour preceding origin.	$\frac{ P(-1, 0, t) }{ P(0, 3, t) } < c$	1	1	1	1	1
Initial retweet volume growth	The growth in retweet volume in a period of time compared to the hour preceding origin (Twitter only).	$\frac{ RT(-1, 0, t) }{ RT(0, 3, t) } < c$	-	-	1	-	-
Initial user count growth	The growth in the number of unique users in a period of time compared to the hour preceding origin.	$\frac{ U(-1, 0, t) }{ U(0, 3, t) } < c$	0.3	0.3, 0.6*	1	0.3	1, 1*
Absolute user count	The number of unique users that have used the term.	$ U(-1, 0, t) < c$	5	5	-	10	5

two or more times on one or more platforms within this 30 day period, it was treated as a single event. If the interval between the detection time was more than 30 days, they were treated as two separate events.

The thirty-day limitation excluded 302 cases (23.9%) from the dataset, leaving 960 events for our empirical analysis. In the analysis, we looked at hashtag and keyword events across several different platforms, namely Facebook, Twitter, Instagram, YouTube, online forums, blogs, and news site comments. The data for these events included 13,777,510 unique social media posts in total. Figure 1 depicts the distribution of these data on different social media platforms. We used custom Python scripts to extract the number of posts, authors, and sources for each day during thirty days including the starting day to study the events. An event was defined to end on the first day during the thirty-day period when there were no posts mentioning the keyword or hashtag in question, a rule that enabled us to determine the duration of each event.

4. Analysis and results

To provide a better understanding of the events under investigation, Table 2 presents descriptive statistics on both hashtag and keyword events. There were 534 hashtag events (56.6% of final data) and 426 keyword events. As keywords are significantly more numerous than hashtags in an average social media post, the chance of one mentioning a trending keyword than mentioning a trending hashtag is higher. Therefore, keyword events lasted on average longer (8.59 days) than hashtag events (2.67 days), and the difference was statistically significant (F=298.38, p<.001). On average, there were 70.02 posts per day for hashtag events and 62.04 posts for keyword events, the difference not being statistically significant. There were more authors per day (33.73)

mentioning the identified keyword than there were authors using the identified hashtag (27.37) with a statistically significant difference (F=9.36, p<.01). Finally, keywords were on average more often mentioned at different social media platforms (here 'sources') (2.40) than hashtags (1.24) and the difference is statistically significant (F=643.48, p<.001). Based on these descriptives, it is concluded that the hashtag and keyword events are inherently different in many regards. Thus, when modeling the duration of a viral social media event, we modeled hashtag events and keyword events separately.

In our analysis, we estimated negative binomial regression models using the number of days a viral social media event had lasted as the dependent variable. The dependent variable consists of non-negative integers and it includes a lot of events that had lasted only one day, making the distribution of the variable skewed. Poisson regression model is generally used to model such non-negative count data but it assumes that the dependent variable follows a Poisson distribution where the mean equals variance [40]. In our case, the mean of event duration for keyword events was 8.59 and variance 55.33, and for hashtag events 2.67 and 5.89, respectively. The over-dispersion problem thus leads us to use negative binomial regression instead, as proposed for example by Greene [27]. Negative binomial regression is a generalized form of Poisson regression, which accounts for over-dispersion.

Independent variables consist of number of posts, authors and sources per day, and the ratio of the second and third day to the first day in terms of posts, authors and sources. The ratios of the second and third day to the first day represent the progress of the event; immediately after the detection, do the events increase or decrease in number of posts, author and sources? These relative measures also enable us to study viral social media events of different

Figure 1: Distribution of data points per platform.

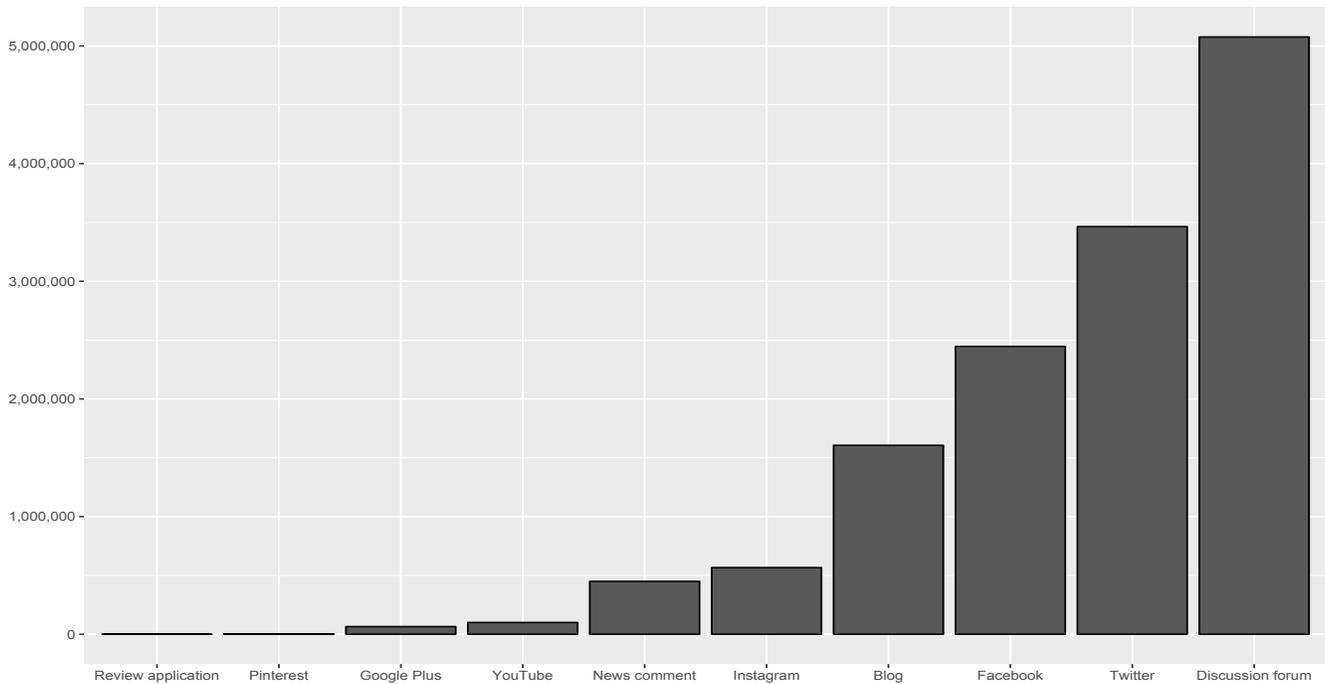


TABLE 2: Descriptive statistics

	Hashtag event (N = 534)	Keyword event (N = 426)
Duration (days)	2.67 (2.43)	8.59 (7.44)
Posts per day	70.02 (129.89)	62.04 (72.23)
Authors per day	27.37 (27.83)	33.73 (36.57)
Sources per day	1.24 (.43)	2.40 (.95)
2nd day posts*	12.1% (29.5)	22.3% (30.4)
3rd day posts*	4.9% (35.7)	9.5% (19.5)
2nd day authors*	16.5% (32.3)	25.0% (28.6)
3rd day authors*	4.9% (20.7)	11.4% (21.1)
2nd day sources*	65.6% (52.0)	82.9% (44.4)
3rd day sources*	30.5% (47.5)	62.8% (48.8)

**Measured as a ratio to the first day's respective variables*

sizes. The biggest event consisted of 6904 posts (duration three days) and the smallest only seven posts (duration two days). Note that both of these events had been detected by the algorithm based on the rules described above. Table 3 displays the regression results for two competing models: in Model 1, there are average number of posts, authors and sources used as predictors. In Model 2, the relative number of posts, authors and sources on the second and third day in relation to the first day are included in the model as main effects.

For both hashtag and keyword events, both of the estimated models were significant as a whole but, in both cases, Model 2 turned out to explain a larger share of the variance in the dependent variable (McFadden Pseudo R-squared .054 for hashtag events, and .046 for keyword

events, indicating a significantly improved data-model fit [44]). Moreover, when comparing the models based on the Log likelihood, the Aikeke Information Criterion (AIC) and Bayesian Information Criterion (BIC) indices, the second models also seem to provide a better fit for the the data. For these measures, lower values are more preferred.

According to the results of Model 2 of hashtag events, the average number sources and the relative number of sources on the second and third day had a positive effect on the duration of the event ($B=.320$, $Exp(B)=1.378$; $pp<.05$; $B=.497$, $Exp(B)=1.643$; $p<.001$; $B=.677$, $Exp(B)=1.969$, $p<.001$). Note that it does not mean the the relative number of sources needed to necessarily increase after the first day to have a positive effect on event duration (even if this happened in some cases); the key is that the relative amount of sources prevails at a high level. Interestingly, neither the average number of authors per day, average number of posts per day nor their relative number of these two variables on the first days of the event had an effect on the hashtag event duration.

For keyword events, the determinants of event duration differed to some extent. In Model 2, the daily average number of posts had, surprisingly, a negative effect on event duration but the effect was very small ($B=-.005$, $Exp(B)=.995$; $p<.01$). As for hashtag events, the average number of sources and the relative number of sources on the third day had a positive effect ($B=.240$, $Exp(B)=1.272$; $p<.001$). The relative number of sources on the second day did not have an effect.

TABLE 3: Regression results

Hashtag events

Model 1

Parameter	<i>B</i>	SE	<i>p</i>	Exp(B)
Average posts	.000	.000	.683	1.000
Average authors	.003	.002	.164	1.003
Average sources	.412	.131	.002	1.511

Fit indices: Log likelihood = -1137.07; AIC = 2282.14; BIC = 2299.26; McFadden Pseudo R^2 = 0.010; Omnibus test: 22.36, df = 3, p = .000

Model 2

Parameter	<i>B</i>	SE	<i>p</i>	Exp(B)
Average posts	-.001	.000	.843	1.000
Average authors	.001	.002	.520	1.001
Average sources	.320	.135	.018	1.378
2nd day posts*	.224	.487	.645	1.251
3rd day posts*	-.084	.295	.775	.919
2nd day authors*	-.095	.430	.825	.909
3rd day authors*	.300	.609	.622	1.350
2nd day sources*	.497	.136	.000	1.643
3rd day sources*	.677	.142	.000	1.969

*Measured as a ratio to the first day's respective variables

Fit indices: Log likelihood = -1086.00; AIC = 2191.99; BIC = 2234.80; McFadden Pseudo R^2 = 0.054; Omnibus test: 124.51, df = 9, p = .000

Keyword events

Model 1

Parameter	<i>B</i>	SE	<i>p</i>	Exp(B)
Average posts	-.008	.002	.000	.992
Average authors	.005	.003	.039	1.005
Average sources	.274	.061	.000	1.316

Fit indices: Log likelihood = -1323.35; AIC = 2654.69; BIC = 2670.91; McFadden Pseudo R^2 = .031; Omnibus test: 85.54, df = 3, p = .000

Model 2

Parameter	<i>B</i>	SE	<i>p</i>	Exp(B)
Average posts	-.005	.002	.001	.995
Average authors	.000	.003	.890	1.000
Average sources	.240	.064	.000	1.272
2nd day posts*	-.085	.647	.895	.918
3rd day posts*	.843	.887	.342	2.324
2nd day authors*	.355	.706	.615	1.426
3rd day authors*	-.351	.831	.673	.704
2nd day sources*	.250	.166	.131	1.284
3rd day sources*	.542	.166	.001	1.720

*Measured as a ratio to the first day's respective variables

Fit indices: Log likelihood = -1303.00; AIC = 2625.99; BIC = 2666.54; McFadden Pseudo R^2 = .046; Omnibus test: 126.24, df = 9, p = .000

5. Discussion

In the present study, we analyzed viral social media events detected by a proprietary algorithm and used a dataset that consisted of events that were based on two different identification strategies (keyword vs. hashtag). The two different strategies produced events that were noticeably distinct, both in terms of what the events looked like (duration, number of authors and number of sources) and in terms of the variables that predicted duration. As the fit indices of the different analyses suggest, the model fits the hashtag data better than the keyword data. However, even when modeling hashtag and keyword events separately, the main result prevails; the higher the number of platforms an event is being discussed at and the higher the relative number of discussion platforms after the first days of the event, the longer the event lasts. For hashtags particularly, the relative growth in the number of platforms where the hashtag is used is the most effective predictor of lengthy virality in social media.

The hybrid media literature and studies on cross-media effects [14, 28, 32, 38] stress that discussions in one media platform are dependent and influenced by other platforms. Thus, according to the literature, discussion topics that are interesting for audiences using different media platforms should last longer, while those that do not attract the attention of the hybrid media system, end fast. This idea is supported by the findings of this study – once discussion topics spread to new platforms, new people get exposed to them. This makes the event last longer. The finding encourages researchers to account for multiple different media platforms

when studying virality as currently a significant share of studies on trending social media topics look inside single platforms, most often Twitter [1, 4, 32, 38]. Simultaneous consideration of different platforms is particularly important in the light of the present findings.

Number of posts or discussants did not have a relationship with event duration. This contradicts some of the existing research findings. For example, Petrovic et al. [50] argued that the number of people discussing a given topic is a better predictor of virality than number of posts. We argue that the same effect is at play in our social media context but the key is the number of platforms rather than the number of people per se. Indeed, single platforms may expose a topic to a very large audience at once but, at that platform, the topic competes with a range of other topics. Therefore, within a single platform, a discussion topic has a high likelihood to die fast. At new platforms, the topic receives a new audience and it is more likely to get discussed again, which adds up to the days the topic is being mentioned. New audiences may also bring new perspectives to the discussion and thereby increase its duration. This interpretation is supported by information diffusion literature [26, 54].

Overall, the current data supports the notion made in previous studies [10, 39] that majority of discussions or events that are abnormally popular at a given time lose their popularity relatively fast. 15.3% of the events identified by the algorithm lasted only one day, meaning that one day after the event was deemed 'viral', there were no mentions of the hashtag or keyword in social media. This finding is naturally dependent on the definition of a viral event but it seems to align with existing ideas on virality in general; a

very small share of all trending discussion topics last long and predicting which topics they might be is very hard [36]. The implications of these results are discussed next.

5.1. Implications

For companies and societal actors the results have several interesting implications. There are various instances that want to get their content spread and understanding the importance of the different platforms is important. In many cases, organizations focus on mastering few platforms and track discussions on those, while the present findings encourage taking a wider perspective and reaching different networks through different platforms. Thus, identifying individual that either connect with different platforms or merely focus on attracting the attention of different people using different platforms is advisable. On the other hand, sometimes organizations do not want discussions about them to spread (think of United Airlines, for example). In those situations, it is good to understand that plain number of posts is not indicative of the duration of the event.

The results also provide implications to optimal information seeding strategies. We support the arguments of Granovetter [26] and Steffes and Burgee [54] as the results suggest that it is essential to get people at different social media platforms to discuss about the topic rather than focus on the overall number of people or overall number of posts mentioning the topic. To achieve this goal, one needs to reach those individuals that connect otherwise non-connected networks and communities. In the present study, we cannot control for the position of the discussants in their network but based on the existing knowledge, the "bridge" strategy seems viable in the case of viral social media events [55]. Notably, some of the other studies on the topic have presented conflicting results [33, 55] but those studies concern situations where an actor intentionally wants to spread the content. In our case, it is not necessarily any specific actor that would benefit from the virality and most often the information spreads without any explicit agendas or strategies.

5.2. Limitations and future research

As any research, this study comes with limitations. The first limitation is related to the database and the given nature of the data. There are arguably many ways to detect viral events or keywords from data streams, and the approach chosen by Futusome is only one and inspired by one extreme hashtag case only. Hence, the data is most likely biased towards certain kinds of events. To reduce the bias, we investigated both hashtag and keyword events. We also chose not to trust the 9-point scoring system generated for the ViralWatch algorithm, and evaluated the success of the identified events using a longer time frame, a cross-platform data collection strategy, and took into account the amount of users and posts in our analysis. However, as pointed out by Choudhury et al. [19], the sampling strategy one uses in social media often affects the results and, therefore,

we encourage future research to use other event detection methods and sampling strategies.

Second, it is extremely difficult to ensure that the posts mentioning a certain keyword really are representative of a certain event or phenomena. As we do not analyze the content of the posts mentioning the hashtags or keywords, it is acknowledged that not all posts discuss the event that was the root cause for virality. With hashtags, it is more likely that people intentionally use them to take part in a certain discussion. In this analysis we aimed to eliminate this problem by using a general restriction of event duration to 30 consecutive days. With this filter we aimed to focus on bursting, peaky events instead of longer trends. Further, this operationalization is likely to leave out long-lasting, exogenous events that are very different kinds of events and often predictable spectacles in the hybrid media system (e.g., elections, Olympic games). However, this restriction is understood to exclude some relevant viral social media events so this creates a limitation to our study.

Third, this analysis is conducted limiting to the Finnish context only. The patterns and durations of viral events might vary in different cultural contexts and in areas of different languages and sizes. Nonetheless, we believe this study is a needed step to better understand the ways how viral events unfold. Also, it is reasonable to expect that many social media-born phenomena are local at least in the beginning, which was the focus of our study. Moreover, Finnish is a very small and distinct language and a great majority of social media posts in Finnish come from Finland. This helps us to understand the context of the events better and the likelihood of posts created by trolls or bots is smaller than when analyzing social media posts in English, for example. Social bots and trolling behavior are more commonly found in languages spoken by very large numbers of people, and in politically sensitive discussion topics [6, 25], none of which is a particular problem in the current context. Future studies in this area should nevertheless focus on tracking global scale events across various platforms. This is a task that requires very sophisticated and powerful services to track and analyze social media content in various languages.

Finally, according to Sampson [51], virality and contagiousness are not only a distinct feature of online environments but a general way of how the world organizes in relational assemblages between different actors. Hence, we must keep in mind that social media events are not born in a vacuum but they reflect some aspects and events in the 'physical' world. By using social media data only it is impossible to fully track all such assemblages and relation between various actors, which still might bear relevant consequences regarding the virality of the events. Therefore, including for example traditional media content and information on the possible consequences of the events to the analysis would further improve our understanding of the dynamics of viral events.

References

- [1] Sitaram Asur and Bernardo A. Huberman. Predicting the Future with Social Media. 3 2010.
- [2] Eytan Bakshy, Brian Karrer, and Lada A. Adamic. Social influence and the diffusion of user-created content. In *Proceedings of the tenth ACM conference on Electronic commerce - EC '09*, page 325, New York, New York, USA, 2009. ACM Press.
- [3] Roja Bandari, Sitaram Asur, and Bernardo a. Huberman. The Pulse of News in Social Media: Forecasting Popularity. 2 2012.
- [4] H Becker, M Naaman, and L Gravano. Beyond Trending Topics: Real-World Event Identification on Twitter. *ICWSM*, 2011.
- [5] Hila Becker, Mor Naaman, and Luis Gravano. Beyond Trending Topics: Real-World Event Identification on Twitter. In *International AAI Conference on Web and Social Media*.
- [6] Leo Benedictus. Invasion of the Troll Armies: From Russian Trump Supporters to Turkish State Stooges. *The Guardian*, 11 2016.
- [7] Jonah Berger and Katherine L. Milkman. What Makes Online Content Viral? *Journal of Marketing Research*, 49(2):192–205, 2012.
- [8] Johan Bollen, Alberto Pepe, and Huina Mao. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. 11 2009.
- [9] Javier Borge-holthoef, Nicola Perra, Bruno Gonçalves, Sandra González-bailón, and Alex Arenas. The dynamics of information-driven coordination phenomena : A transfer entropy analysis. (April), 2016.
- [10] Axel Bruns, Brenda Moon, Avijit Paul, and Felix Münch. Towards a Typology of Hashtag Publics: A Large-Scale Comparative Study of User Engagement across Trending Topics. *Communication Research and Practice*, 2(1):20–46, 2016.
- [11] Axel Bruns and Stefan Stieglitz. Metrics for Understanding Communication on Twitter. In Katrin Weller, Axel Bruns, Jean Burgess, Merja Mahrt, and Cornelius Puschmann, editors, *Twitter and Society*, pages 69–82. Peter Lang, New York, 2014.
- [12] Jean Burgess, Anne Galloway, and Theresa Sauter. Hashtag as Hybrid Forum : The Case of # agchatoz. In Nathan Rambukkana, editor, *Hashtag Publics : The Power and Politics of Discursive Networks*, pages 1–18. Peter Lang, New York, Bern, Berlin, Bruxelles, Frankfurt am Main, Oxford, Wien, 2015.
- [13] Andrew Chadwick. The Political Information Cycle in a Hybrid News System: The British Prime Minister and the "Bullyinggate" Affair. *The International Journal of Press/Politics*, 16(1):3–29, 2011.
- [14] Andrew Chadwick. *The Hybrid Media System. Politics and Power*. Oxford University Press, Oxford, 2013.
- [15] Andrew Chadwick, James Dennis, and Amy P Smith. Politics in the Age of Hybrid Media: Power, Systems, and Media Logics. In Axel Bruns, Gunn Enli, Eli Skogerbø, Anders Olof Larsson, and Christian Christensen, editors, *The Routledge Companion to Social Media and Politics*. London and New York, 2016.
- [16] Kenny K. Chan and Shekhar Misra. Characteristics of the Opinion Leader: A New Dimension. *Journal of Advertising*, 19(3):53–60, 10 1990.
- [17] Pradeep K. Chintagunta, Shyam Gopinath, and Sriram Venkataraman. The Effects of Online User Reviews on Movie Box Office Performance: Accounting for Sequential Rollout and Aggregation Across Local Markets. *Marketing Science*, 29(5):944–957, 9 2010.
- [18] Munmun De Choudhury and Yu-ru Lin Hari. How Does the Data Sampling Strategy Impact the Discovery of Information Diffusion in Social Media ? pages 34–41, 2006.
- [19] Munmun De Choudhury, Yu-Ru Lin, Hari Sundaram, K. Selcuk Candan, Lexing Xie, and Aisling Kelliher. How does the data sampling strategy impact the discovery of information diffusion in social media? *Proceedings of the 4th International AAI Conference on Weblogs and Social Media*, pages 34–41, 2010.
- [20] Angela Dobeles, David Toleman, and Michael Beverland. Controlled infection! Spreading the brand message through viral marketing. *Business Horizons*, 48(2):143–149, 3 2005.
- [21] Peter Sheridan Dodds, Kameron Decker Harris, Isabel M. Kloumann, Catherine A. Bliss, and Christopher M. Danforth. Temporal Patterns of Happiness and Information in a Global Social Network: Hedonometrics and Twitter. *PLoS ONE*, 6(12):e26752, 12 2011.
- [22] Petya Eckler and Paul Bolls. Spreading the Virus. *Journal of Interactive Advertising*, 11(2):1–11, 3 2011.
- [23] M Fazeen, R Dantu, and P Guturu. Identification of leaders, lurkers, associates and spammers in a social network: context-dependent and context-independent approaches. *Social Network Analysis and Mining*, 2011.
- [24] Wei Feng, Chao Zhang, Wei Zhang, Jiawei Han, Jianyong Wang, Charu Aggarwal, and Jianbin Huang. STREAMCUBE: Hierarchical spatio-temporal hashtag clustering for event exploration over the Twitter stream. In *2015 IEEE 31st International Conference on Data Engineering*, pages 1561–1572. IEEE, 4 2015.
- [25] Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. The Rise of Social Bots. *Communications of the ACM*, 59(7):96–104, 2016.
- [26] Mark S. Granovetter. The Strength of Weak Ties. *American Journal of Sociology*, 78(6):1360–1380, 5 1973.
- [27] William H. Greene. *Econometric Analysis*. Prentice Hall, New Jersey, 4th edition edition, 2000.
- [28] Jacob Groshek and Megan Clough Groshek. Agenda Trending: Reciprocity and the Predictive Capacity of Social Networking Sites in Intermedia Agenda Setting across Topics over Time. Technical Report 1, 8 2013.
- [29] L. Guggenheim, S. M. Jang, S. Y. Bae, and W. R. Neuman. The Dynamics of Issue Frame Competition in Traditional and Social Media. *The ANNALS of the American Academy of Political and Social Science*, 659(1):207–224, 5 2015.
- [30] Aniko Hannak, Eric Anderson, Alan Mislove, and Lisa Feldman Barrett. Tweetin ' in the Rain : Exploring societal-scale effects of weather on mood. pages 1–4, 2010.
- [31] Lars Kai Hansen, Adam Arvidsson, Finn Årup Aarup Nielsen, Elanor Colleoni, and Michael Etter. Good Friends, Bad News - Affect and Virality in Twitter. *Future Information Technology, 6th International Conference, FutureTech 2011, Loutraki, Greece, June 28-30, 2011, Proceedings, Part II*, pages 34–43, 1 2011.
- [32] Kelly Hewett, William Rand, Roland T. Rust, and Harald J. van Heerde. Brand Buzz in the Echoverse. *Journal of Marketing*, 80(3):1–24, 5 2016.
- [33] Oliver Hinz, Bernd Skiera, Christian Barrot, and Jan U Becker. Seeding Strategies for Viral Marketing: An Empirical Comparison. *Journal of Marketing*, 75(6):55–71, 11 2011.
- [34] Henry Jenkins, Joshua Benjamin Green, and Sam Ford. *Spreadable media*. New York University Press, New York, 2013.
- [35] Andreas Jungherr. *Analyzing Political Communication with Digital Trace Data*. 2015.
- [36] Andreas M. Kaplan and Michael Haenlein. Two hearts in three-quarter time: How to waltz the social media/viral marketing dance. *Business Horizons*, 54(3):253–263, 5 2011.
- [37] Yonghwan Kim, Shih-Hsien Hsu, and Homero Gil de Zúñiga. Influence of Social Media Use on Discussion Network Heterogeneity and Civic Engagement: The Moderating Role of Personality Traits. *Journal of Communication*, 63(3):498–516, 6 2013.
- [38] Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. What is Twitter, a Social Network or a News Media? In *Proceedings of the 19th International Conference on World Wide Web*, 2010.
- [39] Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. What is Twitter, a Social Network or a News Media ? Categories and Subject Descriptors. *Www 2010*, pages 591–600, 2010.

- [40] Jerald F. Lawless. Negative binomial and mixed poisson regression. *Canadian Journal of Statistics*, 15(3):209–225, 9 1987.
- [41] Yu Ru Lin, Brian Keegan, Drew Margolin, and David Lazer. Rising tides or rising stars?: Dynamics of shared attention on twitter during media events. *PLoS ONE*, 9(5), 2014.
- [42] Adam Lindgreen and Joëlle Vanhamme. Viral marketing: The use of surprise. In Irvine Clarke and Theresa Flaherty, editors, *Advances in electronic marketing*, pages 122–138. Idea Group Pub, 2005.
- [43] Alice E Marwick and Danah Boyd. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience.
- [44] Daniel McFadden. The Measurement of Urban Travel Demand. *Journal of Public Economics*, 3:303–328, 1974.
- [45] Robert K. Merton. The Matthew Effect in Science. *Science*, 159(3810), 1968.
- [46] Nasir Naveed, Thomas Gottron, Jérôme Kunegis, and Arifah Che Alhadi. Bad News Travel Fast : A Content-based Analysis of Interestingness on Twitter. In *WebSci'11*. ACM, 2011.
- [47] Z. Papacharissi. We Have Always Been Social. *Social Media + Society*, 1(1):1–2, 2015.
- [48] Maria Eirini Papadouka, Nicholas Evangelopoulos, and Gabe Ignatow. Agenda setting and active audiences in online coverage of human trafficking. *Information, Communication & Society*, 4462(February):1–18, 2016.
- [49] J. H. Parmelee. The agenda-building function of political tweets. *New Media & Society*, 16(3):434–450, 5 2013.
- [50] S Petrović, M Osborne, and V Lavrenko. Streaming first story detection with application to twitter. *The 2010 Annual Conference of the . . .*, 2010.
- [51] Tony TD Sampson. *Virality: Contagion theory in the age of networks*. University of Minnesota Press, Minneapolis, 2012.
- [52] Sanjay Sharma. Black Twitter? Racial Hashtags, Networks and Contagion. *New Formations*, 78(1):46–64, 2013.
- [53] Lucinda Shen. United Airlines stock drops \$1.4 billion after passenger-removal controversy, 4 2017.
- [54] Erin M. Steffes and Lawrence E. Burgee. Social ties and online word of mouth. *Internet Research*, 19(1):42–59, 1 2009.
- [55] Duncan J. Watts and Peter Sheridan Dodds. Influentials, Networks, and Public Opinion Formation. *Journal of Consumer Research*, 34(4):441–458, 12 2007.
- [56] Gabriel Weimann. The Influentials: Back to the Concept of Opinion Leaders? *Public Opinion Quarterly*, 55(2):267, 22 1991.
- [57] Kim Willsher. François Fillon sinks in polls after ‘Penelopegate’ scandal, 2 2017.
- [58] Weiai Wayne Xu, Yoonmo Sang, Stacy Blasiola, and Han Woo Park. Predicting Opinion Leaders in Twitter Activism Networks. *American Behavioral Scientist*, 58(10):1278–1293, 9 2014.
- [59] Stelios C Zyglidopoulos. The Issue Life-Cycle: Implications for Reputation for Social Performance and Organizational Legitimacy. *Corporate Reputation Review*, 6(1):70–81, 4 2003.