What Triggers Human Remembering of Events? A Large-Scale Analysis of Catalysts for Collective Memory in Wikipedia

Nattiya Kanhabua
L3S Research Center
Leibniz Universität Hannover
Hannover, Germany
kanhabua@L3S.de

Tu Ngoc Nguyen
L3S Research Center
Leibniz Universität Hannover
Hannover, Germany
tunguyen@L3S.de

Claudia Niederée
L3S Research Center
Leibniz Universität Hannover
Hannover, Germany
niederee@L3S.de

ABSTRACT

Going beyond its role as an encyclopedia, Wikipedia becomes a global memory place for high-impact events, such as, natural disasters and manmade incidents, thus influencing collective memory, i.e., the way we remember the past. Due to the importance of collective memory for framing the assessment of new situations, our actions and value systems, its open construction and negotiation in Wikipedia is an important new cultural and societal phenomenon. The analysis of this phenomenon does not only promise new insights in collective memory. It is also an important foundation for technology, which more effectively complements the processes of human forgetting and remembering and better enables us to learn from the past. In this paper, we analyse the long-term dynamics of Wikipedia as a global memory place for high-impact events. This complements existing work in analysing the collective memory negotiation and construction process in Wikipedia directly following the event. In more detail, we are interested in catalysts for reviving memories, i.e., in the fuel that keeps memories of past events alive, interrupting the general trend for fast forgetting. For this purpose, we study the trigger of revisiting behavior for a large set of event pages by exploiting page views and time series analysis, as well as identify of most important catalyst features.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous; J.4 [**Social and Behavioral Sciences**]: [Psychology]

General Terms

Human Factors, Measurement

Keywords

Collective Memory; Social Computing; Time Series Analysis; Wikipedia Page Views; Real-world Events

978-1-4799-5569-5/14/$31.00

1. INTRODUCTION

The way humans forget and remember is a fascinating area of research for both individual and collective remembering. Aspects such as the constructiveness of memories are challenging our intuitive understanding. While forgetting enables us to stay focused and cope with the multitude of our daily experiences, the way past memories are triggered by new experiences is sometimes surprising.

The basis for developing an effective technology that can complement processes in human memory is a deep understanding on how humans remember and forget. Due to its importance for societal processes, it is also important to consider remembering as a crowd phenomenon and investigate what is remembered by communities and societies, e.g., about past events. This is related to the concept of collective memory introduced by Halbwachs [1]. Collective memory is a socially constructed, common image of the past of a community, which frames its understanding and actions. At the same time, collective memory is not static; it is determined by the concerns of the present [1]. With the social Web, the construction and dynamics of collective memory have become an observable phenomenon, which promises new insights. We are especially interested in systematically investigating what triggers (or revives) the memory of past events. Knowledge about such triggering behavior can be used both for recommending related events that are probably remembered by the user, e.g., for enriching a news report about an event and for surprising the reader by reminding the reader of related events she/he (most probably) has forgotten, thus introducing some serendipity.

Web 2.0 offers new rich data sources for a large scale analysis of pattern in human and especially collective remembering and forgetting, which complements qualitative studies from cognitive psychology. One important source for better understanding pattern of collective memory and its construction processes is Wikipedia [2, 3]. The social negotiation and construction processes for example, is reflected by early editing activities on pages referring to events [2, 4] as well as by discussions on the talk pages [3].

In our analysis, we investigate the triggering or reviving of memories of past events using revisiting pattern in English Wikipedia as indicators for what is collectively (actively) remembered and what is rather on the path of forgetting. The content and usage of Wikipedia articles is an important source of information about real-world events [5]. In this study, we focus on exploiting view logs of Wikipedia event pages as the signals of collective memory. From a cognitive point-of-view, access or view logs may not directly reflect how people forget information, e.g., people may remember about an event, but they do not access assets associated to the event. However, we argue that significant patterns found in view activities are a good estimate of public remembering. Such a visit is typically triggered by thinking of this past event and will also refresh the memories on an event by revisiting the information on the respective page. Additionally, analyzing Wikipedia article updates faces scalability issues and this is left for future work.

Generally, individual memories are subject to a forgetting process, which is driven by some form of the forgetting curve first proposed by Ebbinghaus [6]. Especially episodic memory [7], which is responsible for memorizing details of events, is subject to fast forgetting due to interference with memories of new events. Both the effect of proactive interference [8] and retroactive interference [9] make it difficult to remember event details after a while. Various factors can, however, boost human memory of a event or person from one's past, such as, similar events, anniversaries or even a scent. In general, there is a strong relationship between the capability to remember something and the frequency and recency of activating this memory [10]. Such triggering of memories can also be observed for more global events on a cumulative level of communities as the sum of individual remembering re-enforced by information sharing and media coverage. The *2011 nuclear catastrophe in Fukushima* did, for example, trigger the memory of the *Chernobyl event* happened 25 years before raising the Wikipedia event page views from about 9,500 views per day in the first two months of 2011 to up to more than half a million views per day at the time of the Fukushima disaster (around March 15, 2011).

In more detail, we investigate the role of time passed, the type of event, and other factors play in reviving memory. Our work extends the work of [11], who examine collective memory based on its reflection in a newspaper collection, in two directions. Firstly, we analyze the long-term dynamics of collective remembering by looking how forgetting is interrupted by memory revival. This also supplements work on the early memory construction phase in creating Wikipedia articles [2] by looking into long-term temporal development. Secondly, we add an extra perspective by analysing what people actually look at (in Wikipedia), complementing the news coverage perspective of [11].

Our contributions. We analyse over **5500** high-impact events from 11 different event categories. Due to the unique characteristics of every single event especially of the unplanned events, it is very challenging to identify systematic pattern in the revisiting of past events. Therefore, this work just presents a first study in identifying catalysts for event memory triggering.

Using time series analysis, we consider (1) temporal correlations in peaking page visits between events, (2) a surprise score or the residual sum of squares on prediction error, and (3) the skewness of view shapes, as indicators for the capability to act as a catalyst for the memories about the past event. Furthermore, we investigate if there are also other indicators of relationships between the events (e.g., the same types or magnitude of events, same city or country, etc.), by using different features, namely, time, location and impact. To this end, we conduct extensive experiments for identifying promising features.

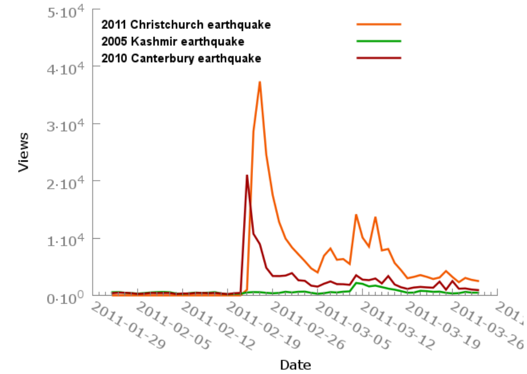

Figure 1: Wikipedia page views triggered by the Christchurch earthquake in February 2011.

2. FORGETTING & MEMORY CATALYSTS

Remembering and forgetting in the context of high-impact events, so called *flashbulb memories*, have been analysed in various studies [12, 13, 14] in cognitive psychology. According to a more recent definition [12], flashbulb memory is "memory about an emotionally impacting event of personal and national importance, which is consequential, socially shared and rehearsed by media". It comprises an autobiographical part, which refers to remembering the personal context, in which one learned about the event and the memory about the event itself. Aspects that have been studied in [14] are the details that people still remember over different periods of time (e.g., 1 week, 11 and 35 months) after the event, the confidence and consistency of their memories over time and the impact of media coverage. However, due to their qualitative nature, those studies are typically limited to a small number of events and a restricted number of users.

Social media analysis has been successfully used in different works for analysing collective attention and awareness [15]. Due to their dynamics, events typically play an important role in such analysis. The transition to analysing remembering of events as a crowd phenomenon relates individual remembering to collective remembering. In social science, the concept of collective memory [16, 1] is used in this context. It refers to the collectively constructed image (memory) of the past, which is shared by a community and, roughly speaking, used by the community for framing their current understanding and activities.

The Web and especially the social Web have a high impact on collective remembering [3]. Due to its popularity as an information reference and the easy and long-term accessibility of information about an event, Wikipedia is a promising subject for analysing collective remembering. In addition to the access numbers, the importance that is assigned to Wikipedia as an information reference for event information is confirmed by the high level of community involvement reflected in the number of editors (19 million registered users and about 30 thousand active editors[1] in English Wikipedia), the fast reflection of new events in Wikipedia [4], and the conflicts and *edit wars* that can be observed on controversial topics. Although religious and political topics are most dominant in edit wars, there is also a considerable number of events in the top-10 lists of controversial topics extracted from Wikipedia in different languages in [17].

[1]Editors with more than 5 edits per month

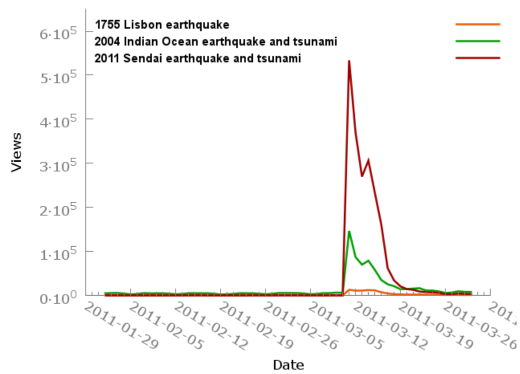

Figure 2: Wikipedia page views triggered by the tsunami in Japan in March 2011.

Figure 1 shows the Wikipedia views of the event page for the earthquake in Christchurch, New Zealand in February 2011 (as a triggering event) and compares it with the view number of two other earthquakes, namely, the earthquake in Canterbury in September 2010 and the large earthquake in Kashmir in October 2005. The strong peak in the views of the Canterbury earthquake around February 2011 suggests a strong influence of the Christchurch earthquake as a catalyst for remembering the Canterbury earthquake. This strong influence can be explained by the facts that a) both earthquakes happened in the same region and b) there is a time gap of just five months between the two events. In contrast, memory for the Kashmir earthquake, which is more distant in time and location, seems to be revived to a much lesser degree by the Christchurch earthquake.

Figure 2 shows page views for the event page of the Japan tsunami in 2011 as the triggering event and views for the page of the Indian Ocean earthquake and tsunami in 2004.

The increasing view numbers suggest that the event in Japan acts as a catalyst for remembering the 2004 Tsunami and the earthquake in Canterbury in September 2010 does also for the event pages of both earthquakes in New Zealand when taking a closer look to Figure 1. Interestingly, there is an increase even for an earthquake, which lays far more in the past, like the Lisbon earthquake in 1755 shown in Figure 2. Of course, an increased number of Wikipedia views is only an indirect signal of memory revival for the considered event. However, we believe that a person, who visits an event page from a past event at least thinks of the event, which brings it back to active memory. Furthermore, visiting a Wikipedia event page on purpose will typically result also in reading some information about the event, such as, refreshing or extending the information memorized about the event.

3. ANALYSING COLLECTIVE MEMORY

In this paper, we leverage the page view statistics of English Wikipedia in analysing the collective memory of a past event. An event e can be represented by a Wikipedia event page, with starting time e_{st} and ending time e_{et}, and it consists of other information, such as, location, impact (e.g., magnitude, fatality and the cost of damage). The time series X of e is created using the aggregated number of daily views of its corresponding Wikipedia page. In the rest of this section, we present our methodological approach for detecting the reviving of memory of past events, which helps in identifying the catalysts for such remembering.

3.1 Remembering Score

For identifying the reviving of memory of past events, we exploit remembering signals based on the event time series and three time series analysis techniques, i.e., cross-correlation coefficient, sum of squared error, and skewness.

1) **Cross-correlation coefficient** (CCF) is a statistical method to estimate how variables are related at different time lags. That is, the CCF value at time t between two time series X and Y indicates the correlation of the first series with respect to the second series shifted by a time amount t, e.g., in days or weeks. A common measure for the correlation is the Pearson product-moment correlation coefficient. The CCF between two time series describes the normalized cross covariance and can be computed as:

$$CCF(X,Y) = \frac{\sum_{i=1}^{N}(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^2}\sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \bar{y})^2}}$$

where x_i and y_i are values at time t_i of X and Y, $\bar{x}$ and $\bar{y}$ are the means values, and σ_X and σ_Y are the standard deviations. In our case, the time series X and Y are corresponding to the time series of two events respectively. The CCF function has values between -1 and +1, where the value ranges from 1 for perfectly correlated results, becomes 0 when there is no correlation, decreases to -1 when the results are perfectly correlated negatively. This measurement can be interpreted as the similarity between two time series in volume, with consideration of time shifts. Hence, the CCF value reflects the remembering of a past event with respect to a studied event.

2) **Sum of squared error** (SSE) is a measure of the accuracy of short-term forecasts for time series data. It has been mentioned in [18] that an unplanned event happened when there is a significant error in the residuals of its predictive model. Our intuition of employing this feature is as follows. Two co-peaking events are not necessarily correlated with each other, e.g., a past event occurs to have an anniversary and simply peaks around the same time. We consider SSE or prediction error is a good feature for surprise detection, rather than just looking at co-peaking. In this work, we use the HoltWinters prediction model computed as the residual sum of squares on prediction errors. The score implies how *unplanned* the value of time series is at the time period of interest. Given a time series $Y = \{y_1, ..., y_t\}$, a predictive model and its fitted values $V = \{v_1, ..., v_t\}$, a SSE score at time t is calculated by:

$$SSE(Y,V) = \sum_{i=1}^{t}(y_i - v_i)^2$$

3) **Kurtosis** is a basic statistic method to measure the *skewness* of a distribution. Intuitively, we seek for a signal of relatedness other than co-peaking. Inspired by the work [19], we use this feature to capture the reviving of past events by considering how much of the probability distribution is contained in the peaks, and how much in the low-probability regions. Kurtosis is calculated as the average deviations of the time series elements to the forth power divided by the standard deviation to the forth power (refer to [19] for more detailed computation).

From the list of candidates for related events, we apply certain filtering criteria to leave out those events with insignificant behavior. We categorize these criteria into two distinct classes: **burstiness** (relied on burst detection) and

basic time series statistics, i.e., *min*, *max* and *average frequency*. In burstiness filtering, we define *overlapRatio*, which is the ratio between the number of bursts occur overlapped with the event period and the total number of bursts within 14 days before/after the event. The time window can be varied. Note that, this burst entropy is an indicator of how *random* is the activeness of the event in the studied period. The other techniques in this class are derived from *burst strength* and *burst duration*. Finally, the **filtering** score is determined as: $Filter = \sum_{f \in \mathcal{F}} \phi \cdot f$, where each feature f comprises the set of filtering techniques $\mathcal{F}$, and ϕ is a mixture parameter.

We compute *remembering* scores based on three main signals, i.e., CCF, SSE and Kurtosis, for quantifying the remembering of past events. The higher the values of CCF and SSE, the better the past events are remembered. On the contrary, Kurtosis values must be low such that a smaller peak around the mean has high remembering scores. Our remembering functions are defined as: (1) a combination of the signals (denoted *a basic remembering score*), and (2) a filtering method applied to a basic remembering score (denoted *a filtered remembering score*). The basic remembering score is calculated as:

$$Remembering = \alpha \cdot CCF + \beta \cdot SSE + \gamma \cdot Kurtosis$$

where the value of each individual feature will be normalized using two methods: (1) zscore - normalize each feature by its mean/standard deviation, and (2) linear - normalize each feature by its min/max values. We report the best results obtained from both normalization techniques.

To this end, a filtered remembering score can computed in two ways: (1) mixture of *remembering* and *filtering* scores, and (2) the multiplication of the two scores.

3.2 Features for Triggered Remembering

3.2.1 Temporal Similarity

We compute the temporal similarity between two events by taking into account a time distance. We adopt the *TSU* metric, proposed in [20], that measures the similarity between a temporal query and a document based on their temporal metadata, i.e., temporal expression(s) in the given query and document timestamps. *TSU* copes with the time uncertainty related to the two entities by relying on a decay function. In our case, we consider the time period of an event; thus $TSU(e_i, e_j)$ between two events e_i and e_j can be computed as follows:

$$TSU(e_i, e_j) = \frac{1}{2} \times \left(DecayRate^{\lambda \cdot \frac{|st_i - st_j|}{\mu}} + DecayRate^{\lambda \cdot \frac{|et_i - et_j|}{\mu}} \right)$$

where $DecayRate$ and λ are constants, $0 < DecayRate < 1$ and $\lambda > 0$. st and et are the starting time and ending time of the event e. μ is the unit of time distance between the two events, e.g., one week or 3 months. In addition to TSU, we also use the *time difference* of two events, i.e., an absolute distance in days, months, or years, as our temporal features.

3.2.2 Location Similarity

We extracted locations where events occurred, as described in its corresponding Wikipedia article, from our annotated dataset. Similar to [21], we create a geographic hierarchy of event locations as follows: *city* $\rightarrow$ *state* $\rightarrow$ *country* $\rightarrow$ *neighbor countries* $\rightarrow$ *continent*. Because our event dataset consists of mostly *high impact* ones, we consider *city* as our finest granularity, and focus more on the *country* level. For instance, if a flood event happened in Thailand (e.g., floods in Thailand 2011), events that took place in the nearby region (floods in Vietnam 2008) will be accounted for our location similarity metric. To obtain higher quality of location expressions (from our annotated dataset, which contains Wikipedia text snippets), we further use Stanford NER[2] for geographical entity extraction. Such location expressions are often short and missing information, we *fully* disambiguate and enrich them (to cover all upper levels) by looking up into data dumps provided by GeoNames[3] (e.g., Chicago ▶ [Chicago (city), Illinois (state), United State (country), Canada, Mexico, Cuba (neighbor countries), North America (continent)]. We define a location similarity metric (based on the Jaccard similarity coefficient) by assigning geographical *weights* to elements in the set. According to our geographic hierarchy, we give higher weights to the lower levels and lower weights to the upper levels. In detail, we give weights in the scale from 1 to 4, where 4 is assigned to *city*, 3 to *state*, 2 to *country* and 1 to *neighbor countries* and *continent*. The location similarity score is given by the weighted Jaccard similarity between two enriched-location sets.

3.2.3 Impact of Events

The impact score of an event is measured based on the impact it causes in different aspects. The aspects include *damaged area*, *damaged properties*, *the cost of damage*, *magnitude* (for earthquake events), *highest winds*, *lowest pressure* (for Atlantic hurricanes) and *fatalities*. The information derived from each aspects are quantified, normalized and consequently accumulated in the final results. We categorize the aspects into two different types: *fatalities* (number of deaths) and *other* (the rest of the aspect), as our empirical studies shown that fatalities often induce people's remembering.

4. EXPERIMENT

4.1 Experimental Settings

We analysed over **5,500** high-impact events from 11 different event categories depicted in Table 1. For computing the temporal similarity, we experimented with 6 different time granularities, i.e., 1 day, 7 days, 1 month, 6 months, 1 year and 10 years, where $DecayRate = 0.5$, $\lambda = 0.5$ and μ is varied according to different granularities defined above. We employed the open source implementation of burst detection by CISHELL[4]) using its default parameters. For the aggregation methods, we set the parameters for *remembering* scores as α is 0.5, β is 0.4 and γ is 0.1, where these values were empirically determined. When applying our filtering technique, we weighed each filtering feature equally, i.e., giving 0.2 to the mixture parameter ϕ for all features. The time series parameters: a time window in days w, lags in days l, and a smoothing sm, are $w \in \{7, 14\}$, $l \in \{3, 7\}$, and $sm = 1$ respectively and these parameters will be studied for their performance in the experiments.

Metrics. In our experiment, we measured the association between two measured quantities *remembering* scores and the proposed catalyst features, i.e., temporal similarity and location-based similarity using different correlation metrics: Pearson product-moment correlation coefficient, Spearman's

[2] http://nlp.stanford.edu/software/CRF-NER.shtml

[3] http://www.geonames.org/export/

[4] http://wiki.cns.iu.edu/display/CISHELL/Burst+Detection

Table 1: Statistics of event categories with time spanning from earliest dates to *October 2013*.

Category	#Events	#Triggers	Earliest Date
Atlantic hurricane	654	134	1900-08-27
Aviation accidents	787	146	1912-05-13
Civil wars	*78*	*7*	*793*
Earthquakes	468	119	426 BC
Floods	*114*	*78*	*1897-04-01*
Mass murder	1136	344	1897-05-27
Pacific typhoon	*253*	*68*	*1944-12-17*
Terrorist incidents	727	295	1950-04-01
Tsunamis	*49*	*5*	*1700-01-26*
Volcanic events	*11*	*7*	*1815-01-01*
Wildfires	74	44	1970-09-26
Total	**4351**	**1247**	

rank correlation coefficient, and Kendall tau rank correlation coefficient. The first coefficient is a measure of linear correlation, whereas the latter two metrics measure rank correlation statistics. Correlation coefficient measures the statistical correlation between two variables, which ranges from 1 for perfectly correlated results, through 0 when there is no correlation, to -1 when the results are perfectly correlated negatively. As observed empirically, Spearman's rank correlation coefficient provides better results than the other two metrics. Hence, we will only report the results for this correlation metric.

4.2 Experimental Results

The expectation in analysing the triggering of related events was that there is a clear correlation with the type of events, time and location. Roughly speaking, when event e happened, people would remember events that are of the same type as e, happened nearby and/or in the recent past of e. Our analysis has, however, shown that there are no such clear pattern. This is partly due to a variety of factors including the unique characteristic of each event, the uneven distribution of unplanned events in space and time, and the dominating influence of very large events. Therefore, our experiments just give first insights into how event memory is triggered as part of collective memory.

One of the factors to be considered is the number of events available in individual categories. Table 1 shows the data set statistics of each event category used in our study. We performed our experiments for the listed 11 event categories. For five of the categories - shown in italics in the table - the number of triggering events as well as the number of events that could be triggered was too low for making any reliable statements. Therefore, we only present results for the remaining six categories in this paper.

For better understanding the impact of the temporal and spatial distribution of the events in the individual categories on event memory triggering, we first take a look on this distribution. For some of the considered event categories, Figure 3 shows the number of events in each year (for the last 100 years) and the distribution of the most frequent locations. The temporal distribution shows in all cases a strong focus on more recent events. This is at least partially due to the development of Wikipedia, which leads to a better coverage of events with the increasing popularity of Wikipedia. For older events, we can observe the typical pattern in collective memory that events with higher impact are better remembered in the collective memory and thus also more readily a-posteriori documented in Wikipedia.

For **Aviation accidents**, for example, there is a high variation in the numbers for individual years, highlighting the random character of those events. The numbers are more evenly distributed for the event categories **Earthquakes** and **Atlantic hurricane**. For the spatial distribution, we see a more or less random distribution for **Aviation accidents** with majority of events in the location group *Others*. In contrast, other events categories, such as, **Earthquakes** and **Terrorist incidents** show the focus on the typical critical regions. For **Wildfires** (not shown) this effect is even stronger. We can see a clear focus on US and Australia (together 79% of all the events). Similarly, **Atlantic hurricane** show a focus on the typical hurricane regions.

As described in Section 3, we conducted several intuitive filtering methods to exclude trivial events from the remembering analysis. To have a clearer view of the possible effects of these filtering on the remembering score, we give an example of one of the chosen methods. Table 2 shows the top remembered events which are triggered by *Hurricane Sandy* event, filtered by the maximum number of views per day. The results are three different ranked lists based on three manually defined thresholds: at least 100, 500 and 1000 as maximum number of views, respectively. *Hurricane Inez* and *1856 Last Island hurricane* which appear high in the first list, are left out when the threshold is increased (from left to right). These events have high remembering score, yet seem to be not publicly attentive (the number of daily views never gets over 200). More interesting events (e.g., *Hurricane Andrew*) are boosted higher in the latter lists.

To address how impact features (i.e., location, time and size) exert influence on collective memory, we present a qualitative analysis on several spotlights from 11 studied categories. The following case studies describe a close examination of some of the experimental results on triggering events throughout different categories. Figure 4 depicts the distribution of Atlantic hurricane events triggered by *Hurricane Sandy (2012)*, and *Hurricane Hanna (2008)*, from top to bottom respectively, with regard to three dimensions: location, time and remembering score. In general, location and time contribute low effect on remembering scores for events in this category. However, the events with significantly peaked scores have clear location similarities with the triggering event, and their happening time is close to the time of the trigger. Figure 5 holds up this claim by delineating top-10 events triggered by the two events. *Hurricane Gustav* is the freshest hurricane toward *Hurricane Hanna* and struck at around the area of Puerto Rico and East Coast of the US. By contrast, *Hurricane Sandy* commemorates old hurricanes decades ago, but location is still a strong indication of the remembering. One interesting finding is that both *Hurricane Sandy* and its triggered *1991 Perfect Storm* were initially formed around Canada areas. Note that, the mentioned events are high-impact (most destructive and costly).

The second category of events that we want to take closer look to are **Aviation accidents**. The *Qantas Flight 32* accident in 2010, for which the top-10 list of remembered events is shown on the top of Figure 7, is a good example on the mix of impact factors that trigger the past remembering. On the first glance there seems to be no clear pattern besides that nearby events are better remembered. This impact is even stronger than it seems, since *Qantas Flight 30* and *British Airways Flight 9* both were on the way to Australia, when the accident happened. A closer look on the individual candidates shows that some of the accidents are

Table 2: Filtering by the maximum number of daily view effects on the ranking of top remembered events for the triggering event Hurricane Sandy.

event	max > 100		event	max > 500		event	max > 1000	
	#view	remember		#view	remember		#view	remember
Hurricane Katrina	106551	0.84	Hurricane Katrina	106551	0.84	Hurricane Katrina	106551	0.81
1991 Perfect Storm	71092	0.54	1991 Perfect Storm	71092	0.52	1991 Perfect Storm	71092	0.51
Great Hurricane of 1780	11492	0.47	Great Hurricane of 1780	11492	0.47	Great Hurricane of 1780	11492	0.45
Hurricane Inez	220	0.45	Hurricane Donna	8565	0.43	Hurricane Donna	8565	0.42
1856 Last Island hurricane	143	0.45	Hurricane Mitch	10026	0.43	Hurricane Mitch	10026	0.40
Hurricane Donna	8565	0.44	Hurricane Frederic	804	0.43	Hurricane Juan	1443	0.39
Hurricane Mitch	10026	0.44	Hurricane Georges	1571	0.43	Hurricane Georges	1571	0.39
Hurricane Isaac (2000)	439	0.44	Hurricane Charley (1986)	620	0.42	Hurricane Andrew	28511	0.39
Hurricane Nicole (1998)	150	0.44	Hurricane Gustav	1576	0.41	Hurricane Gustav	1576	0.38
Hurricane Frederic	804	0.43	Hurricane Alicia	1030	0.41	Hurricane Alicia	1030	0.37
Hurricane Claudette (2003)	174	0.43	Hurricane Juan	1443	0.41	Hurricane Gilbert	4351	0.37
Hurricane Georges	1571	0.43	Hurricane Lili	513	0.41	Hurricane Isaac (2012)	11351	0.37
Hurricane Hilda	146	0.43	Hurricane Andrew	28511	0.40	Hurricane Wilma	17496	0.36
Hurricane Omar (2008)	169	0.43	Hurricane Faith	748	0.40	Hurricane Frances	1708	0.36
Hurricane Bret (1999)	145	0.43	Hurricane Alex (2010)	689	0.40	Hurricane Hugo	9655	0.36

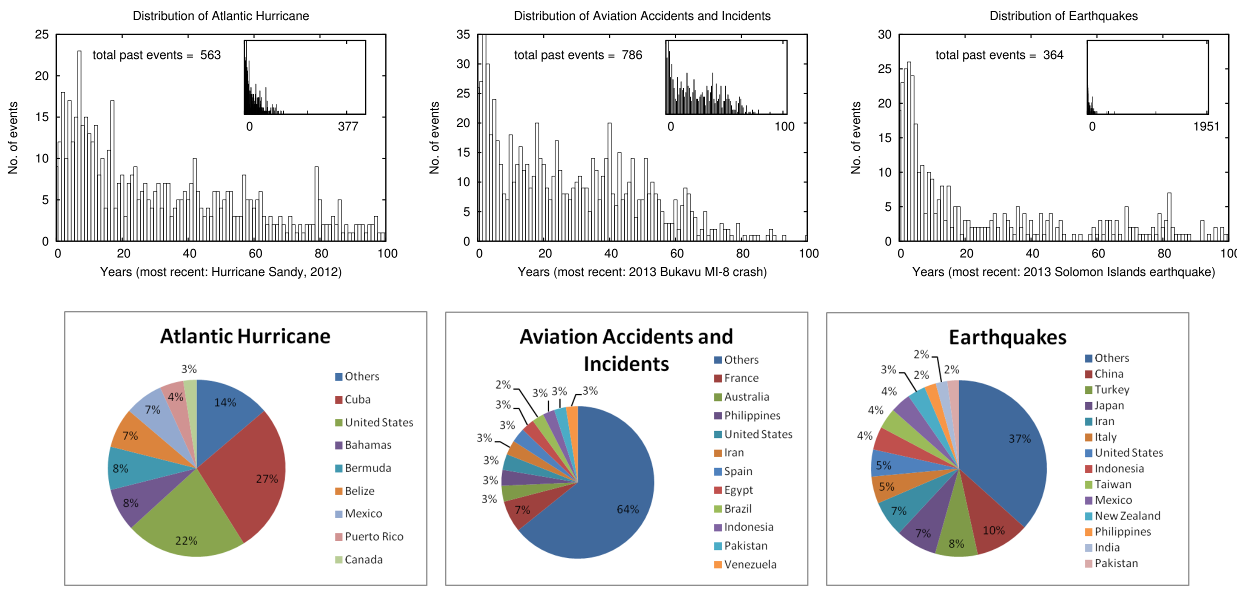

Figure 3: Distributions of highlighted category events over two dimensions: time and location (top to bottom).

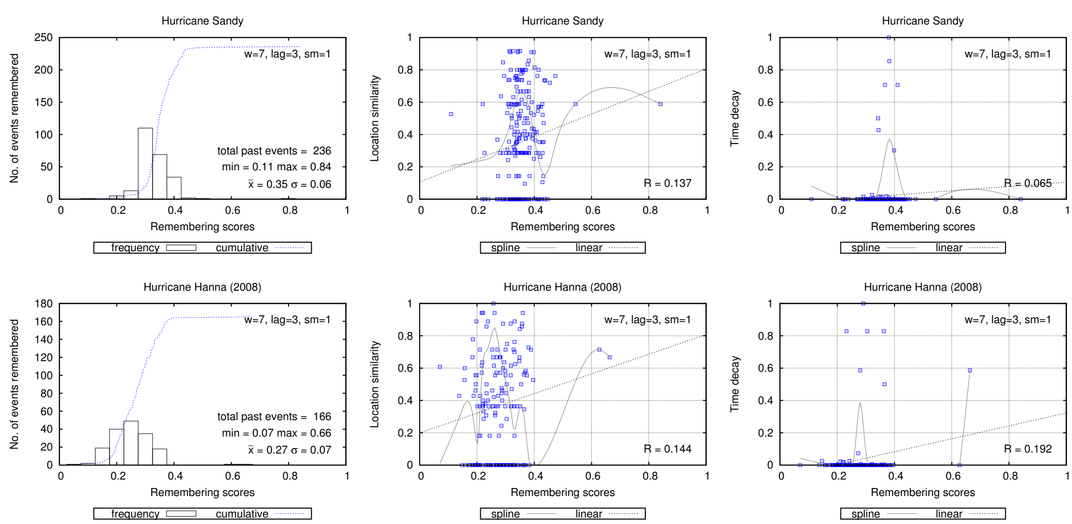

Figure 4: Results for Hurricane Sandy, 2012 (top) and Hurricane Hanna, 2008 (bottom): (left to right) Distribution of remembering scores, the correlation of remembering scores vs. location and time. Spline lines approximate the interpolation of data points weighted by 3.0.

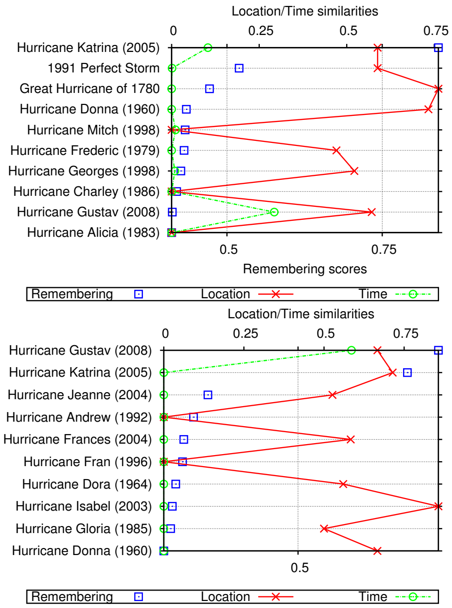

Figure 5: Lists of top-10 past events triggered by remembering of Hurricane Sandy (top) and Hurricane Hanna (bottom).

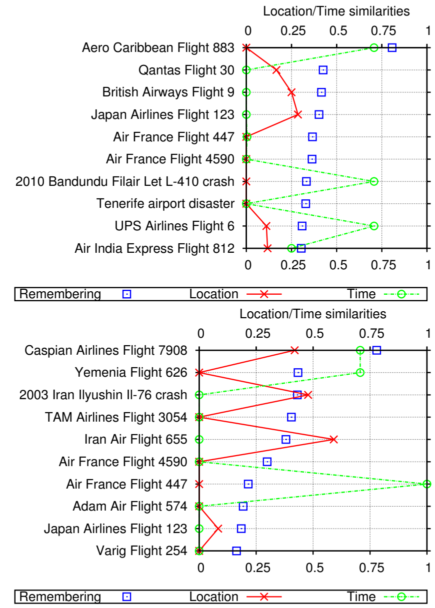

Figure 7: Lists of top-10 past events triggered by remembering of Qantas Flight 32 (top) and Aria Air Flight 1525 (bottom).

probably remembered, since they happened recently before the observed accident, such as, *Aero Caribbean Flight 883 (2010)*, *Qantas Flight 30 (2008)*, which is also the same airline, and *Air France Flight 447 (2009)*. Others are probably remembered because they happened spatially close-by, such as, *Japan Airline Flight 123 (1985)*. *Air France Flight 4590 (2000)* is neither temporally nor spatially very close. However, this was the Concorde accident, which had in general very high visibility and is, thus, also very strongly remembered. It also appears in the top-10 list of the second flight analysed in Figure 7 (appeared bottom). A similar situation occurs for the *Tenerife Airport disaster (1977)*, which is classified as "deadliest accident in aviation history" in Wikipedia, because two aircraft collided. A similar pattern can be observed for the second analysed flight *Aria Air Flight 1525*, an Iranian flight accident in 2009. Although in this case, the impact of location seems to be stronger.

For the event category **Earthquakes**, we want to discuss here an interesting series of events in more detail. These are the *2010 Canterbury earthquake*, the *2011 Christchurch earthquake* and the *June 2011 Christchurch earthquake*. All three of the earthquakes took place close to each other partially affecting the same city, i.e., Christchurch. Furthermore, they also happened to take place in the same time frame, with just some months between the individual events. If we consider the events independent from each other we might expect a very similar set of triggered event memories for all three of them. Figure 8, however, shows a different picture. Of course, the later events have the predecessor event(s) in their top ranked list (close in time and place).

For the first event in the series, the *2010 Canterbury earthquake*, the top ranked events are a recent high-impact event (*2010 Haiti earthquake*) and two close-by events. The rest of the list are high impact and historical earthquakes. In contrast, the *2011 Christchurch earthquake* in February shows a much stronger locality focus. For this second event in the series, people seem to be interested much more in the previous events in the same region. Recent events and high-impact events outside the region are in the minority in the top-10 list. For the third event *June 2011 Christchurch earthquake*, the remembered events are dominated by the two predecessor events. The remembering score drops very quickly after those two events. Within the other remembered events there are mainly historical events and previous high-impact events, of which only one, the *2010 Haiti earthquake* is recent.

In summary, the results suggest that recent events in the same region are good candidates to be remembered. In addition to location and high-impact of earthquakes, recent events that are not close-by do not play a very important role in event memory triggering. Rather there is an interest in high impact events from the past including historical earthquakes. This pattern can also be seen for other earthquakes, which we analysed as triggering events. However, the results also suggest that, for a full analysis, it might also be necessary to look beyond single events, especially, if there are several events in temporal and local proximity.

The last category of events considered here in more detail are **Terrorist incidents**. We have selected four spotlight events from this category for discussion: the *2009 UN guest house attack in Kabul*, the *19 September 2010 Baghdad bombings*, the *9 September 2012 Iraq attacks* and the *28 October Peshawar bombing*. The first and the last event in the list show a quite similar characteristic with the airplane accidents, where the top-10 list is a mix of events very close in time to the triggering event, in near-by places or a generally remembered high-impact event, such as, the *Pan Am Flight 103* event (Lockerbie bombing). The top-10 lists for the second and the third event are however rather surprising, because the remembered events in the lists are neither temporally nor spatially related to the triggering events.

For the *2010 Baghdad bombing*, nearly all triggered events are in the US, the top-ranked events being related to the *September 11 attacks*. This might be the effect of cultural bias of English Wikipedia that we are using as the basis: Terrorism events happening in Iraq are linked to terrorist events in the US rather than to other bombings in Iraq, since the

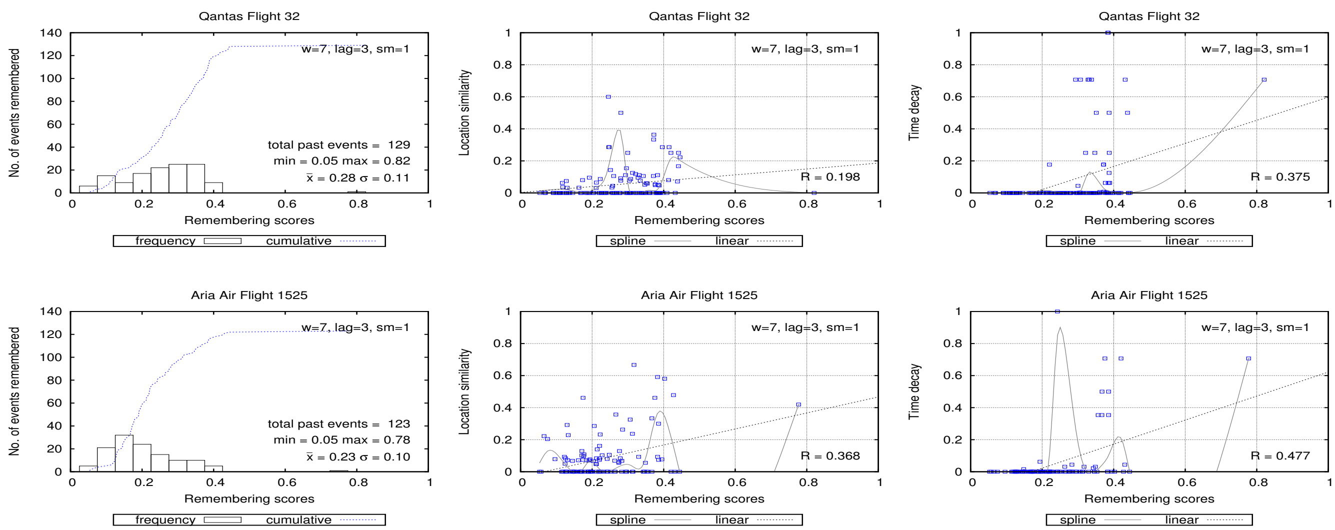

Figure 6: Results for Qantas Flight 32 (top) and Aria Air Flight 1525 (bottom): (left to right) Distribution of remembering scores, the correlation of remembering scores vs. location and time. Spline lines approximate the interpolation of data points weighted by 3.0.

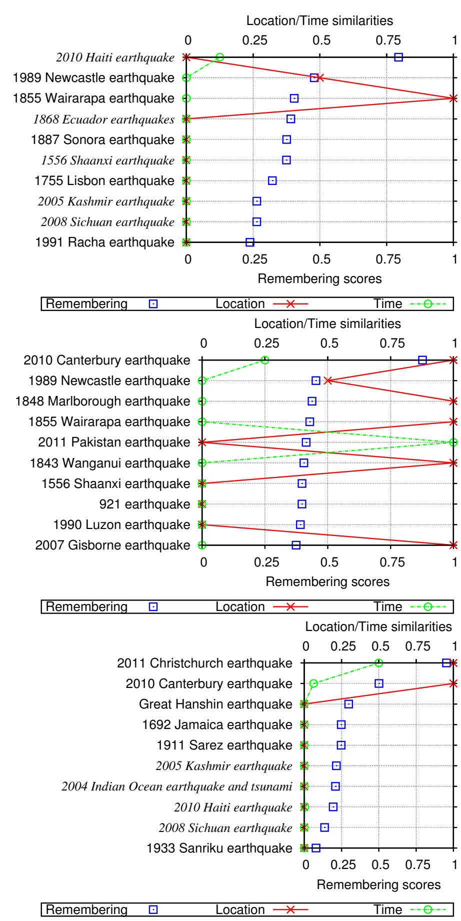

Figure 8: Lists of top-10 past events triggered by remembering of 2010 Canterbury earthquake (top), 2011 Christchurch earthquake (middle) and June 2011 Christchurch earthquake (bottom).

US events are possibly stronger in the typical use of English Wikipedia. The third event the *9 September Iraq attacks* exhibits a similar pattern again with a strong impact on remembering of the *September 11 attacks*.

One interesting observation is that **semantic similarity** between events, beyond a shared category, plays an important role for event triggering in the terrorist attack. E.g., the *June 2012 Kaduna church bombings* (not depicted) triggers the remembering of other religion terror attacks, such as, the *Sarin gas attack* in Tokyo subways (2nd ranked), the *Grand Mosque Seizure* (5th ranked), and the *16th Street Baptist Church bombing* (24th ranked). Another example is the

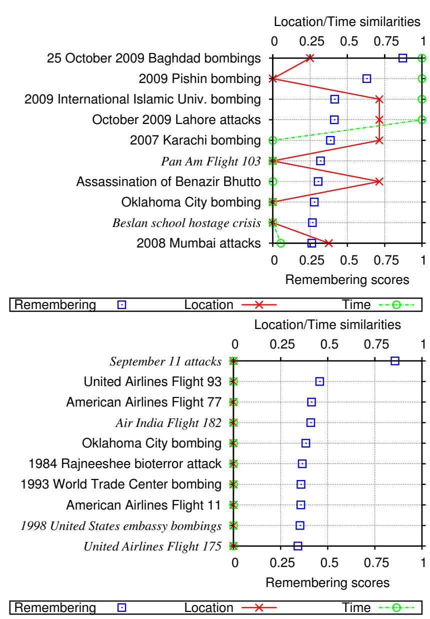

Figure 9: Lists of top-10 past events triggered by remembering of 2009 UN guest house attack in Kabul (top) and 19 September 2010 Baghdad bombings (bottom).

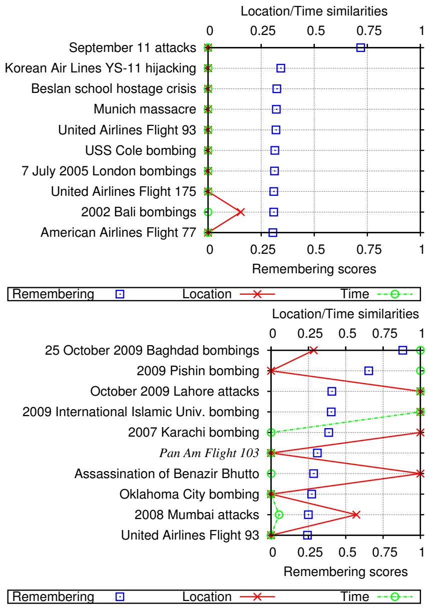

Figure 10: Lists of top-10 past events triggered by remembering of 9 September 2012 Iraq attacks (top) and 28 October 2009 Peshawar bombing (bottom).

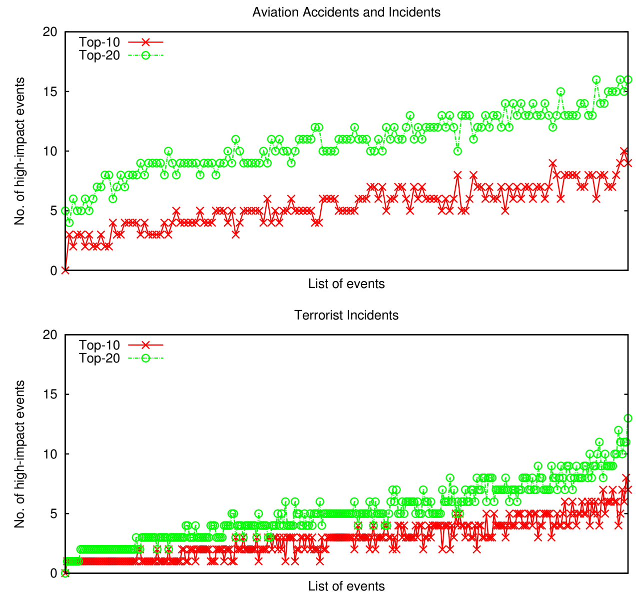

Figure 11: Distributions of high-impact events in the top-10 and top-20 past events triggered by remembering of the list of selected events for two categories.

2008 Mumbai attacks trigger the memory of "terror attacks in business, entertainment or hotel areas", e.g., the *Islamabad Marriott Hotel bombing* (2nd ranked), the *2002 Bali bombings* (7th ranked), and the *Moscow theater hostage crisis* (15th ranked). This finding suggests that besides location and time, semantic similarity between events also influences, which events are remembered. The finer granular classification of events might provide additional factors, although it is crucial to find similar subclasses to those that the human brain uses to associate events with each other.

Influence of High-impact Events. As discussed previously, high-impact events tend to be remembered although they are not strongly correlated in time or place with the considered event. Therefore, we investigated how the impact of an event or its size influences the collective remembering. Our assumption is that events with high impacts, e.g., in magnitudes or perceived damages in terms of casualties or costs are likely to be remembered, regardless of temporal or location distances. For this purpose, we analyse the number of high impact events, which appear in top-10 and top-20 ranked lists of two highlighted categories, namely, **Aviation accidents** (with higher density), and **Terrorist incidents** (with low density), see Figure 11. For assessing event impacts, we identify the reported number of fatalities using a cut-off threshold. Based on the available data, we set the thresholds to be 30 and 100 for both categories. It can be seen that high-impact events comprise between 25% to 50% in the top-10 past events triggered by the studied events. Within the top-20 past events, the percentage of high-impact events is higher, ranging mostly between 25% and 75%. Thus, we conclude that the impact of events is an important factor for remembering events, besides the closeness in time and location.

5. RELATED WORK

So far little work has been done on analysing Wikipedia as a *global memory place* [3]. Most of the work on Wikipedia and collective memory [2, 4] focuses on the early phase of capturing the events in Wikipedia, which is characterized by negotiation and sense making processes. In [4], for example, the collaborative creation of breaking news articles is analysed from a behavioral science perspective, and in [2] the use of language pattern and inferred psychological processes in documenting disasters. In [22], Ciglan and Nørvåg proposed to detect events by analysing trends in page view statistics. In [5], Georgescu et al. extracted event-related information from Wikipedia updates for a given entity using burst detection, temporal and textual features.

With respect to collective memory, the work by Au Yeung and Jatowt presented in [11] is most related to our approach. There, the authors analysed references to the past (as an indicator to what is remembered) in a large news collection from different countries for identifying, which years are most frequently referenced. Our approach differs in two main aspects. Firstly, we rely on actual information access instead of news references for identifying, what is remembered. Secondly, we perform a more systematic analysis of catalysts for triggering memory, which goes beyond the analysis of what is actually remembered in [11].

Our approach is also related to work on analysing peaks of collective attention in other Social Media such as Twitter in the goals and applied methods. In [15], event-related peaks in Twitter are analysed relating their content to temporal activity profiles, especially clustering the fraction of tweets before, during and after the peak.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we studied catalysts for revisiting memories of past events based on Wikipedia view logs. The purpose of this analysis was an improved understanding of collective memory as it is collaboratively constructed in Wikipedia. Identifying pattern of past memory triggering has proven to be more complicated than expected due to the noise and multitude of signals in view logs, due to the multitude of event types in Wikipedia, due to the unique characteristics of every single event and due to the multitude of possible reasons for revisiting a page of a past event.

In spite of this, we managed to identify some first pattern for event memory triggering for diverse event types including natural and manmade disasters as well as accidents and terrorism. For doing this we have combined correlation detection, analysis of the *surprise* aspect (unexpected change) in the distribution of the past event surrounding the peak time of the triggering event and analysis of the skewness of the distribution of the past event at the peak time of the triggering event. Our analysis confirmed the influence of closeness in time and location, but also has shown that these aspects cannot be considered in isolation and that high-impact events and semantic similarity of events also influences, which event memories are triggered by an event.

In our future work, we plan to deepen our systematic analysis of factors for revisiting past events and of the combination of those factors. We also plan to consider more features for the identification of memory catalysts and to verify the predictive qualities of such features in larger experiments. Furthermore, we also plan to investigate external factors for observed memory *revivals* such as media coverage linking new events to past events or reflection of such relationships in other types of social media and how to combine them with our Wikipedia-based analysis.

Acknowledgments This work was partially funded by the European Commission for the FP7 project ForgetIT and the ERC Advanced Grant ALEXANDRIA under the grant numbers 600826 and 339233, respectively.

7. REFERENCES

[1] M. Halbwachs, *On collective memory*. Chicago: The University of Chicago Press, 1950 (Translation).

[2] M. Ferron and P. Massa, “Psychological processes underlying wikipedia representations of natural and manmade disasters,” in *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration*, WikiSym ’12, pp. 2:1–2:10, 2012.

[3] C. Pentzold, “Fixing the floating gap: The online encyclopaedia wikipedia as a global memory place,” *Memory Studies*, vol. 2, no. 2, pp. 255–272, 2009.

[4] B. Keegan, D. Gergle, and N. Contractor, “Hot off the wiki - structures and dynamics of Wikipedia’s coverage of breaking news events,” *American Behavioral Scientist*, vol. 57, no. 5, pp. 595–622, 2013.

[5] M. Georgescu, N. Kanhabua, D. Krause, W. Nejdl, and S. Siersdorfer, “Extracting event-related information from article updates in wikipedia,” in *Proceedings of the 35th European Conference on IR Research*, ECIR ’13, pp. 254–266, 2013.

[6] H. Ebbinghaus, *Über das Gedächtnis. Untersuchungen zur experimentellen Psychologie*. Duncker & Humblot, Leipzig, 1885.

[7] E. Tulving, “Episodic memory: From mind to brain,” *Annual review of psychology*, vol. 53, no. 1, pp. 1–25, 2002.

[8] B. Underwood, “Interference and forgetting,” *Psychological Review*, vol. 64, no. 1, pp. 49–60, 1957.

[9] J. A. McGeoch and W. T. McDonald, “Meaningful relation and retroactive inhibition,” *American Journal of Psychology*, vol. 43, no. 4, pp. 579–588, 1931.

[10] J. R. Anderson and L. J. Schooler, “Reflections of the environment in memory,” *Psychological Science*, vol. 6, no. 2, pp. 396–408, 1991.

[11] C.-m. Au Yeung and A. Jatowt, “Studying how the past is remembered: Towards computational history through large scale text mining,” in *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*, CIKM ’11, pp. 1231–1240, 2011.

[12] E. Coluccia, C. Bianco, and M. A. Brandimonte, “Autobiographical and event memories for surprising and unsurprising events,” *Applied Cognitive Psychology*, vol. 24, no. 2, pp. 177–199, 2010.

[13] A. R. A. Conway, L. J. Skitka, J. A. Hemmerich, and T. C. Kershaw, “Flashbulb memory for 11 september 2001,” *Applied Cognitive Psychology*, vol. 23, no. 5, pp. 605–623, 2009.

[14] W. Hirst, E. Phelps, R. Buckner, A. Budson, A. Cuc, J. Gabrieli, M. Johnson, C. Lustig, K. Lyle, M. Mather, R. Meksin, K. Mitchell, K. N. Ochsner, D. Schacter, J. Simons, and C. Valdya, “Long-term memory for the terrorist attack of september 11: Flashbulb memories, event memories, and the factors that influence their retention,” *Journal of Experimental Psychology: General*, vol. 138, no. 2, pp. 161–76, 2009.

[15] J. Lehmann, B. Gonçalves, J. J. Ramasco, and C. Cattuto, “Dynamical classes of collective attention in twitter,” in *Proceedings of the 21st International Conference on World Wide Web*, WWW ’12, pp. 251–260, 2012.

[16] J. Assmann and J. Czaplicka, “Collective memory and cultural identity,” *New German Critique*, no. 65, pp. 125–133, 1995.

[17] T. Yasseri, S. Anselm, M. Graham, and K. Janos, “The most controversial topics in wikipedia: A multilingual and geographical analysis,” in *Global Wikipedia: International and cross-cultural issues in online collaboration*, Scarecrow Press, 2014.

[18] K. Radinsky, K. Svore, S. Dumais, J. Teevan, A. Bocharov, and E. Horvitz, “Modeling and predicting behavioral dynamics on the web,” in *Proceedings of the 21st International Conference on World Wide Web*, WWW ’12, pp. 599–608, 2012.

[19] F. Diaz and R. Jones, “Using temporal profiles of queries for precision prediction,” in *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR ’04, pp. 18–24, 2004.

[20] N. Kanhabua and K. Nørvåg, “Determining time of queries for re-ranking search results,” in *Proceedings of the 14th European Conference on Research and Advanced Technology for Digital Libraries*, ECDL ’10, pp. 261–272, 2010.

[21] J. Strötgen, M. Gertz, and C. Junghans, “An event-centric model for multilingual document similarity,” in *Proceedings of the 34th ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR ’11, pp. 953–962, 2011.

[22] M. Ciglan and K. Nørvåg, “WikiPop: Personalized event detection system based on wikipedia page view statistics,” in *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, CIKM ’10, pp. 1931–1932, 2010.