

Analysis and Forecasting of Trending Topics in Online Media Streams

Tim Althoff Damian Borth Jörn Hees Andreas Dengel

German Research Center for Artificial Intelligence (DFKI)
D-67663 Kaiserslautern, Germany

althoff@stanford.edu, {damian.borth, joern.hees, andreas.dengel}@dfki.de

ABSTRACT

Among the vast information available on the web, social media streams capture what people currently pay attention to and how they feel about certain topics. Awareness of such trending topics plays a crucial role in multimedia systems such as trend aware recommendation and automatic vocabulary selection for video concept detection systems.

Correctly utilizing trending topics requires a better understanding of their various characteristics in different social media streams. To this end, we present the first comprehensive study across three major online and social media streams, Twitter, Google, and Wikipedia, covering thousands of trending topics during an observation period of an entire year. Our results indicate that depending on one's requirements one does not necessarily have to turn to Twitter for information about current events and that some media streams strongly emphasize content of specific categories. As our second key contribution, we further present a novel approach for the challenging task of forecasting the life cycle of trending topics in the very moment they emerge. Our fully automated approach is based on a nearest neighbor forecasting technique exploiting our assumption that semantically similar topics exhibit similar behavior.

We demonstrate on a large-scale dataset of Wikipedia page view statistics that forecasts by the proposed approach are about 9-48k views closer to the actual viewing statistics compared to baseline methods and achieve a mean average percentage error of 45-19% for time periods of up to 14 days.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval — *Information filtering*

Keywords

Trending Topics, Social Media Analysis, Lifecycle Forecast, Twitter, Google, Wikipedia

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

MM'13, October 21–25, 2013, Barcelona, Spain.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2404-5/13/10 ...\$15.00.

<http://dx.doi.org/10.1145/2502081.2502117>

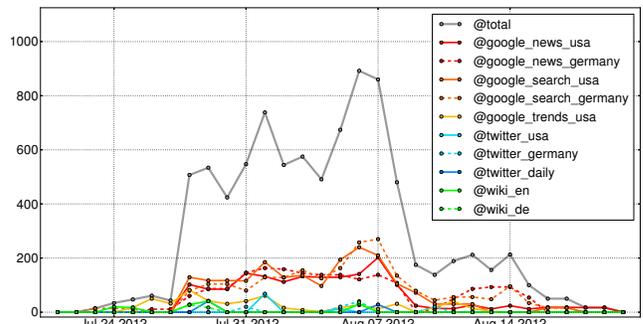


Figure 1: The trending topic “Olympics 2012” during summer 2012. The colored curves represent the contribution of the different online and social media streams.

1. INTRODUCTION

Currently, as multimedia content consumption is rapidly increasing on the internet we can consider ourselves as living in the zettabyte era¹. To a large degree this vast amount of information is created, shared, and exchanged by the users themselves in social multimedia systems. These systems provide us with streams of information that capture how we are spending our time and what we are talking about. Over time, different topics arise in these media streams reflecting changing interests of groups of individuals. In this sense, social media streams mirror our society. These streams contain rich information and immediate feedback about to what people currently pay attention to and how they feel about certain topics.

Such *trending topics* [13] can be associated with a subject (i.e. a textual label) and experience sudden increases in user popularity (“trending”). They usually correspond to real world events such as sport events (Olympics 2012), product releases (iPhone 5), celebrity news (Steve Jobs’ death), incidents (Sinking of the Costa Concordia), political movements (Occupy), and entertainment (Academy Awards). Awareness of trending topics plays a key role in building social multimedia systems satisfying users’ information needs and enhancing the multimedia consumption experience [17].

Today, a large variety of social multimedia systems is available to users on the web such as Twitter, YouTube, Facebook, Flickr, Pinterest, Tumblr, Google News/Trends, and Wikipedia serving different kinds of needs such as information demand, social communication, as well as sharing

¹Cisco, The Zettabyte Era. May 30, 2012.

and consumption of multimedia content. Because of this, many users turn to multiple of these online platforms, depending on their present individual needs, creating a heterogeneous multi-channel² environment within the social media landscape. While researchers have a general intuition about the nature of these channels (e.g. many current event detection systems are based on Twitter feeds [14, 22]), studies of the exact distribution of information across multiple channels have not been in the scope of the multimedia community until now. To the best of our knowledge, there does not exist a study on the behavior and lifecycle of trending topics across multiple media channels and on the correlations between topic categories and media channels over a one year observation period prior to this work. For example, the trending topic “olympics 2012” (the number one trend) manifests itself in multiple media channels which clearly exhibit different behavior (as illustrated in Figure 1).

A second and much more challenging task is *forecasting* the life cycle of a trending topic, i.e. forecasting the amount of user engagement towards it in the very moment they emerge. These forecasts enable us to anticipate changing information needs of users ahead of time and to allocate limited resources in trend aware multimedia systems such as recommender systems [9, 17] or video concept detection systems [5]). For example, such limited resources include time of content consumers and producers, screen estate and content recommendations on web pages [17], processing time for training models in video concept detection [5], and server load and replication capacity [21].

While obviously of great use, forecasting trending topics is a very challenging problem since the corresponding time series usually exhibit highly irregular behavior also known as structural breaks that can lead to large forecasting errors and unreliable models [7]. The complexity of this task will be illustrated in more detail in Sec. 4 (Figure 5).

We present a novel forecasting approach exploiting the assumption that semantically similar topics exhibit similar behavior in a nearest neighbor framework. Semantically similar topics are discovered by mining topics of similar type and category on DBPedia [3], a structured representation of the information on Wikipedia. We evaluate our fully automatic approach on a large-scale dataset of billions of views of several million articles on Wikipedia.

Summarizing, our contribution in this paper is three-fold. We present **(1)** the first comprehensive study of trending topics in three major social and online media channels with an observation period of one year, **(2)** a fully automatic forecasting technique for these trending topics based on a nearest neighbor approach exploiting semantic relationships between topics, and **(3)** an empirical evaluation of our forecasting model over real-world user behavior on a large-scale Wikipedia dataset.

The rest of this paper is organized as follows. In Sec. 2, we outline work related to trending topics in multimedia and social media based forecasting. The multi-channel analysis of trending topics is presented in Sec. 3. Our novel forecasting technique is described in Sec. 4 before evaluating the approach in Sec. 5. Sec. 6 concludes this paper with a short summary and an outline for future work.

²Note that we choose to use the term *channel* instead of *stream* as we believe it to fit the archive-like properties of multimedia collections such as YouTube, Flickr, and Wikipedia better.

2. RELATED WORK

The overview over related work is divided into three parts. We first report work that analyzes or combines information from multiple media channels. Then, we focus on forecasting behavioral dynamics before finally presenting applications of trending topics to social multimedia systems.

Multi-Channel Analyses.

Detecting and tracking topics in news media has been studied since 1997 with the goal of understanding news in multiple languages and across multiple media channels (including television and radio sources) [2].

Ratkiewicz et al. [16] argue that online popularity of specific topics (e.g. “Barack Obama”) cannot in general be characterized by the behavior of individual news-driven events (e.g. “Barack Obama inaugurated as U.S. President”) since the former might subsume many different news stories. They further find bursts in Wikipedia traffic to be correlated with bursts of Google search volume.

Wikipedia article views have also been correlated with behavior on Twitter to analyze the effectiveness of creating new content for breaking news (e.g. Japan earthquake) or updating pages at the moment of news (e.g. Oscar winners) [22] and to filter spurious events on Twitter which indirectly resulted in the finding that Wikipedia lags behind Twitter by about two hours [14]. Adar et al. correlate several behavioral datasets such as general queries from an observation period of one month from MSN and AOL search logs with the general idea of understanding past behavior to predict future behavior [1].

The work of Yang and Leskovec [23] explores patterns of temporal variation in online media by clustering them to analyze the rise and fall of human attention with regard to certain phrases or memes. They report six distinct temporal patterns and find that weblogs trailed mainstream media by one or two hours for most of the considered phrases.

In contrast to the above, we explicitly focus our study on trending topics in online and social media channels over a long observation period. Similar to [8] using trending topics such as “oil spill” and “iPhone” to evaluate tweet selection, we utilize trending topics published by the media channels themselves and focus on multi-channel analysis and forecast. We further analyze the correlation between topic categories and media channels and propose a fully automatic approach of forecasting trending topics in terms of time series (as opposed to visualization tools [1] or predicting temporal clusters [23]).

Forecasting Behavioral Dynamics.

Much research has been devoted to predicting economic variables such as auto sales, unemployment claims [6] (“now-casting”), or opening weekend box-office revenue for movies [10]. In the same way, popularity of online content has often been treated as a single variable (e.g. total number of views) instead of a time series. Either early popularity [19], or content-based features such as publisher, subjectivity, or occurrence of named entities [4] are used to forecast eventual popularity. More recent work treats popularity of queries and clicked URLs of search engines as time series and uses state space models adapted from physics for forecasting [15].

The only work on forecasting Wikipedia article popularity known to the authors of this paper restricts itself to the case

of featured articles on the main page and accounts for daily cycles in viewing behavior [20]. However, trending topics lead to time series with unexpected shifts. These shifts are known as structural breaks in the field of econometrics and can lead to large forecasting errors and unreliable forecasting models [7]. A good example for such a structural break is Whitney Houston’s death in February 2011 that caused 2000 times (!) more people to access her Wikipedia article than usual. Clements and Hendry find that structural breaks (also described as parameter non-constancy) are the main cause of large, unexpected forecast errors in practice. Pooling or combining forecasts from different models has often been found to outperform individual forecasts and to increase forecast robustness.

Forecasting trending topics is different in the sense that we require a fully automatic system and that often there is little historical information available (unlike for many economic variables of interest). For example, few people were aware of “Costa Concordia” before the ship sunk in January, 2012. Similarly, to make predictions about “54th Grammy Awards” one needs to understand the relationship of this event with previous ones such as previous instances of the Grammy Awards. We assume that semantically similar topics share characteristics and behavior and therefore could improve forecasting accuracy. This assumption has not yet been explored in previous work (e.g. [15]). In addition, our proposed approach can forecast popularity of arbitrary articles that exhibit very diverse viewing dynamics. This differs dramatically from the experimental conditions in [20] which are characterized by the authors as “nearly ideal”.

Multimedia Applications.

One application using social media in our target domain of social multimedia systems has been explored in [12]. This work uses the Flickr photo upload volume of specific topics to inform autoregressive nowcasting models for monthly political election results and product sales (i.e. the model requires the Flickr upload volume at time t to produce a forecast for time t). Further, the Flickr queries relevant to the forecast subject of interest are chosen manually (e.g. using “Hillary” instead of “Clinton” to avoid images by Bill Clinton for the 2008 Democratic Party presidential primaries).

Another application that explicitly focuses on trending topics identified in different online and social media channels uses them to dynamically form and extend the concept vocabulary of video concept detection systems [5]. They further find trending topics to be strongly correlated with changes in upload volume on YouTube and demonstrate a visual classification of these trends.

SocialTransfer is another system that uses trending topics obtained from a stream of Twitter posts for social trend aware video recommendation [17]. Learning new associations between videos based on current trends is found to be important to improve the performance of multimedia applications, such as video recommendation in terms of topical relevance and popularity. The popularity of videos is also used in [21] to drive the allocation of replication capacity to serve social video contents. This work analyzes real-world video propagation patterns and finds that newly generated and shared videos are the ones that tend to attract the most attention (called temporal locality). They further formulate the challenge to estimate the videos’ popularity for video service allocation for which they use the number of microblog

Topic	Topic
1. olympics 2012	11. christmas
2. champions league	12. steve jobs
3. iphone 5	13. manhattan
4. whitney houston	14. academy awards
5. mega millions numbers	15. formula 1
6. closer kate middleton	16. justin bieber
7. facebook	17. joe paterno died
8. costa concordia	18. battlefield 3
9. black friday deals	19. muammar gaddafi dead
10. superbowl	20. ufc

Table 1: International Top 20 Trending Topics during Sep. 2011 – Sep. 2012. Please note that the US presidential election was in Nov. 2012.

posts that share or re-share the video. These insights are incorporated into the design of a propagation-based social-aware replication framework. Two other research prototypes that seek to enhance the multimedia consumption experience by extracting trending topics and events from user behavior on the Web are SocialSensor [9] and TrendMiner [18]. The first one emphasizes the real-time aspects of multimedia indexing and search over multiple social networks for the purpose of social recommendations and retrieval. TrendMiner focuses on real-time methods for cross-lingual mining and summarization of large-scale stream media and use cases in financial decision support and political and monitoring.

3. MULTI-CHANNEL ANALYSIS

This section introduces the dataset used for the analysis of trending topics across media channels and presents an analysis of their temporal characteristics as well as insights into the relationship between channels and topic categories.

3.1 Trending Topics Dataset

To facilitate a comprehensive analysis of trending topics across multiple channels, we crawl the top trends from three major online media channels on a daily basis. With the intent to capture people’s communication needs, search patterns, and information demand we choose to retrieve trends from Twitter, Google, and Wikipedia, respectively. We further cluster similar trends across time *and* channels to later be able to compare the different manifestations of a particular trend across multiple channels.

1. Raw Trending Topic Sources.

As motivated above, we use Google, Twitter and Wikipedia as a starting point by retrieving a ranked list of popular terms from 10 different sources on a daily basis: five Google channels (Search and News for USA and Germany as well as the Trends feed), three Twitter channels (daily trends for USA and Germany as well as the Daily Trends feed), and two Wikipedia channels (popular articles in the English or German language). For each of these feeds we retrieve 10-20 ranked topics (110 topics per day). In total the dataset covers the observation period Sep. 2011 – Sep. 2012 and contains roughly 40,000 potentially overlapping topics.

2. Unification and Clustering of Trends.

To connect multiple instances of the same trending topic across time and media channels we perform a unification and clustering of the individual topics obtained in the pre-

	#Topics	#Sequences
Total	200	516
Google	191 (95.5%)	445 (86.2%)
Twitter	118 (59.0%)	232 (45.0%)
Wikipedia	69 (34.5%)	108 (20.9%)
G & T	115 (57.5%)	174 (33.7%)
G & W	66 (33.0%)	86 (16.7%)
T & W	43 (21.5%)	57 (11.0%)
G & T & W	42 (21.0%)	48 (9.3%)

Table 2: Number of analyzed trend topics and sequences for Google (G), Twitter (T), and Wikipedia (W).

vious step. First, we map the individual topic strings to a corresponding Wikipedia URI by selecting the top-most Wikipedia result when performing a Google search for that topic (this approach was found to be more robust than more direct methods on Wikipedia [5]). Then we cluster two URIs together if their Levenshtein or edit distance is below a certain threshold (set to $0.35 \times$ word length). This process can be thought of as a greedy version of a hierarchical agglomerative clustering using the Levenshtein distance. The method allows us to unify topics such as “super bowl time”, “super bowl 2012”, and “superbowl” into a single cluster. Overall this results in 2986 clusters or individual trends. A cluster is now represented by its most prominent member and will be referred to as a *trending topic* for the rest of this work.

3. Ranking Trending Topics.

To reason about the popularity of these trending topics we assign them trend scores based on the following method: For each each day and for each of our 10 feeds, we record the rank at which a topic appears. These ranks are combined using Borda count, obtaining a score for each day that is assigned to the topic’s cluster (from the previous step). To measure the impact of a trending topic (cluster) over its overall lifetime, we define its global trend score by the sum of its daily scores over the observation period. The top trends with respect to this global trend score are shown in Table 1.

3.2 Lifetime Analysis

Some trending topics occur multiple times within our one year observation period. For example sport topics such as “Champions League” are discussed whenever important games are played. To allow for an analysis of when channels begin and stop featuring particular topics, we divide a trending topic into multiple sequences. For the Champions League example, we obtain one sequence per round of the tournament. Formally, topics are broken up into sequences such that their daily trend scores are non-zero for at least two out of any three adjacent dates, i.e. we compensate for “score gaps” of at most one day but form multiple sequences for gaps of two or more days. We perform an analysis of the lifetime of trending topics in the different channels for the top 200 trends (based on their global trend score) that were split into 516 sequences based on the described procedure. Table 2 summarizes the resulting number of topics and sequences for the different channels and their combinations. Note that we map the ten individual observed channels to their respective sources: Google, Twitter, and Wikipedia. The Google channel has the largest coverage of the trend-

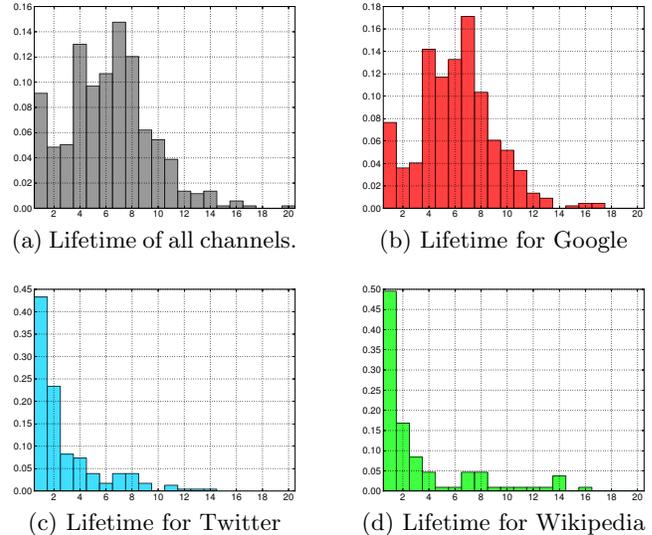


Figure 2: Lifetime in days for top 200 trending topics.

ing topic sequences in our dataset (86.2%). About 9.3% of trending topic sequences occur in all three channels and between 11.0% and 33.7% occur in two channels.

We first deal with the question how long trending topics survive in general and whether there are differences in lifetime in the different media channels. In our case, lifetime is defined as the number of consecutive days with positive trend scores. Histograms of the lifetime of trending topics are shown in Figure 2. We observe that the trending topics in our dataset rarely survive for more than fourteen days with most trending topics having a lifetime of less than nine days. Since Google covers a large share of top trends, the distribution for the channel looks very similar to the overall distribution. The lifetime of topics on Twitter is much shorter complying with our expectations of the ephemerality of trends in this channel – about two thirds of the top trends only survive for one or two days. Interestingly, the distribution looks similar for Wikipedia. This raises the question when and why exactly people are turning to Wikipedia to satisfy their information needs.

For trending topics that occur in at least two of the three channels we record the day on which the trend starts, the day on which the trend peaks (defined by the trend score defined earlier), and the day on which the trend ends again. We then define the *delay* between two channels as the difference between the start/peak/end dates in channel X and the start/peak/end days in channel Y. Note that a positive delay means that the first channel X is slower, i.e. trends tend to start later in channel X than in channel Y. The mean delays for start, peak, and end are summarized in Table 3. Interestingly, there are only marginal differences in starting delays (first of the three numbers) between the three channels with Twitter and Wikipedia being slightly faster than Google. These results are not very surprising considering that Osborne et al. found Twitter to be around two hours faster than Wikipedia [14] – a difference almost impossible to observe in data of daily granularity. A much stronger effect is observed for peak and end delays (second and third number). Both Twitter and Wikipedia tend to peak more

	Twitter	Wikipedia
Google	-0.20 / 2.23 / 3.64	-0.41 / 2.35 / 4.16
Twitter	-	0.07 / 0.36 / 0.21
Google Trends	0.09 / 0.04 / 0.41	-0.42 / -0.32 / 1.25

Table 3: Mean delay in days between pairs of media channels (start/peak/end). Positive delay means that the “row channel” is slower than the “column channel”.

than two days before Google. The picture is even clearer when looking at the end delays where Twitter and Wikipedia lead Google by three and four days respectively. Overall, these results add to the ephemerality of the Twitter and Wikipedia. However, through this analysis we found that some of the Google channels seem to be intentionally delayed or averaged, i.e. still displaying the top stories after several days. Since Google Trends is known not to be artificially delayed we added comparisons with this channel to Twitter and Wikipedia to our analysis. This channel peaks around the same time as Twitter and Wikipedia and trends tend to end around half a day after Twitter and more than one day after Wikipedia.

3.3 Topic Category Analysis across Channels

Considering trend aware multimedia applications it is critical to understand why people turn to specific channels. To shed light on what kind of topics are the most popular in the individual channels we manually annotate the 200 top trends with categories. The categories are chosen by examining the main themes of trends we found in our dataset. Please refer to Table 4 for more information about the individual categories, their descriptions, and examples. Note that a trending topic might be assigned to multiple categories, i.e. the death of a popular music artist would be assigned to celebrity, entertainment, death, and artist to make an extreme example (similar to [1]). The engagement with respect to the different categories in a media channel is measured as follows: For each trend within the channel we assign its score to all of its categories. Finally, we normalize the scores for each channel, e.g. to account for the dominance of Google for the scores overall. The resulting distribution over categories is displayed in Figure 3.

We can observe that the channels tend to specialize in certain topic categories. For example, the most popular category in Google is sports. A large share of the scores further is assigned to celebrity and entertainment categories. Google also has the highest relative share (15%) for politics. Twitter also features many trends in the celebrity and entertainment categories. Interestingly, it has the highest relative shares of trends related to products, companies and technology. One reason might be that many Twitter users are technology affine early adopters that like to share their thoughts on new products. Another interesting finding is that over 20% of the scores on Twitter are assigned to the holidays category. We hypothesize that holiday related trends are big on Twitter as many people tag their posts and pictures with the same hashtags such as #christmas or #thanksgiving. Wikipedia clearly shows a specialization for categories that involve people and incidents such as disasters or the death of celebrities. Contrary to the intuition that Wikipedia is a slowly evolving channel which people use to read up on complicated topics, especially when also consid-

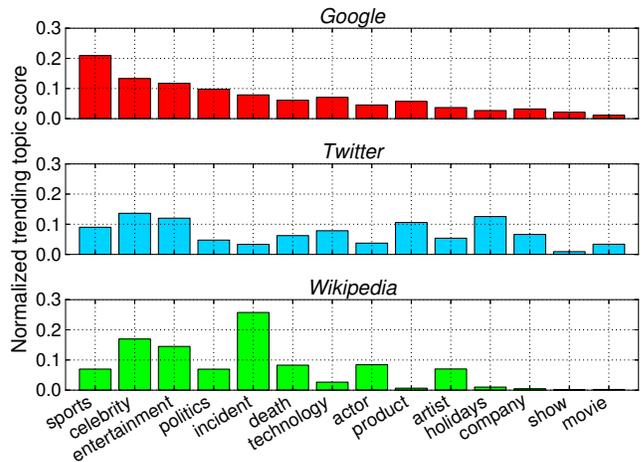


Figure 3: Normalized distribution of trending topic scores over trend categories in the individual channels. Note that channels tend to specialize in certain categories, e.g. Wikipedia for incidents.

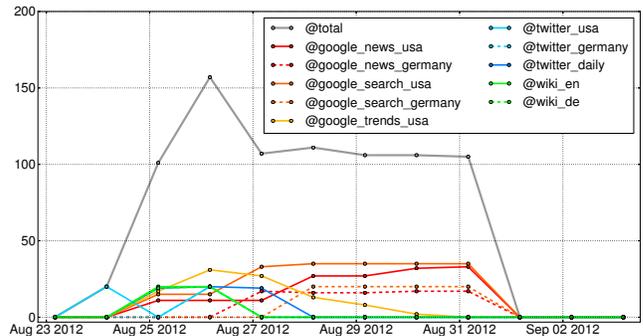


Figure 4: Neil Armstrong received much attention when he died on August 25, 2012. Note that the peak on Wikipedia (green) is at the beginning of the trending topic period whereas news coverage continued for a couple more days.

ering the temporal properties of the Wikipedia channel from the analysis above, we believe that many users use Wikipedia during these trends and events to learn about or remind themselves about related topics. For example, when Neil Armstrong died on August 25, 2012 many people viewed the corresponding Wikipedia article e.g. to learn about who he was (American astronaut), what he was famous for (first person on the moon), and how old he was when he died (82) (see Figure 4). Note that this trending topic also mistakenly subsumes the lifetime ban from sports competitions of Lance Armstrong on August 24, 2012.

4. FORECASTING OF TRENDING TOPICS

As motivated in the introduction, anticipating changing information needs and shifts in user attention towards emerging trends is critical for trend aware multimedia systems. However, forecasting trending topics is a very challenging problem since the corresponding time series usually exhibit highly irregular behavior (structural breaks) when the topic becomes “trending”. In the following, we will be using the multi-channel pipeline described in Sec. 3.1 as a robust trig-

Category	Description	#Seq	Examples
sports	sports events, clubs, athletes	52	olympics 2012, champions league, bayern muenchen
celebrity	person with prominent profile	49	steve jobs, kim kardashian, michael jackson, neil armstrong
entertainment	entertainers, movies, TV shows	39	grammys, emmys, heidi klum
politics	politicans, parties, political events, movements	32	paul ryan, occupy, christian wulff
incident	an individual occurrence or event	27	costa concordia, hurricane isaac, virginia tech shooting
death	death of a celebrity	22	whitney houston, joe paterno died, neil armstrong
technology	product or event related to technology	20	iphone 5, ces, nasa curiosity
actor	actor in TV show or movie	18	lindsay lohan, michael clarke duncan, bill cosby
product	product or product release	15	ipad, windows 8, diablo 3
artist	music artist	15	justin bieber, miley cyrus, beyonce baby
holidays	day(s) of special significance	11	halloween, thanksgiving, valentines day
company	commercial business	10	apple, chick fil a, megaupload
show	TV show	7	x factor, wetten dass, the voice
movie	a motion picture	6	dark knight rises, hunger games, the avengers

Table 4: List of categories of trending topics. Note that a trending topic might be assigned to multiple categories.

ger for trending topics but will perform time series forecasting utilizing historic time series data (see Sec. 5.1 for more details). In this section, we propose a novel forecasting approach that combines time series from multiple semantically similar topics.

A conceptual overview of our system is presented in Figure 5. On the very left you can see the trending topic for the Summer Olympics 2012 along with its Wikipedia page view statistics in 2012. The task to be solved is to forecast the number of page views for a period of 14 days (yellow area) from the day indicated by the red line. This forecast start is triggered by the emergence of a corresponding trending topic in our observed channels (see Sec. 3). Note that the time series exhibits complex behavior such as the smaller peak at the end of the forecasting period which most likely corresponds to the closing ceremony event on this day. Based on the finding that semantically similar events can exhibit very similar behavior, the first step is to automatically discover related topics such as previous Summer and Winter Olympics or FIFA/UEFA soccer championships (as illustrated by the second box). The second step in Figure

5 shows patterns of user engagement for Summer Olympics 2008, the Winter Olympics 2010, and the UEFA Euro championship 2012 that were found to match the historic behavior of the 2012 Olympics best. These patterns show certain commonalities such as a second peak for closing ceremonies or final matches. The identified sequences from the previous step are then combined to a forecast shown at the very right. In the following, these individual steps are explained in more detail.

Please note that we are neither able to nor do we attempt to predict incidents such as natural disasters or sudden deaths of celebrities in advance. However, even for unpredictable events like these, the patterns of user attention once this event has happened can be forecasted e.g. by taking previous instances of natural disasters or celebrity deaths into account.

4.1 Discovering Semantically Similar Topics

The basis for our forecasting approach is the hypothesis that semantically similar topics exhibit similar behavior. One example supporting this hypothesis is given in Figure 5 which clearly shows similar patterns of user engagement during the different sport events. For a given trending topic, we discover semantically related topics using DBPedia [3], a dataset containing structured information about several million named entities extracted from the Wikipedia project. Given category information (via `dct:terms:subject`) and type (via `rdf:type`) from DBPedia, we build a set of (up to 200) semantically similar topics that share the most categories or types with a given trending topic. For example, the Wikipedia URI for “Olympics 2012” is assigned to the categories `Sports_festivals_in_London`, `Scheduled_sports_events`, `2012_Summer_Olympics`, and `2012_in_London`. Its types include `Event`, `SportsEvent`, `Olympics`, and `OlympicGames`, e.g. the 2012 Summer Paralympics share four categories with the Olympics 2012.

Formally, we use a topic set \mathcal{T}_{sim} which includes all discovered similar topics. For later comparisons, we also define a topic set \mathcal{T}_{self} that only includes the trending topic itself and \mathcal{T}_{gen} which contains a wide variety of general topics (in our case the top 200 trending topics). In the following, we will use \mathcal{T} as a placeholder for one of these topic sets. The elements will be referred to as similar topics or $sim_j \in \mathcal{T}$.

4.2 Nearest Neighbor Sequence Matching

Obviously not all time series corresponding to the discovered similar topics look the same. Therefore, we search within these time series for sequences that match historical behavior of our trending topic. For example, historical topics such as the 1896 Summer Olympics have gained very limited attention over the last years and are therefore unlikely to be representative for the large amount of engagement towards the 2012 Summer Olympics. To pick the right instances to inform our forecast, we compare a short history window of the trending topic to be forecasted, i.e. the viewing statistics for the last two months, to all partial sequences of the same length of similar topics in our topic set \mathcal{T} .

To capture this step in formal terms, let $S_{topic}[t] \in \mathbb{R}$ be the time series for the given topic at time $t \in \mathbb{N}$. Further define $S_{topic}^{t_0}[t] := S_{topic}[t_0 + t]$ as the shifted version of the time series (used for aligning multiple series below). In the following, we assume that we want to forecast S_{topic} with a horizon of h days starting at time $t_0 \in \mathbb{N}$.

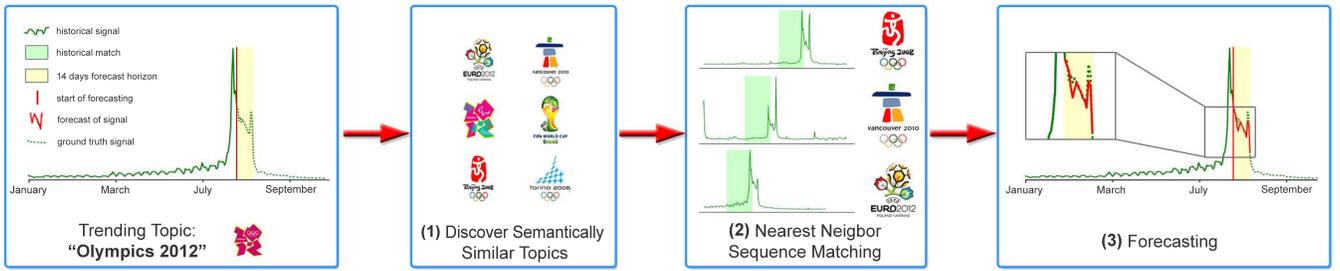


Figure 5: System overview of our proposed forecasting approach. For a given trending topic we first discover semantically similar topics which we search for patterns of similar behavior which are then used to produce a forecast.

Given a topic set \mathcal{T} from the previous step, we define a sequence candidate set $C(t_0)$ that includes all possible shifted time series $S_{\text{sim}_j}^t$: $C(t_0) = \{S_{\text{sim}_j}^t \mid \forall \text{sim}_j \in \mathcal{T}, \forall t : t \leq t_0 - h\}$. The condition $t \leq t_0 - h$ ensures that we never use information more recent than $t_0 - h$, to allow for a forecast of h days, i.e. we do not utilize future information.

Given this candidate set, we now search for the k members $S_{\text{sim}_i}^{t_i}$ ($i = 1, \dots, k$) that are the best matches for our time series of interest ($S_{\text{topic}}^{t_0}$). Note that these nearest neighbors are already correctly aligned through shifting the time series S_{sim_i} by a corresponding t_i . Note that the same similar topic sim_i can occur multiple times (e.g. for repetitive signals).

Formally, the nearest neighbor set becomes $N^k(S_{\text{topic}}^{t_0}) = \{S_{\text{sim}_1}^{t_1}, \dots, S_{\text{sim}_k}^{t_k}\}$ where the $S_{\text{sim}_i}^{t_i}$ are the k distinct elements that are smallest wrt. $d(S_{\text{topic}}^{t_0}, S_{\text{sim}_i}^{t_i})$ for all $S_{\text{sim}_i}^{t_i} \in C(t_0)$. Here, $d(\cdot, \cdot)$ is a distance metric between both time series which, in our case, only depends on a short history window of the time series. An interesting question is whether the metric should be scale invariant and in which form and to what degree. In Sec. 5, we compare three different distance metrics: (1) (squared) **euclidean** distance: $(d(x, y) = \sum_{i=1}^n (x_i - y_i)^2)$, (2) euclidean distance on normalized sequences $x'_i = (x_i - \mu)/\sigma$ (where μ, σ are mean and standard deviation estimated from the respective time series) referred to as **musigma**, and (3) a fully scale invariant metric (called **y_invariant**) that was proposed in [23] ($\min_{\gamma} d(x, \gamma \cdot y)$).

4.3 Forecasting

Even the best matching sequences identified in the previous step might not be a perfect fit for the time series to be forecasted. Therefore, we rescale the matching sequences such that they all match on the last observed value of $S_{\text{topic}}^{t_0}$. This ensures that the forecast will be a continuous extension of past behavior. Now, the forecast \mathcal{F} becomes the median over scaled versions of the sequences from the previous step ($N^k(S_{\text{topic}}^{t_0})$). We represent the days that we want to forecast by $\tau \in [0, \dots, h - 1]$ where again h is the forecasting horizon (usually $h = 14$). Our forecast can then be formally described as

$$\mathcal{F}(S_{\text{topic}}^{t_0})[\tau] = \text{median}_{S_{\text{topic}'}^{t_0'} \in N^k(S_{\text{topic}}^{t_0})} (\alpha(S_{\text{topic}}^{t_0}, S_{\text{topic}'}^{t_0'}) \cdot S_{\text{topic}'}^{t_0'}[\tau])$$

where $\alpha(S_{\text{topic}}^{t_0}, S_{\text{topic}'}^{t_0'}) = S_{\text{topic}}^{t_0}[-1](S_{\text{topic}'}^{t_0'}[-1])^{-1}$ adjusts the scale of nearest neighbor time series based on the last observed score. In practice, we limit α to a limited interval (e.g. $[0.33, 3.0]$) for robustness. We further evaluate our

approach using the average instead of the median for forecasts. However, as presented in Sec. 5 the median proves to be more robust.

5. EVALUATION

In this section, first the Wikipedia page views dataset is described before presenting quantitative results for the forecasting approach proposed in the previous section.

5.1 Dataset Description

We evaluate our proposed forecasting approach on a large-scale dataset of page views on Wikipedia, an online collaborative encyclopedia that has become a mainstream information resource worldwide and is frequently used in academia [16, 22, 14, 20]. Reasons for this particular choice of social media channel were (a) the public availability of historical views data necessary to build forecasting models (hourly view statistics for the last five years), (b) the size of the dataset allowing a comprehensive analysis of our proposed method across a wide range of topics (over 5 million articles), and (c) previous results show that user behavior on Wikipedia (bursts in popularity of Wikipedia pages) is well correlated with external news events [16] and therefore support our proposed method’s ability to generalize to different social media channels. Although we only provide results on Wikipedia our approach could be applied to any online and social media channel for which historic data is available.

The raw Wikipedia viewing statistics are published³ by the Wikimedia foundation. We obtained hourly view statistics starting on January 1, 2008 (2.8 TB compressed in total). We aggregate these logs to daily viewing statistics where we drop URIs that have been viewed less than 25 times on that day. We believe that this does not introduce any bias since trending topics tend to accumulate volumes several orders of magnitude higher. For each day this results in approx. 2.5 million URIs attracting 870 million daily views. In total, the English and German Wikipedia features more than five million articles that can be used for forecasting. During data preparation we noticed that Wikipedia did not record view counts for a few short periods of time during the last five years for which we interpolate linearly.

Note that we use this dataset for forecasting of historical time series data (actual Wikipedia viewing statistics) while the multi-channel pipeline described in Sec. 3.1 serves as a robust trigger for trending topics.

³<http://dumps.wikimedia.org/other/pagecounts-raw/>

Trending Topic	Nearest Neighbor Topics
2012 Summer Olympics	2008 Summer Olympics, UEFA Euro 2012, 2010 Winter Olympics :: 2016 Summer Olympics, 2014 FIFA World Cup, 2006 Winter Olympics
Whitney Houston	Ciara, Shakira, Celine Dion, Brittany Murphy, Ozzy Osbourne :: Alicia Keys, Paul McCartney, Janet Jackson
Steve Jobs	Mark Zuckerberg, Rupert Murdoch, Steve Jobs :: Steve Wozniak, Bill Gates, Oprah Winfrey
Super Bowl XLVI	Super Bowl, Super Bowl XLV, Super Bowl XLIV :: Super Bowl XLIII, 2012 Pro Bowl, UFC 119
84th Academy Awards	83rd Academy Awards, 82nd Academy Awards :: List of Academy Awards ceremonies, 81st Academy Awards
Battlefield 3	Mortal Kombat, FIFA 10, Call of Duty: Modern Warfare 2, Portal, Duke Nukem Forever :: Call of Duty: Modern Warfare 3, Call of Duty 4: Modern Warfare, Pro Evolution Soccer 2011
Joe Paterno	Terry Bradshaw, Joe Paterno, Jack Ruby, Paul Newman, Jerry Sandusky :: Lane Kiffin, Donna Summer, Joe DiMaggio
54th Grammy Awards	53rd Grammy Awards, 52nd Grammy Awards, 54th Grammy Awards :: 51st Grammy Awards, 2012 Billboard Music Awards, 2012 MTV Europe Music Awards

Table 5: Selected trending topics along with their nearest neighbor topics using category and type information on DB-Pedia (step 1). The ones chosen by nearest neighbor sequence matching (step 2) are in bold print. In some cases the topic itself can be used for forecasting, e.g. if the time series contains repetitive patterns.

5.2 Experiments

We structure our experiments along the three main building blocks of our proposed approach to compare design choices for the individual methods independently (recall Fig. 5).

1. Discovering Semantically Similar Topics.

We evaluate the influence of discovering semantically similar topics in two ways. First, we present qualitative results by showing retrieved similar topics for a few trending topics. Second, we evaluate this part indirectly by comparing forecast performance (i.e. through the third step) of using semantically similar topics from DBPedia to using a general set of topics (the top 200 trending topics themselves).

Subsets of semantically similar topics for the top trends are shown in Table 5 (as generated by the approach described in Sec. 4.1). Note that we successfully identify similar events or previous instances for events like the Olympics, the Super Bowl, UFC events, or the Grammy Awards. We are also able to discover similar people like similar music artists, entrepreneurs, athletes, and even people that died from the same cause (such as lung cancer for Joe Paterno, Jack Ruby, and Paul Newman). On average, we retrieve 95 semantically similar topics per trending topic.

2. Nearest Neighbor Sequence Matching.

The main design choice when matching sequences from similar topics to a short history window of our time series is

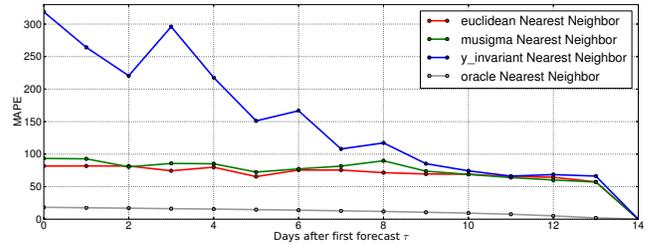


Figure 6: A comparison of Nearest Neighbor Distance Metrics. Euclidean distance performs at least as good as its more complex normalized counterparts.

the choice of the distance metric. As introduced in Sec. 4.2 we compare `euclidean`, `musigma`, and `y_invariant`.

Our experimental setup for this part as well as the forecasting looks as follows. We use the trending topics acquired in Sec. 3 as a trigger for forecasting, i.e. for each of the 516 sequences of the top 200 trending topics (see Table 2), we make a prediction for a horizon of 14 days, starting on the day they first emerge. We chose this window of 14 days since this represents the maximum lifetime for most trending topics (see Figure 2 (a)). Our method is evaluated by beginning with a forecast of 14 days, then a 13 day forecast after one day and so on (X axis in error plot).

To compare the distance metrics, we measure the quality of the (first) nearest neighbor returned by this metric by its similarity to the actual viewing statistics over the next 14 days (similar to the forecast setting but directly using the nearest neighbor as the forecast). The quality of this match is captured by an error metric. In this work, we use root mean squared error defined as $RMSE = \sqrt{\frac{1}{n} \sum_{t=1}^n (A_t - F_t)^2}$ where A_t is the actual value and F_t is the forecast value. As a relative error metric we choose mean absolute percentage error (MAPE) defined as $MAPE = \frac{100\%}{n} \sum_{t=1}^n \left| \frac{A_t - F_t}{A_t} \right|$. Note that unlike [15] we measure error on the actual view counts instead of normalizing by the total views for each day as we believe this glosses over large relative errors (i.e. even larger trends might only account for 10^{-5} of the daily views yielding very small error rates for virtually any forecast).

The results are depicted in Figure 6 and show the average MAPE error (lower is better) between the nearest neighbor and the actual viewing statistics (ground truth). We also show oracle performance by picking the best match for the ground truth from all sequence candidates. First, note that oracle performance is well above 0% error as there might be no perfect match among the sequence candidates. Further, we observe that `y_invariant` can retrieve poor quality matches whereas the simple `euclidean` distance performs equally or better than the other two metrics. Apparently, the invariance of `musigma` and `y_invariant` does not help in this task. Therefore, we use the `euclidean` metric for all following experiments. Further note that the best matches have between 82 and 319% error illustrating the high complexity of the task.

3. Forecast.

As introduced above we measure forecasting performance by forecasting the next 14 days for each of the 516 sequences of the top 200 trending topics at the point in time when they

		RMSEs in 1000		
Method		$\tau = 0$	$\tau = 3$	$\tau = 7$
Baselines	naive	63.2	33.1	17.4
	linear trend	86.9	48.5	23.1
	average trend	49.3	25.9	19.9
	median trend	48.1	24.9	18.1
<i>Self</i>	average	46.0	23.7	18.0
	average_scaled	44.6	21.9	15.5
	median	46.1	23.8	17.7
	median_scaled	44.9	22.3	15.5
<i>Gen</i>	average	45.7	22.9	16.1
	average_scaled	45.7	22.5	14.1
	median	41.4	21.2	15.4
	median_scaled	40.1	19.5	12.8
<i>Sim</i>	average	41.4	18.8	14.0
	average_scaled	39.6	17.1	11.6
	median	42.1	19.9	14.0
	median_scaled	41.0	17.9	11.5

Table 6: RMSE forecasting error for our baselines as well as methods using only the trending topics itself (*Self*), a general set of topics (*Gen*), or similar topics (*Sim*).

first emerge. Our approach is compared to several baselines that use a short history window of the time series itself (similar to [15]): a *naive* forecast (tomorrow’s behavior is the same as today’s) and a *linear trend* based on the last 14 days. We further compare to the *average trend* and *median trend* in our trending topics dataset as a baseline that includes multiple time series. Note that this average and median trend are computed from μ/σ -normalized time series since the average/median of actual view counts are actually very far from most trending topics. To still be able to compute the error for actual view count prediction we then de-normalize the *average trend* and *median trend* baselines with the parameters of the time series to be predicted. Note that while we performed the following experiments for different numbers of neighbors (for the nearest-neighbor-based methods) we only report the results for $k = 3$ which performed best by a small margin.

The RMSE forecasting errors are summarized in Table 6 which reports the results for a number of instances of our proposed forecasting approach. *Self*, *Gen*, and *Sim* refer to the different topic sets \mathcal{T}_{self} , \mathcal{T}_{gen} , and \mathcal{T}_{sim} from which the nearest neighbor sequences are chosen (as described in the last section). On average, the RMSE of our best method is about 9-48k views closer to the actual viewing statistics than the different baseline methods. Further, we observe that taking the median tends to perform about as good or better than taking the average, using scaled nearest neighbor is better than unscaled neighbors, and that using semantically similar topics (*Sim*) is better than using a general set of topics (*Gen*) which in turn is better than restricting oneself to a single time series (*Self*).

However, RMSE error has the disadvantage that it is dominated by the most popular trending topics with the largest view counts. Therefore, we choose the MAPE measure for the remaining analysis as this relative error metric is comparable across trending topics. Additionally, because MAPE errors can become disproportionately large (e.g. forecasting 1000 views when it is actually only 100 results in a 900% relative error), we drop obvious outliers (5%) and report the average error for the remaining sequences. We only show the results for the baselines and the best performing methods

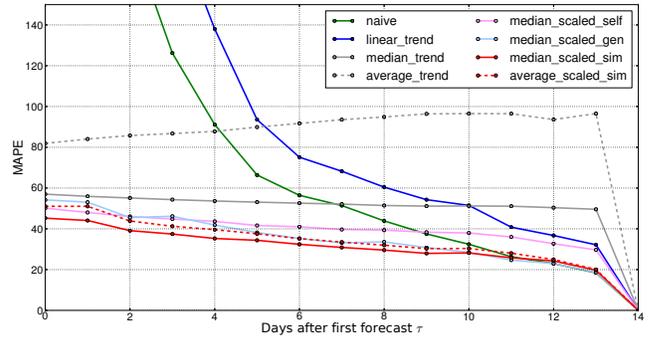


Figure 7: MAPE forecast error moving through a 14 day window, e.g. the error at 4 depicts the forecasting error of the following 10 days.

from Table 6 to be able to visually distinguish them in the error plot. Again, the methods are evaluated by beginning with a forecast of 14 days, then a 13 day forecast after one day etc. as illustrated by the X axis in Figure 7. We can observe that the proposed approach including median and scaling clearly outperforms all baselines as well as other instances of our framework. Our proposed method achieves a mean average percentage error (MAPE) of 45-19%, a relative improvement over our baselines of 20-90%.

We further performed a per-category analysis of the forecasting results (using the categories from Table 4). Summarizing, the forecasting performance across categories is similar. Product, company, artist and technology categories perform slightly better, and death, actor, sports, holidays and politics perform slightly worse. For all categories, the forecast quality improves with time (τ).

Example forecasts for two trending topics, “Battlefield 3” (computer game) and “The Hunger Games” (novel and movie), are given in Figure 8. Each column depicts multiple forecasts at different points in time (as indicated by the vertical red line). Our proposed method *median_scaled_sim* (*Sim*) is compared to its variants only using the trending topic itself (*Self*) or using a set of general topics (*Gen*). Note that in many cases, only using information from a single time series is not enough. Further, using semantically similar topics (*Sim*) leads to more accurate forecasts that e.g. are able to capture multiple peaks (such as in the Hunger Games example).

6. CONCLUSION

We present a comprehensive study of trending topics in three major social and online media streams with an observation period of one year. Our findings include that trends on Twitter and Wikipedia are more ephemeral than on Google, both rising and declining rapidly for newly emerging topics. Further, we find that the observed media channels tend to specialize in specific topic categories. The second key contribution is a fully automatic forecasting technique of trending topics exploiting semantic similarity between topics. Lastly, we evaluate our approach on a large-scale dataset of Wikipedia viewing statistics and show empirically that forecasts by the proposed approach are about 9-48k views closer to the actual viewing statistics than baseline methods and achieve a mean average percentage error

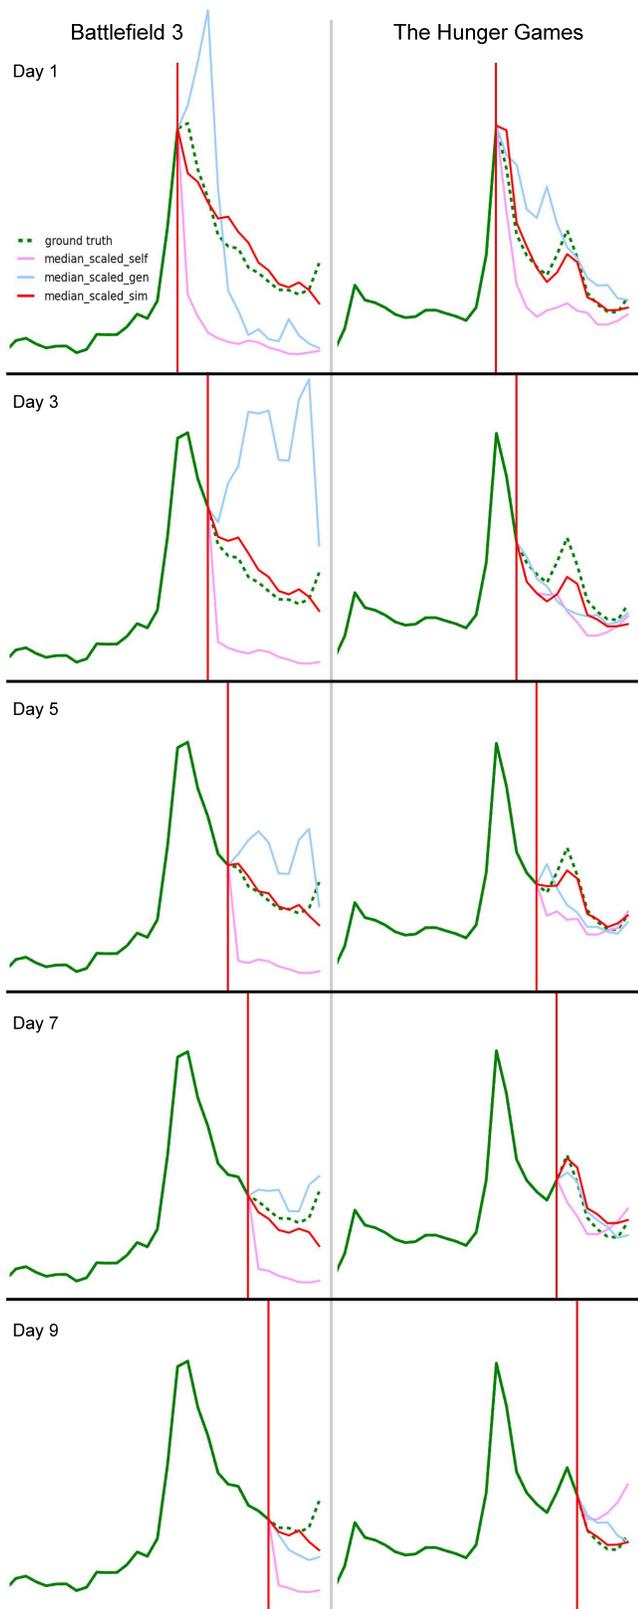


Figure 8: Visualization of trending topic forecasts for “Battlefield 3” and “The Hunger Games”. Each column depicts multiple forecasts at different points in time (as indicated by the vertical red line).

of 45-19% for time periods of up to 14 days, a relative improvement over baselines of 20-90%.

As future work, the proposed model could be extended to explicitly detect and exploit seasonality as well as incorporate global changes in viewing statistics. In a more general ranking framework, the nearest neighbor sequence matching part could further be extended by e.g. the number of common semantic types and categories.

7. REFERENCES

- [1] E. Adar et al. Why we search: visualizing and predicting user behavior. In *Proc. of WWW*, 2007.
- [2] J. Allan. Introduction to topic detection and tracking. pp. 1–16. Springer, 2002.
- [3] S. Auer et al. Dbpedia: A nucleus for a web of open data. *The Semantic Web*, pp 722–735, 2007.
- [4] R. Bandari, S. Asur, and B. Huberman. The pulse of news in social media: Forecasting popularity. *Proc. of ICWSM*, 2012.
- [5] D. Borth, A. Ulges, and T. M. Breuel. Dynamic vocabularies for web-based concept detection by trend discovery. In *Proc. of ACM MM*, pp. 977–980. ACM, 2012.
- [6] H. Choi and H. Varian. Predicting the present with google trends. *Economic Record*, 88(s1):2–9, 2012.
- [7] M. P. Clements and D. F. Hendry. Forecasting annual uk inflation using an econometric model over 1875–1991. *Frontiers of Economics and Globalization*, 3:3–39, 2009. Forecasting in the presence of structural breaks and model uncertainty.
- [8] M. De Choudhury, S. Counts, and M. Czerwinski. Find me the right content! Diversity-based sampling of social media spaces for topic-centric search. In *Proc. of ICWSM*, 2011.
- [9] S. Diplaris et al. Socialsensor: sensing user generated input for improved media discovery and experience. In *Proc. WWW companion*, pp. 243–246. ACM, 2012.
- [10] S. Goel et al. Predicting consumer behavior with web search. *Proc. of NAS*, 107(41):17486–17490, 2010.
- [11] R. J. Hyndman and Y. Khandakar. Automatic time series forecasting: The forecast package for r. *Journal of Statistical Software*, 27(3), 2008.
- [12] X. Jin et al. The wisdom of social multimedia: using flickr for prediction and forecast. In *Proc. of ACM MM*, pp. 1235–1244. ACM, 2010.
- [13] M. Naaman, H. Becker, and L. Gravano. Hip and trendy: Characterizing emerging trends on twitter. *Journal of ASIST*, 62(5):902–918, 2011.
- [14] M. Osborne et al. Bieber no more: First story detection using twitter and wikipedia. In *SIGIR 2012 WS on Time-aware Information Access*, 2012.
- [15] K. Radinsky et al. Modeling and predicting behavioral dynamics on the web. In *Proc. of WWW*, 2012.
- [16] J. Ratkiewicz, A. Flammini, and F. Menczer. Traffic in social media I: paths through information networks. In *Proc. of SocialCom*, pp. 452–458. IEEE, 2010.
- [17] S. D. Roy, T. Mei, W. Zeng, and S. Li. Socialtransfer: cross-domain transfer learning from social streams for media applications. In *Proc. of ACM MM*, 2012.
- [18] S. Samangooei et al. Trendminer: An architecture for real time analysis of social media text. In *ICWSM*, 2012.
- [19] G. Szabo and B. A. Huberman. Predicting the popularity of online content. *Comm. of the ACM*, 53(8):80–88, 2010.
- [20] M. t. Thij, Y. Volkovich, D. Laniado, and A. Kaltenbrunner. Modeling and predicting page-view dynamics on Wikipedia. *arXiv*, 2012.
- [21] Z. Wang et al. Propagation-based social-aware replication for social video contents. In *Proc. ACM MM*, 2012.
- [22] S. Whiting and O. Alonso. Hashtags as milestones in time. In *SIGIR 2012 WS on Time-aware Info. Access*, 2012.
- [23] J. Yang and J. Leskovec. Patterns of temporal variation in online media. In *Proc. of WSDM*, 2011.