

Towards Supporting Search over Trending Events with Social Media

Sanjay R. Kairam¹, Meredith Ringel Morris², Jaime Teevan², Dan Liebling², and Susan Dumais²

¹Stanford University
353 Serra Mall, Stanford, CA
skairam@cs.stanford.edu

²Microsoft Research
One Microsoft Way, Redmond, WA
{merrie, teevan, danl, sdumais}@microsoft.com

Abstract

Many search engines identify bursts of activity around particular topics and reflect these back to users as Popular Now or Hot Searches. Activity around these topics typically evolves quickly in real-time during the course of a *trending event*. Users' informational needs when searching for such topics will vary depending on the stage at which they engage with an event. Through a survey and log study, we observe that interaction with content about trending events varies significantly with prior awareness of the event. Building on this observation, we conduct a larger-scale analysis of query logs and social media data associated with hundreds of trending events. We find that search and social media activity tend to follow similar temporal patterns, but that social media activity leads by a few hours. While user interest in trending event content predictably diverges during peak activity periods, the overlap between content searched and shared increases. We discuss how these findings relate to the design of interfaces to better support sensemaking around trending events by integrating real-time social media content with traditional search results.

Introduction

Trending events are events that serve as novel or evolving sources of widespread online activity. Such events range in nature from anticipated events (e.g., Summer Olympics) to breaking news (e.g., Aurora shooting), and topics can vary widely from politics to sporting events to celebrity gossip. In the last few years, popular Web search engines have begun reflecting these patterns of activity back to users in the form of *Trending Queries* (e.g., Bing *Popular Now*, Google *Hot Searches*, Yahoo! *Trending Now*). In this paper, we aim to improve support for searchers issuing these types of queries by studying how their information needs evolve during the course of a trending event.

Research on crisis informatics has demonstrated that social media users can generate and synthesize valuable information in a real-time, distributed manner (Starbird et al. 2010). Users already appear to utilize Twitter search for

finding and monitoring information about time-sensitive topics (Teevan, Ramage, and Morris 2011). However, research has shown that the topics discussed on Twitter can change quickly (Kwak et al. 2010; Lin and Mishne, 2012), so it is not clear for how long information about these topics will persist. We pose the questions: *For what types of trending events will real-time information be useful, and for how long will it continue to align with the information needs of users searching about these events?*

This paper explores these questions, engaging in what we believe to be the first systematic exploration of trending events through the lens of search activity. We identify differences in user information needs, particularly with respect to the consumption of real-time content, and the applicability of social media for satisfying these needs. We explore these questions by examining hundreds of events that trended during the summer of 2012, using (1) qualitative survey data, (2) query logs from Bing, and (3) Twitter updates from the complete Twitter Firehose. Our findings reveal that:

- **Searchers who click *Trending Queries* links engage less and with different result content than users who search manually for the same topics.** Survey results indicate that this may be due to a preference for real-time information that is perhaps not currently being satisfied.
- Search query and social media activity follow similar temporal patterns, but **social media activity tends to lead by 4.3 hours on average**, providing enough time for a search engine to index and process relevant content.
- **User interest diverges during the peak of activity for a trending event**, as reflected by a spike in the entropy of content searched and shared; however, a corresponding increase in content overlap highlights opportunities for supporting search with social media content.

We discuss implications of these findings for the design of systems to leverage social media content and support sensemaking around novel, widespread phenomena such as trending events.

Related Work

We begin by describing three relevant lines of research: 1) trending events in search, 2) trending events in social media, and 3) social information seeking.

Trending Events in Search. We study search activity surrounding trending events by analyzing search logs. Search logs allow us to observe patterns of behavior across millions of users, and have provided insight into the types (Broder 2002) and topics (Spink et al. 2001) of events for which users search. Following prior recommendations (Grimes, Tang, and Russell 2007), we complement our log analysis with qualitative data from users.

Our analysis of temporal patterns in search behavior draws on prior study of long-term temporal query dynamics. We adopt methods from Kulkarni et al. (2011) for categorizing events according to these patterns, and we extend methods from Adar et al. (2007) for comparing patterns across information streams. Our work differs both in scale (our focus is on hours and days rather than weeks and months) and scope (we focus on a specific class of events). Prior work has also aimed at characterizing query dynamics by examining query result content (e.g., Jones and Diaz 2007; Kotov et al. 2010). This work informs ours, but does not directly address our goals of characterizing correspondences between content searched and shared in real-time over the course of a trending event.

Trends in Social Media. As the largest source of public social media activity, Twitter is a popular target for the study of trends. Kwak et al. (2010) compared 4,000 Twitter trends to the top keywords from Google *Trends* revealed little overlap in the topics surfaced by. Manual inspection of the trends found that 85% of the topics represented “headline” or “persistent” news. This observation is comparable to prior efforts (Zubiaga et al. 2011) in which manual classification identified 73% of Twitter *Trends* to be related to “news” or “current events.”

Naaman, Becker, and Gravano (2011) present a more detailed taxonomy, separating trends into *exogenous* (breaking news, broadcast events, holidays, and local events) and *endogenous* (memes, retweets, and fan activity) events and identifying temporal, content, and other features characteristic of various trend types. We extend this line of research to examine events trending in queries on a major search engine, conducting what we believe to be the first large-scale study of query activity with respect to trending events.

Automatic identification of trends in web and text data is an interesting and challenging problem (Gabrilovich, Dumais, and Horvitz 2004; Kleinberg 2006; Marcus et al. 2011; Vlachos et al. 2004). In our analysis, we rely on the trends identified by the online services that we studied in

order to focus specifically on user interactions with trends that have been surfaced and reflected back to users.

Social Information Seeking. Socially-generated content is often used to address users’ information needs. Efron (2011) describes two types of search in social systems such as microblogs: (1) asking questions to one’s network, and (2) searching over social repositories. We focus on the latter, drawing on observations about the complementary benefits of searching and asking to support sensemaking (e.g. Morris, Teevan, and Panovich 2010). Posing questions to one’s social network, for instance, has been shown to produce less task-relevant information while stimulating engagement and sensemaking (Evans, Kairam, and Pirolli 2010).

Prior research comparing queries issued to search engines with those issued on Twitter (Teevan, Ramage, and Morris 2011) and blogs (Mishne and de Rijke 2006; Sun, Hu, and Lim 2008) has found that queries over social resources tend to focus more on people, named entities, and temporally-relevant content. Topics searched on Twitter change quickly; Lin and Mishne (2012) recently showed that churn rates for top Twitter queries are up to four times higher than those for search, with these rates increasing during major events, such as the trending events we study. Our analysis differs in that we compare web queries directly against social media content, providing insight into how such content can better support patterns existing already in major search engines.

Collecting Trending Events

To study people’s experiences with trending events in search and social media, we collected trending events from two sources, Twitter *Trends* and Bing’s *Popular Now* queries (referred to from here as *Trending Queries*), over a six-week period starting July 19, 2012.

For each trending event, we also collected a dataset of matching queries and tweets from users within the United States. We stemmed and removed stop words from the *Trends* and *Trending Queries* shown to users; we then matched those tokens against all queries issued via the search engine homepage and all public tweets for a period starting one week before the trend appeared and continuing one week afterwards. If all tokens appeared within a query or tweet, it was considered a match; word-order, case, and non-alphanumeric characters were not considered. For example, “Toyota Recall” matched the query “Toyota Camry recall,” but not the query “toyota recal [sic].” We chose this technique because it captured more content than strict keyword matching without introducing some of the complexities associated with more sophisticated approaches, such as topic modeling (cf. Ramage, Dumais, and Liebling 2009; Teevan, Ramage, and Morris 2011).

Entry Point	% Click on Answer	% Click on Result	Click Entropy
Link	17.98%	4.64%	2.93
Typing	31.73%	29.28%	4.13

Table 1. Post-search behavior for users who click a *Trending Queries* link and those who type queries manually. Columns show percentage of users for whom the first click is on an *Instant Answer* or a standard search result, as well as the click entropy. All differences are significant ($p < 0.001$).

Preliminary analysis revealed that many single-word *Trends* reflected topics internal to the Twitter community (e.g., memes like *#MostShareWorthyMovies*); given our focus on exogenous events, we filtered all single-word trends. To mitigate the number of overlapping trends, we also removed any trend that was a superset of another (e.g., “Hurricane Isaac Forecast” was removed if “Hurricane Isaac” was a trend). This resulted in 763 trending events (370 *Twitter Trends* and 393 *Trending Queries*). We further filtered out 415 trends without sufficient activity in both sources. We used a simple trend-detection algorithm similar to that used by Marcus et al. (2011) to remove 17 additional events with no detectable “spike” of activity. These filtering steps left us with 331 trending events (113 *Twitter Trends* and 218 *Trending Queries*), each with a two-week corpus of associated queries and tweets.

Trending Events and User Search Needs

Using these trending events, we engaged in two studies aimed at relating users’ prior awareness of a trending event to their search behavior. The first identifies quantitative differences in post-search behavior by comparing people who search for trending events by typing queries directly into the search engine and those who click on *Trending Queries* links. The second utilizes qualitative survey data to extend and explain these findings, particularly with respect to preferences for real-time information.

Engagement with Search Result Content

To explore how search behavior varies with prior awareness, we studied users’ interactions with web search results for trending event queries. As a proxy for awareness, we looked at whether users typed queries manually into the search engine or clicked *Trending Queries* links. We assumed that users typing queries were, on average, more likely to be aware of an event than users clicking *Trending Queries* links, who may be new to an event and prompted to click by the search engine.

Method

From the search engine logs, we extracted post-query behavior for queries associated with each trending event. To control for variation, we restricted our analysis to

queries initiated from the search engine homepage, either via typing or via a *Trending Queries* link. For 233 (74.9%) of our trends, we observed search queries issued from the home page using both methods. Query volumes per trend ranged from tens to tens of thousands (median: 22,229).

As search engine interaction behavior can vary greatly by task, we compared post-query behavior on a per-trend basis (e.g., users typing queries associated with “Honey Boo Boo” were compared directly with users clicking a “Honey Boo Boo” *Trending Queries* link). The same results were returned regardless of how the query was issued, allowing for direct post-query comparisons. For trending queries, result pages often consist of both standard results and *Instant Answers* (i.e., summary content shown above the results, usually news results for trending events). Significance was calculated using a two-tailed pairwise *t*-test. All differences reported are significant ($p < 0.001$).

Results

Overall, we observe less interaction with result content when a trending query is issued via link than by manual entry. Table 1 shows differences in post-query behavior according to how the query was issued. The percent of manual queries for which users click any content (61.01%) is almost three times that for link queries (22.62%).

We observe less diversity in post-query behavior from users who click trending query links. These users are almost four times as likely to click on an instant answer than a standard search result (17.98% vs. 4.64%), while users who query manually click these options with similar frequencies (31.73% vs. 29.28%).

Click entropy captures the variability in results clicked in response to a query q . It has been used to measure query result diversity (Dou, Song, and Wen 2007; Clough et al. 2009) and user satisfaction (Weber and Jaimes 2007), and is defined as:

$$\text{Click-entropy}(q) = - \sum_{u|u} p(u | q) \times \log(p(u | q))$$

For users who do click after searching, the click entropy is higher for manual queries (4.13) than for link queries (2.93), indicating higher variability in clicked results.

We observe that users behave quite differently depending on how they initially engage with trending event queries. Together, these results suggest that users who click *Trending Query* links may be less engaged with these events, have needs currently unmet by the search engine, or may be satisfied with the limited content available in the result snippets. When they do click, the content they engage with is more homogenous and more likely to be satisfied by an *Instant Answer* than the algorithmic results. This may indicate an opportunity to better support and engage these users with additional real-time content.




Figure 1. Information sources used for searching information about trending topics, as reported by survey respondents. Non-social sources (Online News, Broadcast Media, Search Engines) were reported with higher frequency than social sources.

User Motivation and Search Strategies

To support these observations from query logs, we also conducted a survey to examine how user motivation and prior awareness influenced search strategies and needs.

Method

Using Amazon Mechanical Turk, we issued surveys daily from Monday, August 27 to Friday, August 31, 2012. In the survey, we asked participants about a current trending event, including their familiarity with the event, sources used, and information found. Participants were shown a list of 17 trending events that had appeared as *Twitter Trends* or *Trending Queries* within the previous 24 hours and asked to select one with which they had recently engaged (or choose “None” where applicable). Eight of these events were trends appearing as *Trending Queries*, and nine were *Twitter Trends* (excluding promoted trends).

Participation was restricted to residents of the U.S. and Canada, and participants were paid \$0.20 per survey completed. Although they could not complete the same day’s survey multiple times, they were able to participate across multiple days. Low-quality results were mitigated where possible by randomizing answer order for multiple choice questions and by including short free-text response questions which allowed for easy manual flagging of off-topic or irrelevant answers. 453 surveys were initiated in total; below, we discuss data from the 288 fully completed surveys in which respondents reported engaging with one of the trending events (e.g. did not choose “None”).

Participants. Excluding the six participants who declined to provide demographic information, participants were evenly split by gender (48.8% female) with a median age range of 21-29. The majority (83.8%) had completed at least some college, and roughly half (47.8%) had obtained a degree. These demographics roughly match Quantcast (<http://quantcast.com>) statistics for top search engines and social media sites, such as Bing, Google, and Twitter.

Almost all participants (97.9%) reported using search engines at least daily. The proportion of respondents who read social media content at least weekly (Facebook: 76.2%; Twitter: 35.5%) was roughly twice the proportion posting content at least weekly (Facebook: 39.0%; Twitter:




Figure 2. Reported utility of information types. N/A indicates that participants did not find this type of information.

19.8%). Most participants were not frequent consumers of explicitly “trending” content; the majority indicated that they clicked on *Twitter Trends* (78.7%) or search engine *Trending Queries* (60.9%) less than once a month.

Results

Survey responses covered 49 of the 85 trends about which we inquired. The most frequently-chosen events centered on aspects of two salient real-world events that occurred during the study period: Hurricane Isaac (*Tropical Storm Isaac, Hurricane Isaac Path*) and the Republican National Convention (*GOP Convention, Clint Eastwood*). Below, we focus on results regarding participants’ prior awareness of the trending event, sources used to learn about the event, and perceived utility of various types of information.

Prior Awareness. Most respondents (73.3%) indicated having looked for information about the chosen trend within the prior 48 hours. Participants generally chose trends of which they had recently become aware and with which they were not familiar. The majority (80.9%) indicated being aware of the chosen trend for less than a week, and less than a third (33.0%) reported being very or expertly familiar with it.

Information Sources. Participants indicated whether or not they had used each of several information sources for finding information about the chosen trends. Figure 1 shows the percentage of participants reporting using each source. The most frequently reported sources were non-social in nature (e.g., online news, broadcast media, search engines); social sources (e.g., forums, blogs, Twitter) were used much less frequently. The median number of sources participants reported consulting was two, indicating that many users currently combine information from multiple locations to learn about trending events.

Information Needs. We also asked participants to indicate the utility of each of the following types of information in learning about trending topics: *Real-Time/Breaking Updates, Public Opinion/Sentiment, Friend Commentary, Expert Commentary, and Background Information About Relevant People/Places/Organizations*. Figure 2 shows the responses. Real-time information appeared most valuable, with 86.1% reporting they found it “somewhat” or “very” useful. Expert commentary was also judged useful, with 77.7% of respondents finding it at least “somewhat” useful.

Kendall's τ , a measure of correlation between ordinal variables, was used to assess the relationship between the reported utility of each of the found information types and the measures of trend awareness listed above. We find that respondents who had searched more recently about an event rated real-time information as more helpful ($\tau = -0.213, p < 0.001$). Similarly, respondents who had become aware of the event more recently rated real-time information ($\tau = -0.193, p < 0.001$) and expert commentary ($\tau = -0.153, p < 0.005$) as more useful.

Chi-squared tests of independence were performed to examine the relationships between reported utility of information and the information sources used; to avoid data sparsity issues, we focused on the four most frequently used sources (online news, broadcast channels, search engines, and Facebook). Respondents who used Facebook ascribed significantly higher utility to commentary by friends ($\chi^2(3, N=288) = 22.87, p < 0.001$). Respondents who found information through broadcast channels valued real-time information ($\chi^2(3, N=288) = 11.38, p < 0.01$) and expert commentary ($\chi^2(3, N=288) = 12.01, p < 0.01$) more. Respondents who used online news to find information also highly rated the utility of real-time information ($\chi^2(3, N=288) = 18.44, p < 0.001$).

Discussion

We observe differences in information needs as a function of a user's prior awareness of a trending event. While real-time information appears valuable to all consumers of trending event information, it appears especially so for users new to the event. In our analysis of search logs, we observe that users who click *Trending Queries* links engage less overall with result content and focus more on "up-to-the-minute" content than users who are aware enough of an event to manually enter related queries. Further investigation might examine how user behavior adapts to changes in result presentation, such as promoting a standard result to an *Instant Answer*. These differences point to opportunities for introducing more real-time content into search results for trending event queries, as well as tailoring search results based on measures of users' prior engagement with trending events and use of different classes of online media sites.

What Trends Where, and When?

Trending Queries and Twitter *Trends* are each prompted by a wide variety of triggering events. Our hypothesis that social media content can be leveraged to support real-time search needs rests on an assumption that content is being produced for the same types of events that are being heavily searched and at roughly the same time. In this section, we zoom in from general search behavior to

specific aspects of trending events, comparing events reflected as *Trending Queries* with those appearing as Twitter *Trends*. We compare user activity over time for individual trends across both search and social media. We aim to identify classes of events where social media may be particularly suited for supporting trending event search.

Categorizing Trending Events

In order to explore differences in the kinds of events which are surfaced as Twitter *Trends* or *Trending Queries*, we categorized each trending event according to two schemes: *type* and *topic*. For each event, we used web, social media, and other search tools to find relevant content authored near the trend date to aid in identifying the corresponding real-world event underlying the observed trend.

Method

Two coding schemes were each developed iteratively from the data using a conventional content analysis approach (Hsieh & Shannon, 2005). From a small sample of events, three authors developed two sets of mutually exclusive codes (*type* and *topic*) to apply to each event. The same authors then used each coding scheme to categorize a larger set of 99 events, at which point each scheme was revised. Calculation of Fleiss' κ revealed substantial agreement among the raters for both *Event Type* ($\kappa = 0.71$) and *Event Topic* ($\kappa = 0.82$). One author then manually categorized the remaining events using each scheme.

Event Type. With this coding scheme, we aimed to characterize the nature of the triggering event, capturing aspects such as whether it was anticipated or whether it was continuing while users discussed it. The scheme developed was analogous to the categories proposed by Zubiaga et al. (2011): *News* (breaking news, renamed *Breaking* in this work for clarity), *Meme* (viral conversation topics), *Commemorative* (e.g., birthdays, anniversaries) and *Current Event* (events being discussed as they happened, renamed *Ongoing* in this work). We add an additional label *Unknown* for cases where the triggering event could not be identified or categorized.

Event Topic. We developed a second scheme to represent high-level topical categories. The categories iteratively developed were: *News*, *Entertainment*, *Politics*, *Sports*, *Holiday*, *Deaths*, and *Unknown*.

Results

Table 2 shows the percentage of events trending in each stream by type, with relevant examples. We explored the relationship between trend origin (Twitter *Trend* vs. *Trending Queries*) and event type; pooling low-volume event types (*Meme*, *Commemorative*, *Unknown*) into a single category, a Chi-squared test of independence revealed an association ($\chi^2(2, N=331) = 41.09, p < 0.001$). For events appearing as *Trending Queries*, the vast

Event Type	Examples	Search Engine	Twitter
Breaking	<i>colorado shooter; sherman helmsley; toyota recall</i>	80%	63%
Ongoing	<i>ufc 150; olympics schedule; medal count</i>	17%	11%
Meme	<i>hi boyfriend; hakuna matata; stevie j</i>	0%	9%
Commemorative	<i>selena 20; amelia earhart; national tequila day</i>	2%	7%
Unknown	<i>8 mile; big mac; dear john</i>	1%	11%

Table 2. “Event Type” categories with examples and breakdown for search engine *Trending Queries* vs. Twitter *Trends*.

Event Type	Examples	Search Engine	Twitter
News	<i>aleppo syria; paterno statue</i>	26%	12%
Entertainment	<i>toy story 4; kendrick lamar</i>	9%	18%
Politics	<i>todd akin; paul ryan</i>	9%	4%
Celebrity	<i>stevie wonder divorce; jodie foster</i>	20%	20%
Sports	<i>aly raisman; blind archer</i>	30%	18%
Holiday	<i>national cheesecake day; eid mubarak</i>	1%	2%
Deaths	<i>neil armstrong; binchy maeve</i>	5%	13%
Unknown	<i>hi crush; new zealand</i>	1%	13%

Table 3. “Event Topic” categories with examples and breakdown for search engine *Trending Queries* vs. Twitter *Trends*.

majority were *Breaking* (80%) or *Ongoing* (17%). The majority of Twitter *Trends* were also *Breaking* (63%) or *Ongoing* (11%) events, but this was tempered by the substantial number of *Meme* (9%) and *Unknown* (11%) events. The proportion of events in these categories would likely be higher without our initial filtering step.

Table 3 shows a breakdown of events by topic. Again pooling low-volume topics (*Holiday*, *Unknown*) into a single category, a Chi-squared test revealed a relationship between event origin and topic ($\chi^2(6, N=331) = 46.69, p < 0.001$). *News* and *Politics* occupied twice the portion of *Trending Queries* (26%, 9%, respectively) as they did Twitter *Trends* (12%, 4%). *Entertainment* and *Death* events were much more prevalent on Twitter (18%, 13%, respectively) than in *Trending Queries* (9%, 5%). Again, the large number of Twitter *Trends* marked as *Unknown* points to the relative noise in these trending events.

Temporal Patterns in Search and Social Media

Given a trending event surfaced by either service, we are also interested in the extent to which temporal patterns of real-time information production (from Twitter) and information consumption (seen in search activity) overlap.

Method

For each trending event, we use the associated search query and social media logs to construct two time-series representing hourly query and tweet volumes. We employed a LOESS-based seasonal decomposition algorithm (Cleveland et al. 1990) to isolate the longer-term temporal patterns from the daily variation.

Our categorization method draws specifically on prior work on temporal query dynamics (Kulkarni et al. 2010). Each time-series was first categorized on the basis of whether it contained zero, one, or multiple “spikes.” Those

containing a single spike were further categorized into “wedges” (where popularity rises and falls at similar rates over time), “sails” (left or right) (where popularity’s rise is slow and drop-off is sudden, or vice-versa), and “castles” (where popularity stays at a new level for an extended period after a change) (see Kulkarni et al. (2010) for more detailed descriptions of these categories). We identified an additional category consisting of a wedge occurring within a period of already-elevated activity, which we labeled a “chimney.” Figure 3 illustrates examples of these patterns. For each trending event, we coded the temporal patterns for query activity and for social media activity separately.

Results

Query Dynamics. As shown in Figure 4, the majority of trending events tended to trigger query activity that was wedge- (32.0%) or castle-shaped (27.5%), or flat (14.8%). We performed a Chi-squared test to examine the possibility of an association between event type and temporal query pattern; we pooled event types with low volumes (*Meme*, *Commemorative*, *Unknown*) into a single *Other Type* category, and did the same for the temporal query patterns not mentioned above. We found an association between event type and query pattern ($\chi^2(6, N=331) = 71.90, p < 0.001$). Examination of cell frequencies showed that this difference may have been driven by the *Other Type* events, which were predominantly wedge-shaped or flat.

We similarly tested for an association between event topic and query pattern, again pooling low-volume topics (*Death*, *Politics*, *Holiday*, *Unknown*) into *Other Topic*, observing an association ($\chi^2(12, N=331) = 34.00, p < 0.001$). Examination of cell frequencies revealed that about 40.3% (27 out of 67) of *Celebrity* events were wedge-shaped, and that only 10.3% (4 out of 39) of *Entertainment* events took on a castle-shape. Identifying the typical

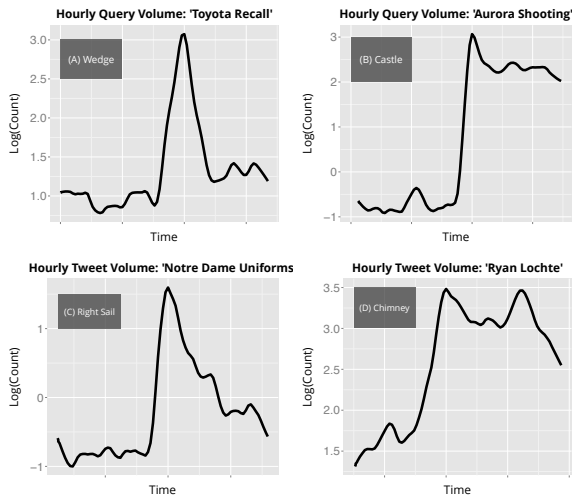


Figure 3. Examples of common temporal patterns for trending events: (a) *Wedge*: Queries for “Toyota Recall”; (b) *Castle*: Queries for “Aurora Shooting”; (c) *Right Sail*: Tweets for “Notre Dame Uniforms”; (d) *Chimney*: Tweets for “Ryan Lochte”

temporal dynamics associated with different topics may be helpful in designing appropriate search experiences; for example, castle-shaped trends (representing sustained interest in an event) might merit creation of richer, extended search “Instant Answers” due to an anticipated sustained volume of interaction.

Social Media Dynamics. As shown in Figure 4, social media patterns were similar to search; most events were associated with wedge- (29.3%) or castle-shaped (28.7%) activity, or were flat (18.4%). Using a Chi-squared test and the same pooling strategy as above, we found an association between event type and social media activity pattern ($X^2(6, N=331) = 67.74, p < 0.001$); we observed that a high number of *Ongoing* events (39.6%, or 19/48) exhibited a wedge-shaped pattern of Twitter activity.

An association was also observed between event topic and social media activity pattern ($X^2(12, N=331) = 51.24, p < 0.001$); we observed that 18 out of 49 (46.2%) *Entertainment* events follow a wedge-shaped pattern. Knowing likely temporal patterns for key event types could enable prioritization of resources or screen real estate for creating trend-specific search result pages. Topics that are *Entertainment*-related and *Ongoing* (such as broadcast events), for instance, may be especially likely to have users contributing simultaneously during a short period, affording additional possibilities in terms of user interaction around content being created.

Correspondence. For most trends (56.8%), patterns of query activity match exactly the patterns observed in social media activity. Of the cases that differ, the most interesting are the events in which activity is flat in one stream but not the other. Of the 49 events with no spike in query activity,




Figure 4. Trending event patterns for query volumes and tweet volumes. The majority of trending events exhibited *wedge* or *castle* patterns of activity with respect to search queries and Twitter updates.

11 (22.4%) exhibited a wedge of Twitter activity. These were predominantly *Celebrity* or *Entertainment* topics with waning popularity (e.g., *Drake and Josh*, *Laguna Beach*, *Limp Bizkit*). Of the 61 events with flat Twitter patterns, 20 (32.8%) exhibited a wedge of query activity; many of these were *Celebrity* topics concerning individuals currently in the limelight (e.g. *Britney Spears*, *Joe Biden*, *Lolo Jones*). Understanding factors that may cause differential trending on search engines versus social media, such as the aforementioned examples, is an interesting area for further research that likely requires supplemental data (such as user demographics) that are beyond the scope of this particular investigation.

The observation that the majority of topics share the same patterns in search and social media activity is encouraging as it lets us know that real-time content production peaks in a similar manner to the information needs of potential real-time content consumers.

Alignment of Social Media and Search Activity

Above, we found that trending events trigger similar temporal patterns of user activity with respect to seeking and sharing information. In order to leverage social media content to support real-time sensemaking about trending events, we must also ask whether these peaks of activity are occurring at the same time.

Method

We find the maximum cross-correlation between the two time-series for an event using a method similar to that developed by Adar et al. (2007). For all values of h in a given range, we shift one time-series by h hours and calculate the correlation; the value of h maximizing the correlation represents the “delay” of one stream relative to another. We set a window for h of 48 hours before and after the beginning of the day when the event first begins to trend in order to avoid matching unrelated peaks. We removed 16 (4.8%) events for which the maximum correlation corresponded to a value of h outside this range.

Results

Social media and search activity patterns aligned strongly, as shown in Figure 5. For optimally chosen values of h ,




Figure 5. Distribution of cross-correlations across trending events for optimal values of h . At the intersection of the pink lines, we see that 90% of trends have $r > 0.496$. The red lines show that 50% of the trends have $r > 0.847$.

50% of trends had Spearman rank correlation coefficient $r > 0.847$, and 90% had $r > 0.496$. As shown in Figure 6, optimal values of h (marked by a red dashed line) followed a roughly normal distribution, with $\mu = -4.3$ and $\sigma = 14.4$ indicating that social media activity around trending events tends to lead query activity by a small margin.

We identified no significant differences with respect to events that originally appeared as *Trending Queries* and those that appeared as *Twitter Trends*. In addition, we did not observe any significant differences between *Breaking* and *Ongoing* events. These patterns tended to be fairly consistent across event topics, with the exception of *Holidays*, which appear to be the only category where query activity peaks before social media activity. This is likely because these are the only category of events that are completely anticipated; users may search when making plans and then share content during or after those activities.

We manually examined the 21 events for which search significantly led social media activity and correlations were high ($r > 0.8$). Five of these represented events from the 2012 Summer Olympics (*Closing Ceremony*, *Medal Count*, *Michael Phelps*, *Nigeria vs. US*, *Watching Olympics*); these events were broadcast in the United States with a tape delay of several hours, possibly introducing noise into our results. We also found that 4 (19.0%) of these events were categorized as “flat” in both streams, meaning that they likely corresponded to unimportant or non-events, despite the strong alignment.

Discussion

In this section, we observed differences in the types of events that appear as *Trending Queries* or as *Twitter Trends*. We observe that *Trending Queries* are more likely to reflect *News* and *Politics* topics than *Twitter Trends*, while the opposite is true for *Entertainment* and *Deaths*. It is important to remember that these comparisons pertain to events which are presented rather than the larger set of trending events which might be identified in each system, as each system may employ some amount of editorial




Figure 6. Distribution of delays (in hours) that maximize the correlation between query and social media activity. Negative values of h indicate that Twitter precedes search. The dotted red line shows the mean $h = -4.3$ hours.

control over which trends to show. In addition, our analysis here focuses primarily on exogenous events, differing from prior analyses (Naaman, Becker, & Gravano 2011).

Looking at the temporal patterns of activity in each stream, we observe that events that are *Ongoing* and/or *Entertainment*-related are more likely to take on a wedge-shaped pattern of activity on Twitter; it is likely that this corresponds with the phenomenon of “live-tweeting” events. In addition, the observation that peaks of search activity generally follow a few hours after peaks of social media activity suggests that this class of events is particularly amenable for providing different experiences during and after an event. One might take advantage of the large numbers of users simultaneously producing content during an event to facilitate interactions and then leveraging the delay before peak search activity to index and process content to better meet information needs.

How Interest in an Event Changes over Time

In order to best support users during trending events, it is critical to understand how user interest shifts over time and when content sought overlaps most with the real-time content being produced. In this section, we explore how user interest in a trending event changes over time. We look at variation in terms used in queries and tweets associated with trending events, and similarly at the variation in domains searched and shared.

Method

We restricted our analysis to trends that exhibit a single peak in both search and social media activity. Specifically, we chose trends that fit either a wedge, castle, or right sail patterns in each stream. We lined each trend up on a common timeline, where $t = 0$ represented peak activity in the medium of origin (i.e., if it was originally a *Twitter Trend*, it was aligned such that $t = 0$ represented the hour when tweet volume peaked for that event). This method allowed for comparisons among many different events.

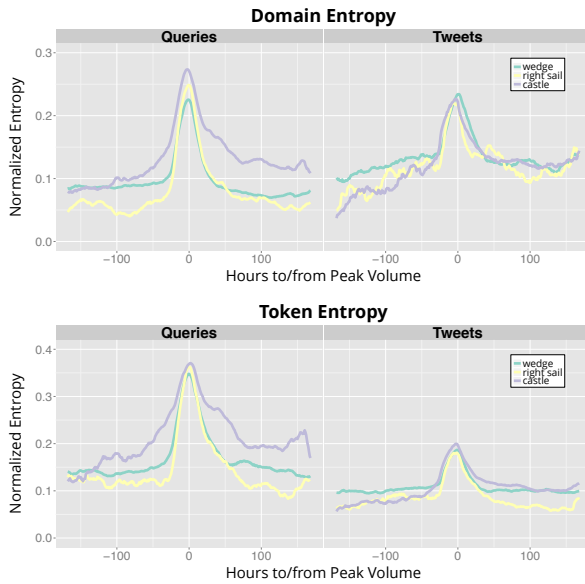


Figure 7. Normalized entropy of content associated with trending events as a function of hours from peak volume. The top charts show entropy for domains clicked in search results (left) and shared on Twitter (right). The bottom charts show entropy for terms in queries (left) and Twitter updates (right).

Entropy. We first examined how variation in the terms and URL domains associated with trending events changed over time within search and social media. We calculated the entropy of this content for each hour, normalizing each value by dividing by the maximum possible entropy in an hour, thus controlling for the greater amount of content shared during the peak time. As before, higher entropy is an indication of higher variability in the domains or terms seen in queries or tweets.

Overlap. In addition to looking at the variation of domains and terms used within social media or search, we also looked at the overlap across the two sources. For a given trend in a single hour, we calculated overlap by taking a set of items (either terms or domains) and dividing the number of common elements by the sum of all such elements referenced in that hour. Because we are making such direct comparisons across streams, we restrict our set of topics further to the 58 topics that exhibited wedge shapes and the 71 topics that exhibited castle shapes in both streams.

Results

Figure 7 shows how the normalized entropy of trending event content changes over the course of a typical trending event with a single peak. We see that variety in tweet content (domains and terms) starts to increase about 24 hours before peak activity, peaks shortly before peak volume, and returns to normal levels soon afterwards. This pattern appears relatively robust across the type and topic of event. We also see in Figure 7 that the pattern is similar




Figure 8. Overlap between content shared and content searched with respect to trending events as a function of hours from peak volume. 8(a) shows overlap between domains clicked in search results and domains shared on Twitter. 8(b) shows overlap for terms appearing in search queries and in Twitter updates.

for castles; even when the conversation continues, it may converge with respect to the topics discussed.

Figure 7 also illustrates differences in the variety of clicked domains and co-occurring terms for search queries, depending on the query volume patterns. Entropy for wedge and castle-shaped trends starts to increase 24 hours before and then decreases symmetrically after peak volumes, returning to normal levels roughly 24 hours afterwards. We observe that for castle-shaped trends, the entropy of associated content appears to remain slightly elevated for an additional 24 hours or so; this prolonged entropy elevation may explain why these trends experience continued activity; conversation may be prolonged by the introduction of new tidbits of information.

In Figure 8, we see the overlap of content searched with content shared on Twitter. Looking at domain overlap, we see a pronounced pattern for castles, where there is a slow drop-off in overlap; this indicates that domains shared on Twitter remain relevant for users simultaneously searching for a longer period of time after the trend has peaked and been identified. This pattern is similar for keywords appearing in tweets and queries, supporting the idea that it will be easier to match tweeted content to search queries in the days after an event has trended.

Discussion

We observe a characteristic spike in entropy of searched and shared content centered around the peak of activity, representing a divergence in user interest during this time. However, during this same period, topics discussed in social media overlap particularly closely with search needs, as well; together these findings point to a need for diverse result content during the period of peak interest in a trending event and for continued potential for support from social media content hours or days after this period, especially for events following a castle pattern of continued activity.

Conclusion

In this article, we investigated the online information dynamics surrounding trending events, as reflected by large-scale search and social media activity. Analysis of post-query behavior logs for queries about trending events indicated possible differences in user preferences for post-result content based on prior awareness of the event. A follow-up survey of users who had recently engaged with trending event content supported this finding, showing that users who were more newly aware of a trending event had stronger preferences for real-time updates.

We hope to invite future work in identifying how informational needs evolve over the course of a trending event. By pairing toolbar data with automated methods for assessing knowledge via browsing behavior (e.g., Pirolli & Kairam 2012), for instance, one might conduct a more fine-grained and large-scale analysis of this relationship.

Using queries from Bing search logs and updates from the Twitter Firehose, we examined the temporal patterns of activity around trending events. We found that information-seeking and information-sharing activity around these events tend to follow similar temporal dynamics, but that social media tends to lead search activity by a small margin of 4.3 hours, on average. We also studied how content associated with trending events changes over time, observing that user interest appears to diverge around the period of peak activity, but that the overlap of content searched and content being shared in real-time appears to increase accordingly.

Together, these findings paint a rich picture of how social media might better serve as a source of real-time content for users searching about trending events. Many current search interfaces which incorporate social media content tend to provide a reverse-chronologically ordered list of keyword-matched updates. Our findings show that there is a lag of several hours between peak content production and peak search activity, meaning that there may be time for more complex indexing and ranking computation to present more relevant “near-real-time” content in search results. This observation may also help search engines in identifying trending events earlier to allow for algorithmic or manual interventions to proactively meet the needs of the majority of searchers.

References

Adar, E., Weld, D.S., Bershad, B.N., and Gribble, S.D. 2007. Why we search: Visualizing and predicting user behavior. *WWW 2007*.

Broder, A. 2002. A taxonomy of Web search. *SIGIR Forum*, 36(2): 3-10.

Cleveland, R.B., Cleveland, W.S., McRae, J.E., and Terpenning, I. 1990. STL: A seasonal-trend decomposition procedure based on Loess. *Journal of Official Statistics*, 6(1): 3-73.

Clough, P., Sanderson, M., Abouammoh, M., Navarro, S., and Paramita, M. 2009. Multiple approaches to analysing query diversity. *SIGIR 2009*.

Dou, Z., Song, R., and Wen, J-R. 2007. A large-scale evaluation and analysis of personalized search strategies. *WWW 2007*.

Efron, M. 2011. Information search and retrieval in microblogs. *JASIST*, 62(6):996-1008.

Evans, B.M., Kairam, S., and Pirolli, P. 2010. Do your friends make you smarter? An analysis of social strategies in online information seeking. *Info. Proc. & Management*, 46(6):679-692.

Gabrilovich, E., Dumais, S., and Horvitz, E. 2004. Newsjunkie: Providing personalized newsfeeds via analysis of information novelty. *WWW 2004*.

Grimes, C., Tang, D., and Russell, D.M. 2007. Query logs alone are not enough. *WWW Workshop on Query Log Analysis*.

Hsieh, H.F., and Shannon, S.E. 2005. Three approaches to qualitative content analysis. *Qual. Health. Res.*, 15(9):1277-1288.

Jones, R., and Diaz, F. 2007. Temporal profiles of queries. *TOIS (ACM Transactions on Information Systems)* 25(3): Article 14.

Kleinberg, J. 2006. Temporal dynamics of on-line information streams. In *Data Stream Management: Processing High-Speed Data*. Springer.

Kotov, A., Kolari, P., Duan, L., and Chang, Y. 2010. Temporal query log profiling to improve web search ranking. *CIKM 2010*.

Kulkarni, A., Teevan, J., Svore, K.M., and Dumais, S.T. 2011. Understanding temporal query dynamics. *WSDM 2011*.

Kwak, H., Lee, C., Park, H., and Moon, S. 2010. What is Twitter, a social network or a news media? *WWW 2010*.

Lin, J., and Mishne, G. 2012. A study of “churn” in Tweets and real-time search queries. *ICWSM 2012*.

Marcus, A., Bernstein, M.S., Badar, O., Karger, D.R., Madden, S., and Miller, R.C. 2011. TwitInfo: Aggregating and visualizing microblogs for event exploration. *CHI 2011*.

Mishne, G., and de Rijke, M. 2006. A study of blog search. *ECIR*.

Morris, M.R., Teevan, J., and Panovich, K. 2010. A comparison of information seeking using search engines and social networks. *ICWSM 2010*.

Naaman, M., Becker, H., & Gravano, L. 2011. Hip and Trendy: Characterizing emerging trends on Twitter. *JASIST*, 62(5).

Pirolli, P. and Kairam, S. 2012. A knowledge-tracing model of learning from a social tagging system. *UMUAI*, 22(4-5)

Ramage, D., Dumais, S., and Liebling, D. 2010. Characterizing microblogs with topic models. *ICWSM 2010*.

Spink, A., Wolfram, D., Jansen, B.J., and Saracevic, T. 2001. Searching the Web: The public and their queries. *JASIST*, 52(3).

Starbird, K., Palen, L., Hughes, A., and Vieweg, S. 2010. Chatter on The Red: What hazards threat reveals about the social life of microblogged information. *CSCW 2010*.

Sun, A., Hu, M., and Lim, E-P. 2008. Searching blogs and news: A study on popular queries. *SIGIR 2008*.

Teevan, J., Ramage, D., and Morris, M.R. 2011. #TwitterSearch: A comparison of microblog search and Web search. *WSDM 2011*.

Vlachos, M., Meek, C., Vagena, Z., and Gunopoulos, D. 2004. Identifying similarities, periodicities, and bursts for online search queries. *SIGMOD 2004*.

Weber, I., and Jaimes, A. 2011. Who uses Web search for what? And how? *WSDM 2011*.

Zubiaga, A., Spina, D., Fresno, V., and Martinez, R. 2011. Classifying trending topics: A typology of conversation triggers on Twitter. *CIKM 2011*.