Narratives: A Visualization to Track Narrative Events as they Develop

Danyel Fisher[*], Aaron Hoff, George Robertson
Microsoft Research

Matthew Hurst
Microsoft Live Labs

ABSTRACT

Analyzing unstructured text streams can be challenging. One popular approach is to isolate specific themes in the text, and to visualize the connections between them. Some existing systems, like ThemeRiver, provide a temporal view of changes in themes; other systems, like In-Spire, use clustering techniques to help an analyst identify the themes at a single point in time. Narratives combines both of these techniques; it uses a temporal axis to visualize ways that concepts have changed over time, and introduces several methods to explore how those concepts relate to each other. Narratives is designed to help the user place news stories in their historical and social context by understanding how the major topics associated with them have changed over time. Users can relate articles through time by examining the topical keywords that summarize a specific news event. By tracking the attention to a news article in the form of references in social media (such as weblogs), a user discovers both important events and measures the social relevance of these stories.

KEYWORDS: blogs, events, trends, time series, topic detection and tracking.

INDEX TERMS: I.7.m [Document and Text Processing]: Miscellaneous. I.3.8 [Computer Graphics]: Applications.

1 INTRODUCTION

A standing challenge in Visual Analytics research is the analysis of unstructured text streams, such as news stories and blog entries. There have been a wide variety of approaches to these problems, each of which has emphasized various aspects of the data. Natural-language processing approaches try to bring out the actors and events of the stories. Other approaches extract keywords, cluster concepts, or arrange stories and themes along timelines.

One particularly interesting area of analysis is the news, in part because it has implications both for analysts and news readers. News stories are a relevant source of current information when taken one at a time; *en masse*, they become a reflection of culturally-important information. Yet examining individual news stories, or even groups of stories, loses out on important aspects of news context. Reading an article in the paper gives little information about several critical, related areas: how the topic of that article has changed over time, and how readers are reacting to the article. As readers of the news, we are interested in the evolution of stories. What we might call the "narrative" around certain themes is shaped by the appearance of articles, and evolves over time: a company releases a new product, and is featured in the news; a presidential candidate enters a race, competes, and weathers scandals. All of these separate stories come together in a unified narrative of the candidate's trajectory.

[*]{danyelf, aaronho, ggr, mhurst}@microsoft.com

The reactions of readers to the news also help us understand the context of the information we are reading.. One of the most accessible sorts of responses to news can be found in blogs, which have recently gained prominence within the VAST research community (and many others): the 2007 Contest, for instance, leveraged blogs as a critical portion of the solution.

In this paper, we present Narratives (Figure 1). Narratives presents a way to view temporally-changing data. It works from a corpus of blog entries that talk about news stories, and so both reflects the articles about a topic and the blogs that comment on these articles. Despite its fairly simple visualization technique, based around a line graph, Narratives allows users to see what additional concepts are most associated with a selected term by displaying closely related terms in several ways.

In Figure 1, the Narratives display compares the fortunes of four presidential primary competitors over the first three months of 2008. The number of references to each candidate's name is shown as a line on the graph; the lines share axes, and so can be compared.

The contribution of this paper is to show a way to piece together this complex of information. By viewing each response to a news story as a single event with multiple keywords, we can visualize the sequences of keywords as a series of simple (but related) line graphs. Unlike much past research, which has largely emphasized a single variable changing over time, our particular challenge is to examine multiple possibly-related variables. We wish to both examine the continuity of themes over time, and also find correlations between themes.

2 RELATED APPROACHES

There have been a variety of approaches to looking at how ideas evolve over time. The information retrieval topic of topic detection and tracking, for instance, looks at how discussions of topics change. We examine the topic detection and tracking

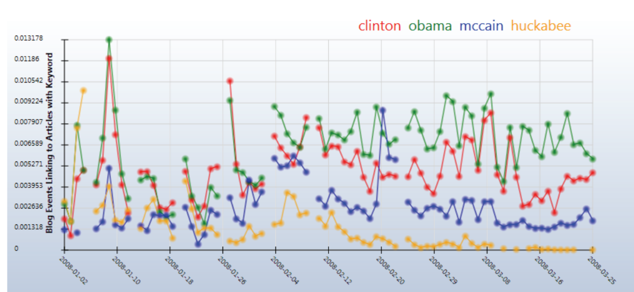

Figure 1. Narratives, showing daily references to four US presidential candidates from January 1 – March 26, 2008. Time passes along the x axis for each candidate; number of mentions of the term along the y. Note that Huckabee (orange) falls off as his campaign ends.

subfield, before turning to temporal visualizations.

2.1 Topic Detection and Tracking

Research in topic detection and tracking [11][16] emphasizes tracking a single topic over time, finding new articles that are related to, but not the same as, a seed article. For instance, NewsJunkie [5] finds articles on the same topic that introduce new concepts over time. Taglines [4] tracks popular keywords, and presents the user with a cascading view of the newest keywords. In those systems, however, we do not know of an attempt to explicitly extract correlation between topics—rather, they approach the problem from the other direction, finding new material that on a related topic. It is up to the user to figure out *what* is new about the incoming material. These systems also do not prioritize popularity—indeed, much of their research claim is that no matter how many new articles pile on, they only report the newest material, rather than the volume of messages.

BlogPulse [6] addresses the issue of volume by allowing a user to generate a visual search of a keyword. BlogPulse presents a line graph of keyword frequency over time, allowing a user to see how often a keyword, or set of keywords, has appeared over time and to compare their volumes. However, BlogPulse is a search-only interface; the visualization cannot help users know what additional keywords to pursue. Other popular web sites devoted to the aggregation, mining, and indexing of social media include Technorati, and Boardtracker, as well as Google's Trends. In an effort to track the topical "buzz" of the source community, these sites tend to offer recent top articles, tag clouds of popular terms, and temporal views of term frequencies. Google Trends offers histories of the volume of keyword searches on Google.

While these projects all explore different ways of extracting topics over time, and several produce line graphs, our approach considers the ways that users can explore this data further. In many ways, we are taking the next step from many of these systems. The closest analog is TimeMines [13], which looks for closely-related stories on a single topic across timelines. Narratives' use of social media, and our flexible notion of correlation, extends the basic notion from TimeMines.

2.2 Temporal Visualizations

The temporal axis is a natural one for many datasets, and a number of visualization techniques have examined ways that data changes over time. Several types of research have examined a single variable changing over time, often looking for periodicity in the variable's values. Carlis and Konstan [3] note that a spiral visualization can help visualize cycles that may occur in data; van Wijk et al [14] cluster short sequences of time with well-known periods, so that the membership of each cluster are largely similar. In their paper, the authors cluster different types of days—holidays, weekends, etc—across fine-grained data. Aigner et al [1], in a broad overview of temporal visualization techniques, warn that such systems can be very sensitive to frequency parameters.

In our system, we do not attempt to locate cycles or periodicity in the data. Rather, we are interested in correlations between streams of data. We are interested in event frequencies for a variety of correlated and uncorrelated events. These events are fairly frequent—at the more than one-a-day level—and so can be evaluated meaningfully through aggregation. A useful taxonomy of these alternatives is articulated by Aris et al [2]. However, this taxonomy allows users to compare independent time series—such as eBay auction prices. Hochheiser [7] uses a timeline-based system to interactively query multivariate time series. In contrast, we are interested in locating correlations that are intentionally collocated in time, as opposed to finding similar patterns between different time series.

Other tools are tuned to visualizing a corpus of text, changing over time. One prominent tool is PNNL's ThemeRiver [9]. ThemeRiver visualizes "themes" from a collection of documents over time as a stacked area chart, smoothly curved to suggest a "river" of topics. Topics that have become prominent are visualized as larger, while topics that have disappeared from the public eye become thin. ThemeRiver can present two topics that are simultaneously interesting at a particular time by showing they had bulges at similar times. For example, [9] reports that the word "earthquake" appears at the same time as the word "turkey"; the analysts must then investigate those themes to find out whether the linkable is real or spurious; they learned that the simultaneous spikes coincide with an earthquake in Turkey. PaperLens [8] visualizes topics from a digital library and shows how each topic changes over time. The topics are derived from clustering based on the text in paper titles and abstracts, as well as other metadata. Both of these tools are intended for fairly small sets of themes, and show them all at once: ThemeRiver depends on a hand-coded list of relevant themes in the data; PaperLens derives its themes from clustering.

NameVoyager [15] represents changes over time as an area chart, allowing users to compare the popularity of names over time. It is able to deal with a much larger corpus by incorporating a smooth animation that allows users to filter through the corpus to the portions they find most interesting. However, this zooming is largely alphabetical, making it very difficult to discover whether Octavian and Adolphus were popular names at the same time.

Our approach begins with a seed term, and then highlights corresponding high points; this would thus clarify the correlations between trends.In-Spire [10], another text processing project from PNNL, clusters articles on a two-dimensional plane based on extracted keywords. Every article is located in exactly one place, based on the keywords in the article. Thus, an article is represented by its keywords, and (conversely) keywords are, by their nature, clustered together. This fundamentally means that an article can only sit in one cluster; there is no mechanism for an article having sets of keywords that are not already clustered.

In-Spire has a notion of time; the document corpus can be sliced by time. Each time slice can be viewed separately or strung together into an animation. This provides a limited sense of continuity, and makes comparison difficult between remote times.

Our approach continues In-Spire's notion of identifying keywords associated with related groups of articles; however, we add a dimension of time, and allow users to explore keywords that are decoupled from each other, and to explore relationships that may not surface in a clustering approach. Thus, if an article joins two disparate topics, it can still be closely associated with the shared terms.

3 Design Requirements and Scenarios

In order to describe our approach, we begin with a pair of simple scenarios.Our first scenario is based around an analysis task. At our organization, we are working with business analysts who have expressed interest in trying to understand business trends and public relations. These analysts want to understand how certain branding terms have changed over time: they want to know what the history of a particular product has been in the marketplace. Has the product been closely associated with other products, or with particular announcements? What sorts of events tend to generate buzz—and what promotions sank without a trace? Seeing what the top concepts associated with the product over time would help them decode their product's history, and would allow them to carefully target future marketing efforts.

Our second scenario follows "Chris," who follows the news closely for recreation. Chris finds out that the CIA (the American Central Intelligence Agency) is one of the most discussed terms in the current blogosphere, and wants to understand what is driving the story. Chris sees several times when discussions about the CIA have become more vigorous, and sees that many of the discussions of the CIA mention "interrogation." Each of these peaks, however, is associated with a different controversy, and thus has different keywords with it. He checks several of the past high points in detail, and learns what news stories have driven them. Last, he can look at the most recent peak to find out what blog entries and news stories are driving the most recent attention.

We want to systematize this discovery and make it possible to know what concepts have co-occurred in the past. Our design has been developed in collaboration with analysts, who need to understand how a topic has developed over time, but also has aspects oriented toward news enthusiasts, who want to see the evolution of the latest and most interesting stories.

4 System Design

In order to address these scenarios, we have designed Narratives. Narratives is a visualization of the changes in stories over time. It allows users to view news stories in both temporal and social context by:

- Showing how the social media responses to these stories and they underlying concepts have changed over time
- Showing other important concepts associated with the story
- Showing who is commenting on the stories.

Narratives combines keywords from news articles with reactions from social media, visualized over a timeline, to show the evolution of a story over time. It them provides several tools to allow users to investigate correlations between keywords, and suggests closely-related keywords.In this section, we discuss the Narratives infrastructure and interface, and articulate the ways that Narratives displays correlation between topics..

4.1 Data Acquisition

Narratives is based on the Social Streams architecture, a platform being developed at Microsoft's Live Labs. One of the core components of the platform is a real time data acquisition system which crawls many social media content types, including weblogs. This component monitors ping servers and crawls feeds in response to ping events. For blogs that do not provide regular pings, it performs additional, scheduled crawling. Partial feeds are augmented with an intelligent scraper that reads complete posts based on permalinks. In addition, the system monitors other social media sources (such as Usenet articles and Twitter feeds).

The data acquisition system includes a news article scraper. It examines each element in the streams, looking for links that may

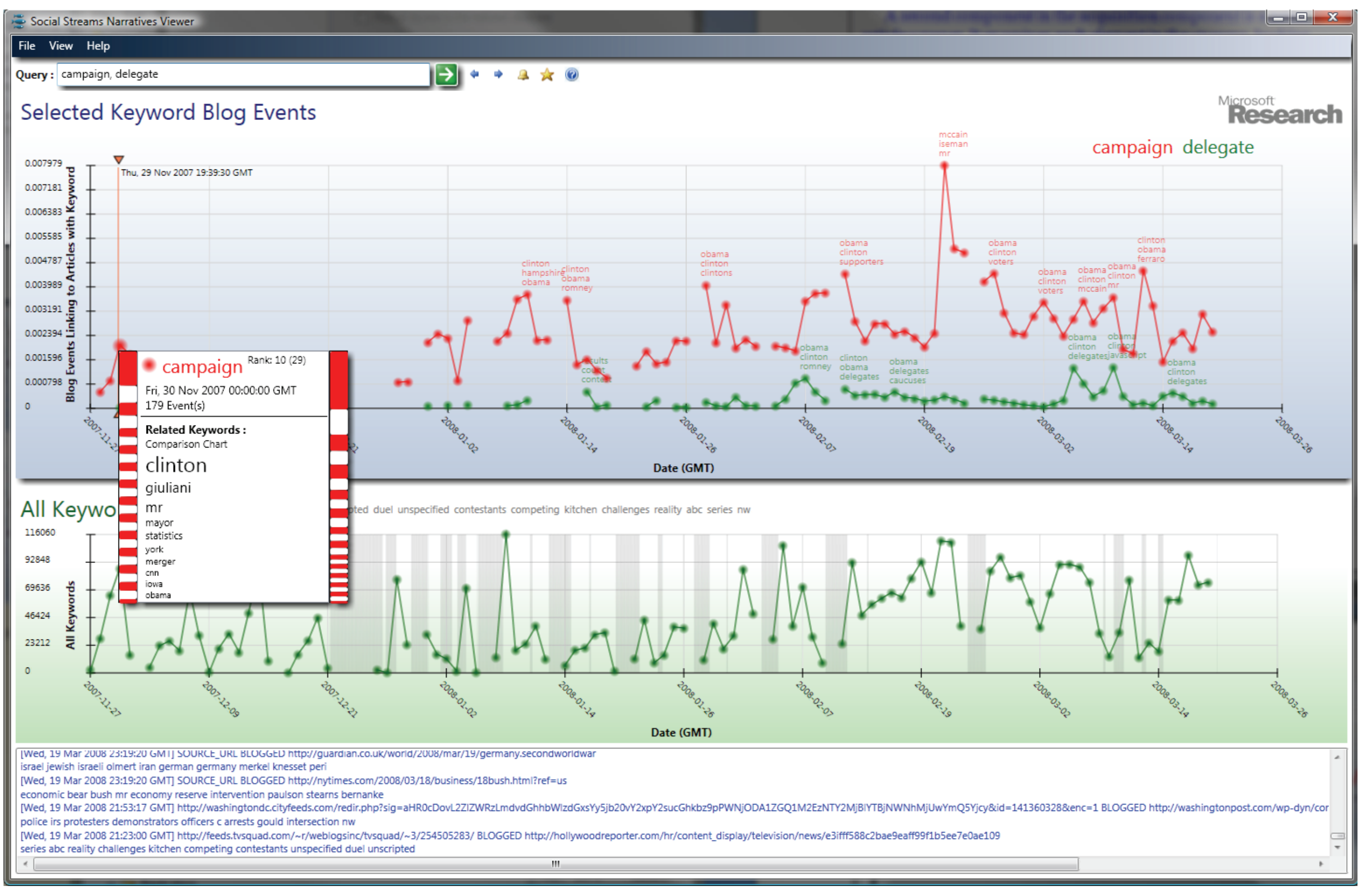

Figure 2. The *Narratives* user interface is divided into three parts. *Top*, the search results window, showing the selected blog events that mention articles with this keyword. *Middle,* the total flow of articles into the system. *Bottom*, the current stream of new articles.

refer to news stories. A classifier determines if a URL refers to a news article or some other type of content; news articles are downloaded. An article summarizer (based on a unigram likelihood model **Error! Reference source not found.**) parses and examines the downloaded articles and extracts significant keywords from each. We refer to these as *topic keywords*.

The source for Narratives, then, is a stream of input of *events*; each event consists of the URL to a *referring blog*, the URL of a *target article*, and a series of up to ten of the most significant topic keywords. Narratives stems and archives these keywords, along with the blog and news links that generate them.

Narratives collects between 50,000 and 100,000 new keyword events from the Social Streams back end over the course of each day: an average of 3,000-5,000 blog posts and 4,000-7,000 article URLs. (The number of articles is an over-estimate, as Narratives does not currently remove duplicate articles with identical text.)

4.2 Temporal Visualization

The core visualization is a line graph showing fluctuations in popularity, labeled "Selected Keyword Blog Events" in Figure 2. The graph portrays the fraction of that days' traffic that refers to the specific keyword. Increases are times when the keyword is more popular; decreases are times when the keyword is less so. In Figure 2, we compare the traffic for the keyword "campaign" to the keyword "delegate." We interpret high points on the line as places where bloggers have commented more frequently on articles about that topic keyword.

Each point is a percentage of all counted events for its day. For instance, on February 21, the term "campaign" (the peak of the red line) spiked to 0.8% of all events in the database. At that particular high point, we detected 893 different events with the word "campaign." The New York Times had just run a story on John McCain's relationship with a lobbyist; the spike consisted mainly of commentators discussing the story.

This comparison allows us to see how a small group of user-chosen words have changed over time relative to each other.

4.3 Four Forms of Correlation

Several other tools, such as ThemeRiver, also allow side-by-side comparisons between terms. Narratives is distinct in that it offers ways to distinguish when a pair of terms are co-occurring by coincidence, and when they are correlated.

Date-based correlation. The first type of correlation that Narratives displays is that of *date based* correlation. It is intended to allow a user to ask "what else was happening on this date around this particular keyword?"

Every point in the data is caused by a set of events; each event has a number of keywords associated with it. Narratives orders the list of keywords from these events by frequency, and presents the top ten. The text is sized slightly to suggest relative frequency (a technique borrowed from the popular "tagclouds"). An example of this is in the pop-up on Figure 2. These are visible only interactively. In addition, the graph's peaks are also annotated with the top three associated words. This allows the user to read the graph rapidly, and see at a glance which peaks have promising events associated with them. Figure 3 illustrates the way that spikes are annotated with their most-associated words.

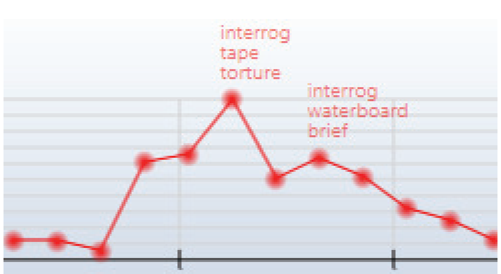

Figure 3. Detail from *Narratives* search for cia*:* the top three terms associated with spikes.

Numerical correlation. Second, Narrative displays numerical inter-concept correlation by computing a correlation coefficient between every pair of terms visible onscreen. For each term displayed on the line chart, the event counts are binned by GMT day, using the timestamp for each event. A Pearson correlation coefficient is computed for each possible pair of vectors by

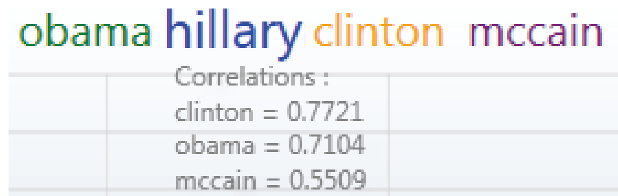

Figure 4. Correlations of "Hillary" with "Clinton", "Obama", and "McCain".

treating the daily event count for a term X as a random variable. A correlation coefficient of zero suggests that two terms are likely to be unrelated; a correlation of one indicates that any change in the daily event count for one term is matched by a proportionate change in the other. Words that have a high correlation coefficient tend to appear in the same stories, or share a fate. An example in Figure 4 follows the primary campaign: "Hillary," the first name of one candidate, is well-correlated with her last name, "Clinton," and her primary opponent "Obama", but less so with "McCain." Narratives displays the correlation between the set as a dropdown.

Most-correlated terms. Narrative can also display the most-correlated terms for a given word. Fortunately, this can be done fairly efficiently. The database collects a list of all articles that the requested keyword is associated with, and then collects the frequencies of all keywords associated with those articles. This operation is far less arduous than calculating the correlation coefficient, but produces the same result: the terms that co-appear the most often with the search term.

The system then displays up to the ten most-correlated terms with that term. As of March 26, the top ten words associated with "marriage" are "same-sex", "couples", "gay", "state", "court", "married", "divorce", and "unions." Needless to say, there are many stories about "court" or "state" that have nothing to do with "marriage;" visualizing all of these lines on a shared coordinate system helps identify their separate spikes. We might expect, for instance, that a dramatic legal case would drive a great deal of traffic around the word "court," but not "marriage." The last three terms in Figure 4 are the most-correlated terms with "Obama" across the database.

Dependant correlations. In the "marriage" case above, the term "court" would add a rather bumpy line to the curve, preventing the user from easily detecting places where the term "court" and "marriage" both co-occurred. Narratives includes an ability to visualize dependant correlations: for each correlated term it returns, it presents the curves for the conjunction of the two terms: thus, it shows only articles where both terms occurred.

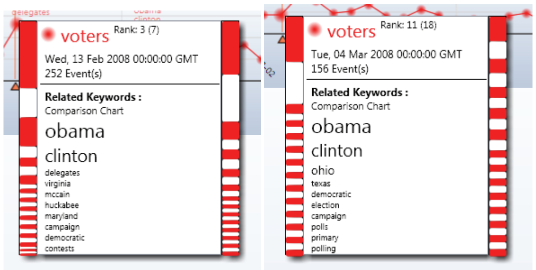

Figure 5. Two different source maps for the term "voters." In the first, two stories (one each from CNN and MSNBC about the "Potomac Primary") are responsible for half the events; in the second, no single story has the lion's share.

For example, in response to the query "marriage," the system displays a curve for only stories where "gay" and "marriage" co-occurred; where "state" and "marriage" co-occurred, and so forth. An example is below, in Figure 7: the term "google" is most correlated with "microsoft", "yahoo", and "search" during the relevant time period. In order to display *only* those occasions of "search" that are linked to Google, the system is showing times when the words co-occur. Figure 6 shows this process computed manually, by requesting co-occurring pairs of terms.

This distinction between types of correlations allows us to discover when co-occurrences between a pair of words are coincidental, and when they are a part of the story. This is a critical difference from tools such as ThemeRiver, where these relationships can be hard to discover.

4.4 Understanding Readership

The last core feature makes it possible to understand the distribution of readers who are commenting on the news story. The lines in Narratives visualize the count of links from blog readers to news stories. As a result, they reflect the degree of interest that these stories, rather than reflecting the number of stories published. Many newspaper stories may yield few, or no, events; on the other hand, a single influential story may yield hundreds of events.

Clicking on a single point yields the list of events—blog entries and the news stories to which they refer—that generate that point (see Figure 5). The candy-striping suggests the distribution of the sources (on the left side) and the distribution of targets (on the right). The striping is limited to the top twenty-five sources or targets, sorted by their contribution to the total event count for that day. Note that in our system, all Usenet and Twitter posts appear as one entry, usually the largest of the set.

Broad stripes on the left side indicate a disproportionate number of links coming from a small number of blogs. Broad stripes on right side display a distribution of which newspaper URLs are being discussed. In Figure 6, on the left, nearly half of the mentions of the word "voters" are coming from two stories (about an upcoming primary). These two stories are sufficiently interesting to drive most of the links. In contrast, on the right, the sources are more-broadly distributed.

4.5 Additional Features in Narratives

Across the top of the screen are several useful tools that help ease interaction with the system. Several of these are useful for our scenarios, and came out of discussions with users.

Alarms. The Social Streams architecture used in this study notifies the application of social media events in near real-time. A user may wish to study the evolution of a concept, or the response of the blogging community to a significant media event on a time scale of minutes or hours. We allow the user to create and manage a collection of terms of interest. When the application receives an event that contains one of these terms in its topic keyword list, a brief audio alert is sounded, along with a visual indicator of the event.

Top 20 List. For the news junkie scenario above, it is critical to know which terms are at the top of the chart. A drop-down top 20 list is triggered by a star icon. The user can choose to look at the top 20 terms for various time intervals, from most recent half hour up through the most recent week. These can be substantially different lists: during the daytime, the last half-hour of information tends to follow small stories and spikes, while the last seven days tends to track larger trends.

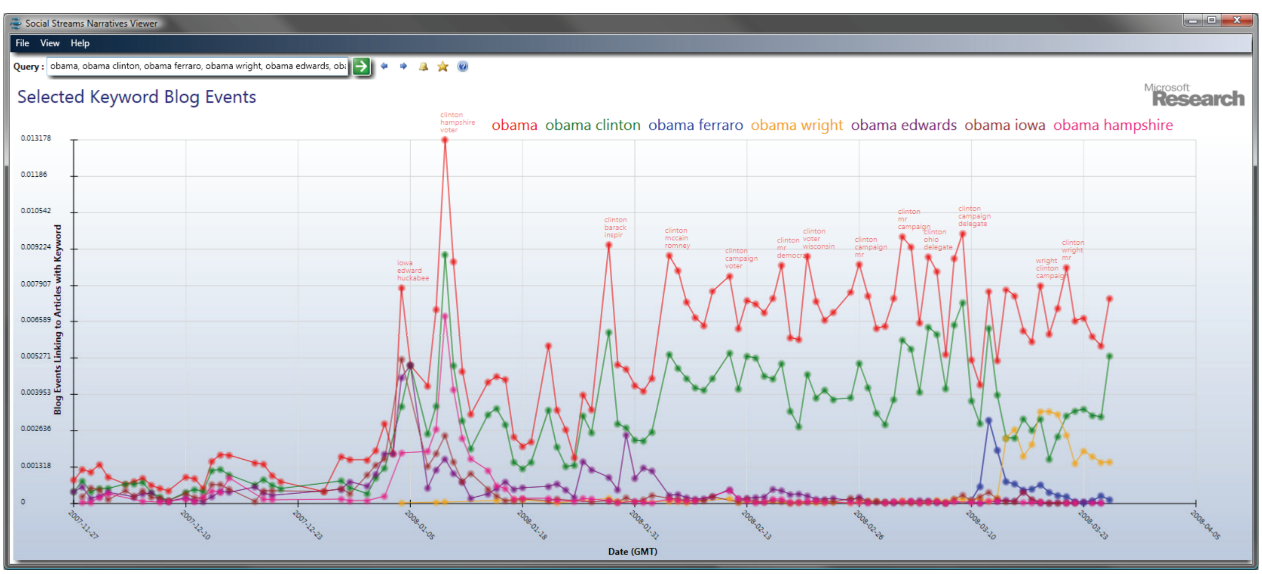

Figure 6. The Narratives interface showing Obama and some relevant keywords, sharing axes.

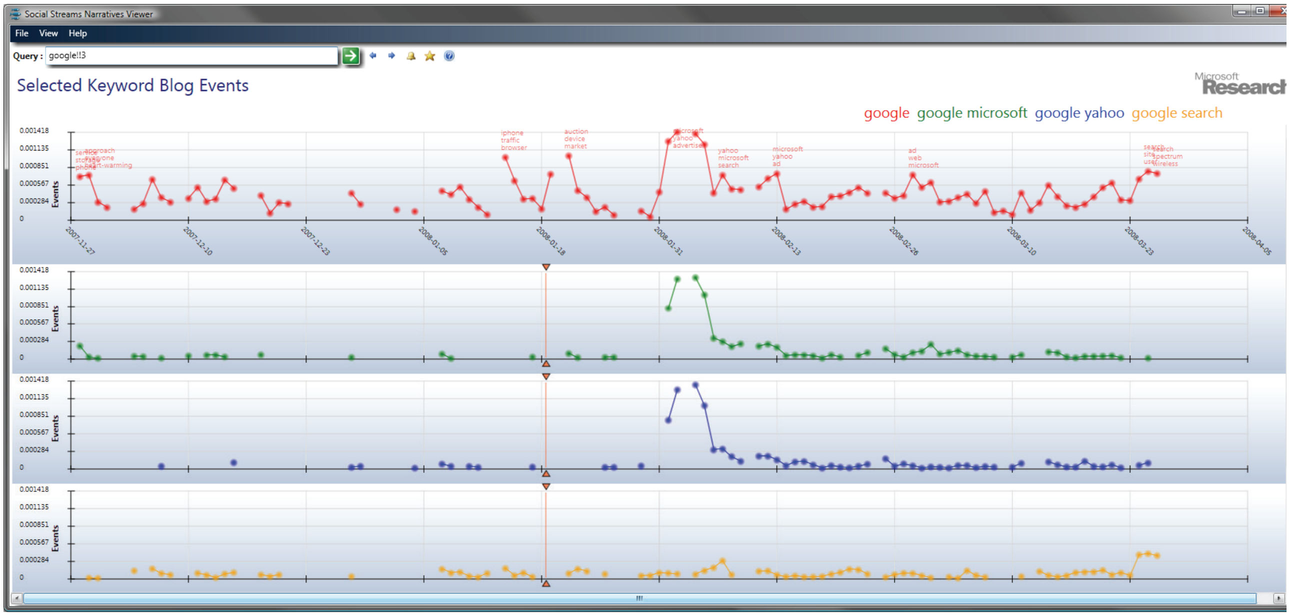

Figure 7. The top three words associated with “google” are,“microsoft,” “yahoo,” and “search.” Note that Google’s biggest spike in attention was in reaction to Microsoft’s attempted purchase of Yahoo. (The other spikes all relate to an FCC auction for spectrum)

Stacking Related Terms. We have experimented with both plotting related terms on the same axes (as in Figure 6) and plotting them on separate parallel plots (as in Figure 7). We have found that they can each be useful, and so we allow the user to select which to use.

5 Evaluation

We discuss the evaluation of Narratives from two different approaches. First, we discuss several interesting findings that have come out of the Narratives data that help better understand the particular value of the system. These findings suggest that Narratives can show novel aspects of the data, and come in part from our discussions with analysts who are trying to learn about ways to interpret blog data.

As we are trying to compare our ability to cluster over time to other solutions, we ran the same dataset through PNNL’s In-Spire [10]. Our results show interesting differences between In-Spire and Narratives, and suggest that the social media emphasis brings a different set of articles to light.

5.1 Illustrative Findings

The story of Obama. One of our key goals was to be able to visualize a narrative—the story of how news articles have shaped a topic, and how social media responses to the topic have changed over time. In figure 6, we trace the history of one particularly prominent topic, the primary run for the presidency by Barack Obama. In this description, we illustrate how Narratives can be used to show how a story does, in fact, progress over time.

Narratives has data starting from November 2007, shortly before the election season started heating up. In late 2007 and early 2008, Obama was just one of many candidates: the top few words those days changed day to day, and Obama never ranked above the twentieth-most-popular name.

At the time of the first primaries (in brown and pink: “Obama Hampshire” and “Obama Iowa”), Obama’s surprise win pushed him to near the top of the stack—and while he has never since covered as large a percentage of the conversation (1.3% of the news), he started being an important part of the story. With Edwards’ campaign’s collapse (purple), Obama surged to the top of the news pack and stayed there. Throughout the primary, the top word after Obama has been consistently Clinton (in green): few mentions of Obama do not also mention his chief competitor. (Their correlation is 0.913).

In March of 2008, several striking events shook the campaign: “Ferraro” (blue) made comments that were perceived as racist, and comments from Obama’s controversial former “Wright” (yellow) pastor came to light, which Obama defended.

In this example, we have manually selected words to view; we populated those choices from the words most associated with peaks. We did not use the mot correlated terms, because we wanted to reconstruct the time-specific features of the story.

While Narratives does not currently implement a tool for picking out stories that are relevant at more particular periods of time, a straightforward relevance algorithm, such as TF/IDF or Kleinberg’s bursts [11], could assist in identifying candidates.

Google gets swept up. As in the first scenario, above, our conversations with business analysts suggested that one important task is to understand how companies are perceived. One way of examining that is to assemble sets of ideas that are closely linked. It is not uncommon to assemble “clipping files,” for example, which highlight prominent media and blog coverage of important events. Narratives can act both as a clipping file, and can also aggregate the relative importance of events. We examine the amount of discussion around one company.

Google is a prominent search engine company that produces a wide variety of applications. We initially expected Google’s most related terms to be linked to the company’s products; however, what seems to have truly excited bloggers who linked to news are two sequences of events.

The smaller was an FCC spectrum auction, at which Google first bid, then lost. These are visible from the small labels at the peaks of Figure 7, which are labeled with terms like "spectrum" and "auction". One peak labels the release of Apple's iPhone, with built in maps (keywords: "iphone", "traffic", "browser").

We were surprised at the larger, and more sustained, peak of attention. This came as a result of Microsoft's attempt to acquire Yahoo, at the end of January. This possible takeover was seen as a competitive threat to Google, and so Google was invoked frequently in these discussions. The search in this case is simply a dependant correlation for the three words most associated with Google: 'google microsoft', 'google yahoo', and 'google search'.

5.2 Comparison with In-Spire

Narratives is not directly comparable to other visualization systems; none the less, we wanted to try to see how it compared to In-Spire, which one of the analysts we worked with was exploring.

We stored archival copies of all news articles that Narratives used in a database; we then fed this archive into a copy of PNNL's In-Spire. For this experiment, we used only article content (not titles), and disregarded article titles. While we could not directly simulate the effect of social media popularity, only articles that appeared were referred to at least once within the blogs was included in the In-spite dataset. Figure 8 shows the diagram generated by In-Spire, based on about 100,000 news articles from February 1 through March 25, 2008 As expected, what is revealed as important in In-Spire is somewhat different from what is revealed as important in Narratives. In this view, time is stripped away: this view shows the set of all articles in the archive.

In-Spire and Narratives cluster the data somewhat differently. In-Spire is able to disambiguate terms: thus, for instance, the terms "development, environment" are visible; together, those terms imply software engineering. Narratives is unable to separate these terms from the construction and nature contexts, respectively.

On the other hand, In-Spire lends great salience to several central clusters that can be difficult to interpret. After some investigation, the hundreds of articles in "people, time, re" seemed to have little to do with each other.

Narratives' analysis is based both on news stories and on reactions to those news stories in blogs. For example, Narratives indicates much more interest in the primaries than the raw news feeds suggest: the top few words within *Narratives* have consistently referred to current political events. This social view acts as a powerful filter on the data, and the immediacy of the temporal information allows us to see what garnered the attention.

In addition, as Narratives is based on vocabulary, rather than clustering, it does not run the risk of users being unable to interpret where peaks come from: a peak is a direct count of articles with a given word. On the other hand, Narratives is unable to view the entire dataset at once—a feature that the analysts found desirable.

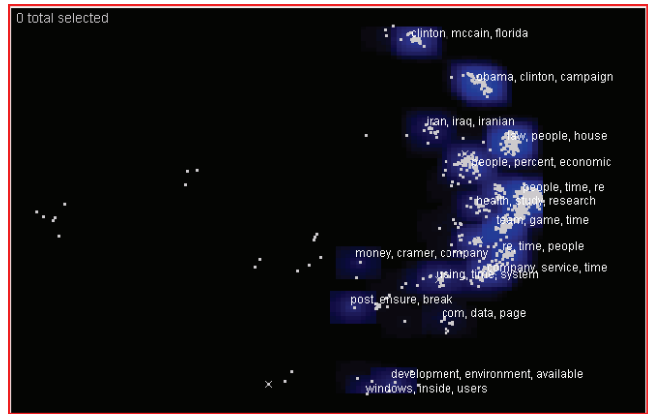

Figure 8. In-Spire analysis of news articles from February 1 to March 26, 2008.

6 Conclusion

Narratives is a simple interface that straightforwardly presents trends in keywords over time. It maintains a broad selection of correlation types: highlighting related keywords at peaks, presenting keywords at all points, and displaying correlated terms both dependently and independently. This selection makes for a powerful exploratory tool that presents a novel interface into the dataset. Its flexibility allows users to explore broadly through a dataset, and to better understand the relationships between articles, bloggers, and the evolving narrative. In particular, dependant queries are a novel contribution.

Narratives continues to be a work in progress as we work with business analysts and news enthusiasts to better understand their needs.

It is clear that the domain of data in blogs and news stories needs to be expanded. Narratives bypasses topics that blogs are discussing that have not made the news, and has a limited view of the breadth of news stories.

Through its four types of correlation, Narratives manages to present word clustering in a way that other tools have not managed: to show relationships between terms over time. This cross-linked analysis has turned out to be a useful way to watch how social media changes over time.

Acknowledgements

We thank Jim Holt, of Microsoft's Public Safety Initiative, who helped us understand how this application might be used, and who explored In-Spire with our dataset. We thank the Social Streams team for their help keeping our application running; and Bongshin Lee, Roland Fernandez, and Lev Nachmanson for their advice. Last, we thank the anonymous reviewers for their useful suggestions. In-Spire is copyright PNNL/Battelle.

References

[1] W. Aigner, S. Miksch, W. Müller, H. Schumann, and C. Tominski. Visual Methods for Analyzing Time-Oriented Data. *IEEE Transactions on Visualization and Computer Graphics* 14, 1 (Jan. 2008), 47-60.

[2] A. Aris, B. Shneiderman, C. Plaisant, G. Shmueli, and W. Jank. Representing Unevenly-Spaced Time Series Data for Visualization and Interactive Exploration. *Proceedings of the International Conference on Human-Computer Interaction (INTERACT 2005)*, 835-846, 2005.

[3] J. Carlis and J. Konstan. Interactive Visualization of Serial Periodic Data. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST)* 1998.29-38.

[4] M. Dubinko, R. Kumar, J. Magnani, J. Novak, P. Raghavan, and A. Tomkins. Visualizing tags over time. *ACM Trans. Web* 1, 2 (Aug. 2007), 7.

[5] E. Gabrilovich, S. Dumais and E. Horvitz (2004). NewsJunkie: Providing personalized newsfeeds via analysis of information novelty. In *Proceedings of WWW'2004,* 482-490.

[6] N. Glance, M. Hurst, K. Nigam, M. Siegler, R. Stockton, and T. Tomokiyo. 2005. Deriving marketing intelligence from online discussion. In *Proceedings of the Eleventh ACM SIGKDD international Conference on Knowledge Discovery in Data Mining* (Chicago, Illinois, USA, August 21 - 24, 2005). KDD '05. ACM, New York, NY, 419-428.

[7] H. Hochheiser, B. Shneiderman. Dynamic Query Tools for Time Series Datasets: Timebox Widgets for Interactive Exploration. *Information Visualization (2004)* 3*,* 1-3. Palgrave MacMillan.

[8] B. Lee, M. Czerwinski, G. Robertson, and B. Bederson, Understanding Research Trends in Conferences Using PaperLens. *CHI'05 Extended Abstracts on Human Factors in Computing Systems*, 1969-1972, 2005.

[9] S. Havre, E. Hetzler, P. Whitney, P., and L. Nowell. ThemeRiver: Visualizing Thematic Changes in Large Document Collections. *IEEE Transactions on Visualization and Computer Graphics* 8, 1 (Jan. 2002), 9-20.

[10] E. Hetzler and A. Turner. Analysis Experiences Using Information Visualization. *IEEE Comput. Graph. Appl.* 24, 5 (Sep. 2004), 22-26.

[11] J. Kleinberg. Bursty and Hierarchical Structure in Streams. *Data Mining and Knowledge Discovery.* 7, 4 (Oct. 2003), 373-397.

[12] M. Marcus, G. Kim, M. Marcinkiewicz, R. MacIntyre, A. Bies, M. Ferguson, K. Katz, and B. Schasberger. The Penn treebank: Annotating predicate argument structure. In *ARPA Human Language Technology Workshop*. 1994.

[13] R. Swan and D. Jensen TimeMines: Constructing timelines with statistical models of word usage. Papers of the ACM SIGKDD 2000 Workshop on Text Mining. pp. 73-80.

[14] J. van Wijk, E. van Selow. Cluster and Calendar based Visualization of Time Series Data. In *Proceedings of the 1999 IEEE Symposium on Information Visualization* (October 23-26, 1999*)*. INFOVIS. IEEE Computer Society, Washington, DC.

[15] M. Wattenberg. Baby Names, Visualization, and Social Data Analysis. In *Proceedings of the 2005 IEEE Symposium on information Visualization* (October 23 - 25, 2005). INFOVIS. IEEE Computer Society, Washington, DC, 1.

[16] J. Yi, Detecting Buzz from Time-Sequenced Document Streams. In *Proceedings of the 2005 IEEE international Conference on E-Technology, E-Commerce and E-Service* (March 29 - April 01, 2005). IEEE Computer Society, Washington, DC, 347-352.