

News Sync: Enabling Scenario-based News Exploration

V.G.Vinod Vydiswaran¹, Jeroen van den Eijhof², Raman Chandrasekar³,
Ann Paradiso⁴, Jim St. George⁵

¹University of Illinois, 201 N. Goodwin Ave, Urbana, IL 61801, USA

²Information School, Box 352840, University of Washington, Seattle, WA 98195, USA

³Evri.com, 71 Columbia Street, Suite 300, Seattle, WA 98104, USA

^{4,5}Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA

¹vgvinodv@illinois.edu, ²jeroen@uw.edu, ³mickeyrc@hotmail.com,

⁴annpar@microsoft.com, ⁵jamesg@microsoft.com

ABSTRACT

News consumption patterns are changing, but the tools to view news are dominated by portal and search approaches. In this paper, we suggest using elements of search, visualization, natural language processing, and machine learning to provide a more captivating, sticky news consumption experience. We propose a novel use-case driven approach to present news and present *News Sync*, a system that addresses three specific news exploration scenarios where a user wants to catch up on news from a particular time period, keep in touch with news from specific locations, or follow the lives of celebrities. The news experience is enhanced by clustering news articles and allowing users to interact with and share stories of interest, and filter results on specific dimensions such as time, location, and key entities. User deployment studies suggest a distinct preference for an interface that supports exploration and visualization of news articles.

Keywords

News summarization, clustering, news adaptation, news exploration, news interfaces.

INTRODUCTION

The news landscape has undergone major changes with the advent of online media. While the readership of traditional newspapers has declined over the past few years, the consumption of news over the Internet has increased significantly. In a March 2010 survey of US Internet users (Gather, 2010), it was found that the Web/Internet is by far the most popular source to find news (49%), as compared to television (32%) and newspapers (9%). Further, as with other kinds of online information, the dominant mode of accessing news online is through search. According to the

June 2010 Pew Research biennial news consumption survey (Pew Research Center for the People and the Press, 2010), henceforth referred to as the Pew survey, 34% of general public use search engines for news, and three of the top four most frequented websites for news are search engine websites. It suggests that, even though there are several dedicated news portals, consumption of online news is triggered primarily through search queries. Search engine companies take advantage of this user behavior by integrating relevant news results with Web search results for news-related queries and providing dedicated news verticals and topic-specific news web pages.

Treating a news query as just another query, however, restricts users to just a set of ten (usually recent) news articles. News, however, is a living entity with a rich past and future. It deals with interactions between other entities, such as people, places, events, and topics. Any good news presentation must encourage users to explore these dimensions. In this paper, we propose a novel use-case driven approach to present news and list key design goals for such a system. We describe the design decisions we took to build a prototype, called *News Sync*, to explore an archive of twenty years of News York Times articles. We study how users interact with the system, using a combination of usability studies and user feedback, and show that users prefer using an exploratory interface in terms of task success and user satisfaction.

SURVEY OF ONLINE NEWS SOURCES

Traditional search companies form the largest and most frequently visited online news destinations. According to the Pew survey, *Yahoo!* (28%) is the most used website for news, ahead of television news leader, *CNN* (16%). Other search portals, *Google* (15%) and *MSN* (14%) compete for the top news websites. Over the past two years, search engine websites have gained in popularity over traditional news media conglomerates. Two-thirds of those polled in the Pew survey said they used search engines to find news on a particular subject, but only 10% regularly got news from customized webpages (e.g., *iGoogle* or *My Yahoo!*) or through RSS feeds. Social networking sites such as *Facebook*, *MySpace*, and *LinkedIn* also played a role in disseminating news headlines – 19% of those polled got

This is the space reserved for copyright notices.

ASIST 2011, October 9-13, 2011, New Orleans, LA, USA.
Copyright notice continues right here.

news from social networking sites regularly or sometimes, and some also got news occasionally from *Twitter*.

In the following sections, we briefly survey different online news sites, and present a summary of some of their features.

News Aggregators

Traditional search engines aggregate news from many sources and categorize them by news categories, such as World (News), Business, Sports, etc. Aggregating news from multiple sources helps them present a multimodal view of a news story that includes videos, photos, and live feeds. Sites such as *Yahoo! News* prioritize news sources and differentiate news coming from different sources. Users can select sources for each category, and news stories from just these sources are displayed. Other sites, such as *Bing News* and *Google News*, present news from many sources as clusters. This allows them to also incorporate stories from other sources, such as blogs, *Twitter*, and *Wikipedia*.

Traditional News Media Online

While readership and viewership of traditional newspapers and television news diminished over the past decade, their online versions have seen an increase. The average monthly reach of web newspapers among Internet households has increased from 27.4% in 2004 to 40.9% in 2008 (Nielsen/Net Ratings). Many television channels now make news clips available online either on their sites or on other video-sharing sites like *YouTube*. Newspapers have augmented their online content with videos and photos to visually appeal to younger readers on the Web (e.g., see Figure 1). With the news consumption moving away from print media, some newspapers, such as the *Seattle Post-Intelligencer*¹, have gone web-only, while other news outlets, such as *The Huffington Post*², *Newser*³, and *Seven-Sided Cube*⁴, present editorial and blog content as independent news online. In addition, a significant portion of news is community-generated. Sites such as *NewsVine*⁵ and *GlobalReporter*⁶ allow users to post (and rate) news local to their community, while news outlets, such as *CNN iReport*⁷, have accepted this notion of grass-roots journalism and allow users to post news videos.

News on-the-fly

News consumption patterns have also changed over the past decade. Rather than setting apart time to access news, users



Figure 1. The Washington Post online edition (retrieved June 10th, 2010). Note the integration of (a) video, (b) images, (c) live feeds, and (d) social networking sites.

consume news throughout the day. The Pew survey found that 76% of all online users (62% of public) say they come across news online when they have been on the Web for another purpose. This is especially true of younger consumers, who typically follow links to news stories, rather than going directly to news sources by themselves. Social networking sites have played an important role in pushing news content to the web users. Sites such as *StumbleUpon*⁸, *Digg*⁹, *Reddit*¹⁰, *Yahoo! Buzz*¹¹, *Delicious*, *Facebook*, and *Twitter* allow users to tag news stories and recommend or share them with friends on their social network. Such news recommendation services not only rate mainstream news, but also help users stumble upon unique news stories that they may not otherwise have a chance to see. Many traditional news sites have plug-in interfaces to these bookmarking sites to enable readers to share news freely. The bookmarking also helps present “popular” or “upcoming” news, as shown by news ranking sites like *NewsPulse*¹², *BuzzFeed*¹³, *Digg*, and *Reddit*.

Analysis of Online News

As illustrated above, many online news sites augment news with videos, images, and user comments to enrich the news consumption experience. Search-based news aggregator sites cluster and categorize news stories and allow users to

¹ Seattle Post-Intelligencer, <http://www.seattlepi.com/>

² The Huffington Post, <http://www.huffingtonpost.com/>

³ Newser, <http://www.newser.com/>

⁴ Seven-Sided Cube, <http://www.sevensidedcube.net/>

⁵ NewsVine, <http://www.newsvine.com/>

⁶ Global Reporter, <http://globalreporter.com/>

⁷ CNN iReport, <http://ireport.cnn.com/>

⁸ StumbleUpon, <http://www.stumbleupon.com/>

⁹ Digg, <http://digg.com/>

¹⁰ Reddit, <http://www.reddit.com/>

¹¹ Yahoo! Buzz, <http://buzz.yahoo.com/>

¹² NewsPulse, <http://newspulse.cnn.com/>

¹³ BuzzFeed, <http://www.buzzfeed.com/>

customize and personalize what they want to read. Services such as email/mobile news alerts and RSS feeds, and customized web pages, such as *My MSN*, *My Yahoo!*, and *iGoogle*, allow users to get news on demand. Readers are encouraged to share news with others, either in their social network or the online community at large. News portals use these to assess popularity of news stories and surface articles that seem to be generating a lot of interest (“buzz”).

The algorithmic aggregation of news across sources used by news aggregators, however, appears to treat all news sources equally, especially when selecting which news item to show. Often, recent updates supersede earlier reports, even if the earlier reports were from reputable sources. On the other hand, users may have other preferences, like local news sources for community news, or in general, specific news sources for preferred genres of news. Local news or news from semi-urban regions are often ill-represented either because of fewer sources reporting on regional news or limited space being allocated to local news on news sites.

The “one size fits all” approach of using keyword query based retrieval is not optimal for news. Query-based triggering is often imperfect, and searching for news using just keyword queries often limits expressivity. News demands a ranking different from web search. News dissemination is more than just selecting a list of news articles about popular events from well-known news sources. Online news must instead cater to specific use-cases and should ideally be personalized to users.

BEYOND SEARCH: NEWS EXPLORATION SCENARIOS

In this section, we propose techniques to select and present news according to user needs and preferences. We present three specific scenarios for a user-driven news digest to illustrate our ideas.

Scenario 1: Catching up on News

Consider the following scenario: Katie is an avid news reader who tracks news on a daily basis, often following up on specific news events several times a day. At times, Katie may be cut off from news, for example, when she goes on a long vacation. When she is back online, she may want to know what happened while she was away. She may want to skim through the major news stories that took place, including updates on the news she was following regularly before going on vacation. This caters to a common, specific need of a news consumer wanting to catch up on news.

Scenario 2: Diaspora Digest

It has become fairly common for people to migrate to another country or city for work or studies. Though most of these expatriates try to keep abreast with the news from the country of origin, they lose touch with traditional sources of news. They may visit news websites from their home country / city periodically to do so, but this may not be easy or convenient. If Katie is from Berlin and residing in the US, she may be interested in a summarized view of key events in Germany from the past week. She might be

interested in German soccer team’s performance round the year and in country-wide soccer competitions such as the German Cup. This caters to the need of expatriates who either live in another country or migrate to a new city within the country and access news about their home city.

Scenario 3: Following Celebrities

A longitudinal look at news is of great value for specific needs, such as following the activities of celebrities. Assume Katie is an admirer of Princess Diana, and wants to get a perspective of Princess Diana’s life history as described in the news. She would be interested in key events such as her marriage, her time as the princess, her divorce, and her death and subsequent investigations. This caters to the user need to get a historic perspective on key people or events in an archival news corpus.

REQUIREMENTS FOR NEWS SYNC

To address the above scenarios, we propose a system called *News Sync*. This system allows Katie and similar news consumers to get adaptive, personalized news digests on a topic, region, time period, or a combination of these.

We list the following requirements for *News Sync*:

1. **Choice of news categories, topics, and sources:** Users should be able to specify the time period of interest. In addition, users may specify if they are interested in news from particular sources, specific news categories, locations/regions, and/or specific topics.
2. **Personalized news feed:** The system should identify stories that are currently the most relevant to the user, based on past user behavior and user preferences, similar in spirit to work by Billsus & Pazzani (2000).
3. **Variety in news content:** The system should show a variety of content across diverse categories, instead of, say, returning a list of ten “most popular” news links which may be restricted to one or two topics. Users can thus get an overall picture of key events first, before they delve into specific stories.
4. **Multimodal and adaptive news presentation:** The news interface needs to be adaptive to the nature of news topic presented, and availability of multiple modes of news content. For example, a search for news about “Harry Potter” over summer 2007 should result in, among other stories, movie trailers (video) of “Harry Potter and the Order of Phoenix”, book reviews (blogs) of “Harry Potter and the Deathly Hallows”, which were both released in July ’07, and news about the Harry Potter theme park announced in May ’07.
5. **Interactive and exploratory user interface:** Users should be able to interactively modify time, location, and other parameters and have the system respond immediately with updated views of relevant news.
6. **Parameterized interface design:** Users should be able to set parameters to get results at different specificities.

7. **Support source-tracing and finding related news:** The system should allow users to go from a news summary to the original news article. Further, the system should suggest other related articles based on the news items viewed.
8. **Ability to share news:** Users should be able to comment on and share interesting news articles over their social network or over the Web via email.
9. **Support news analyses by sentiment and points of view:** Users should be able to view stories summarized by sentiment or different points of view.
10. **Support the familiar “list view” as back-off:** Finally, even as the news interface gets a facelift, it may be prudent to support the list view as a back-off option to take advantage of users’ familiarity with the concept.

Although these requirements have been specified with news domain in mind, they are generally applicable and relevant to other domains such as legal search, patent search, and in intelligence gathering tasks that aim at gleaning information from significant amount of archival text corpora.

THE NEWS SYNC SYSTEM

In this section, we present a brief description of *News Sync*, the system we developed based on requirements we listed in the previous section. We identified three key dimensions for news which users have control over during search (namely, the time period, location, and category of news articles), in addition to search via keywords. The following sections describe system features, assuming control over only these dimensions. It must be noted that it is possible to easily extend the system to other dimensions. Further, while the individual techniques used may not themselves be new, our proposed integration leads to a better news experience.

System Description

Figure 2 gives a schematic diagram of the *News Sync* system. The key steps in the system are:

1. **Deciding on a news corpus:** In this prototype, we use the New York Times corpus, released as part of the HCIR 2010 Challenge¹⁴. It consists of all articles published (or posted online) by New York Times from 1987 to 2007. This is analogous to an archival dump of news articles that can be augmented incrementally, if desired. The corpus also contains fairly rich meta-data annotations like normalized names of people, locations, organizations, and key concepts found in the articles.
2. **Indexing the corpus:** The New York Times corpus was indexed using Lucene.Net¹⁵, with support for field-level queries. This required the removal of stop words

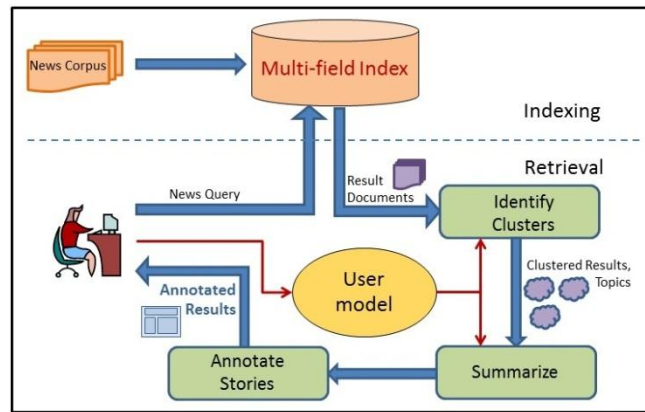


Figure 2. System diagram of *News Sync*.

and additional pre-processing to normalize some fields (such as publication date) to make them searchable.

3. **Retrieving relevant news results:** When a user issues a news query, the system converts it to a suitable Lucene¹⁶ query and retrieves several hundred relevant news results. If a date range is specified, only results from that date range are retrieved. If a category or location is specified, it must appear in all result articles.
4. **Grouping news articles:** News needs to be presented in a manner that is easy to consume. This involves selecting the content to present and deciding how best to present it. In this work, we cluster articles returned by the search system to find related groups of articles. Each group may not be a single story thread, but this clustering-based dimensionality reduction offers a more structured view into the articles. Recursive clustering can help us get to *news stories*, which are collections of strongly related articles. We currently cluster on key concepts from articles, including named entities, descriptors, categories, and section headings obtained from the article meta-data. We use the Hierarchical Agglomerative Clustering algorithm (Hastie et al., 2009) to find clusters and threshold it based on similarity between news articles computed over the key concepts listed above. These news clusters may be created adaptively based on the individual user models built from the user profile, explicit user preferences, and implicit interest tracking.
5. **Summarizing news clusters:** We also adaptively summarize the clusters to provide some insight into the articles within a cluster. Summarization is performed using a modified version of SumBasic (Nenkova & Vanderwende, 2005). SumBasic is an extractive summarization system which iteratively selects a few most significant sentences from one or more articles. The set of sentences extracted do not necessarily form a cogent paragraph. However, just as result snippets

¹⁴ Fourth Workshop on Human-Computer Interaction and Information Retrieval, 2010, <http://www.hcir2010.org/>

¹⁵ Lucene.Net, <http://lucene.apache.org/lucene.net/>

¹⁶ Lucene, <http://lucene.apache.org/java/docs/index.html>

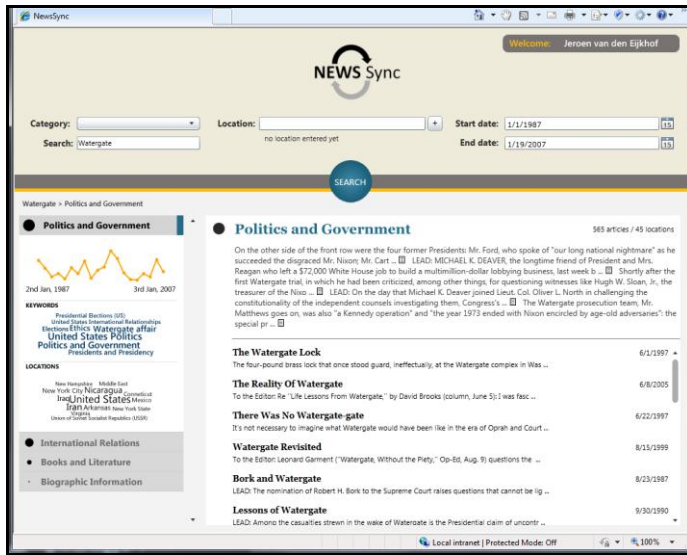


Figure 3. Screenshot of News Sync showing results for the query “Watergate” in the catching-up scenario.

provide some insight into individual results on a search result page, we expect these summaries to provide an indication of what the created clusters are about.

6. **Adding aggregated meta-data about the clusters:** Each news cluster is annotated with additional meta-data such as the news timeline, relevant categories, locations, and key concepts from the articles.
7. **Presenting and visualizing news:** Once the news clusters are annotated, they are presented to the user along with relevant meta-data. The meta-data, presented in the form of sparklines and tag clouds, can be used by the users to further explore and refine the clusters. This is described in detail in the next section on User Interaction. Figure 3 shows a screenshot of the results for a catching-up scenario query, “Watergate”.

The system is developed in C#. The interface is developed using Microsoft Silverlight¹⁷, since it gives us browser independence and access to animation and interactivity.

USER INTERACTION

The system interaction flow is sketched below:

1. **Providing search parameters:** When Katie logs in, she is shown a tag cloud of key topics from the corpus. She can browse for other news by providing one or more of four input parameters – the news category, topics of interest (keywords), location(s), and a date range of interest in the input panel (see Figure 4).
2. **Viewing news clusters:** When Katie enters a news query consisting of one or more parameters, several hundred (currently 1000) relevant results are retrieved

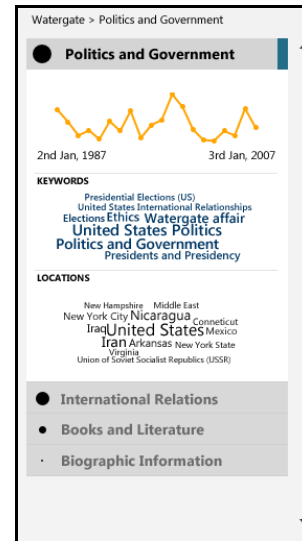


Figure 5. The left panel is used to browse relevant clusters, interact with the sparkline and keyword and location clouds.

from the indexed article store and the articles are dynamically clustered. The left panel of the result screen (see Figure 5) lists the clusters, ordered by popularity and relevance. The top-most cluster is highlighted and the left panel displays additional properties about the cluster, such as tag clouds of key concepts and locations mentioned in the news articles. A sparkline shows the distribution of articles with time.

The right panel (see Figure 6) gives additional information about the highlighted cluster. It shows a brief summary, followed by a list of relevant articles. The list shows the publication date, headline, and the lead paragraph for each article.

3. **Browsing news results:** Katie can either explore the articles in the current cluster or can look into other clusters from the left panel. If she clicks on the article headline, the article is displayed with all relevant meta-data (see Figure 7). If she clicks on another cluster from the left panel, the section with additional information on the first cluster shrinks, and the newly selected cluster expands to show its information. The right panel also shows results for this selected cluster.

Katie can also interact with the sparkline by hovering over or clicking it. The sparkline is divided into 25



Figure 4. Top portion of the interface, showing input panel for entering search parameters.

¹⁷ Microsoft Silverlight, <http://www.silverlight.net/>

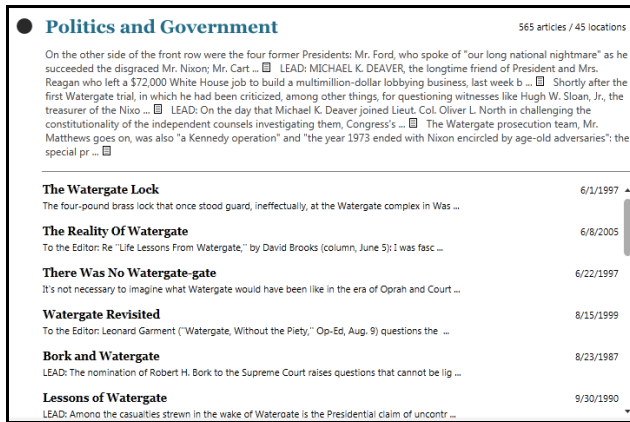


Figure 6. Right panel shows abstract for the selected cluster, followed by a list of articles in the cluster.

time periods, ranging from the date of the first (earliest) story in the cluster through the date of the last (latest) story in the cluster. The time period between the first and last dates is divided into 25 segments, and mapped on to the closest logical time period (viz., day, half-week, week, half-month, month, 3 months, and year). As the hover point is moved across the sparkline, different date ranges can be selected and the number of articles from that date range is displayed. The articles from the selected date range are also highlighted in real-time in the right panel. Clicking on a time segment filters the result set to articles only from this time period, allowing Katie to zoom into news from specific time periods. Typically, this is used to focus on articles corresponding to the peaks or valleys in the sparkline.

Katie can also click on one or more keyword or location tags and reduce the result set to only articles that contain the selected tags. If Katie is interested in exploring a particular cluster in detail, she can select a cluster and choose to dig deeper. A new query is then issued based on the chosen cluster, the news articles in the cluster, and the original query to obtain a refined, second-level clustering that Katie can further explore.

4. **Sharing results:** The interface also allows Katie to share the summary, articles, or stories with her friends on popular social networking sites. She can also save her queries and results for future recall.
5. **Following user actions:** As Katie interacts with the system, her actions, queries, and parameter settings are stored. When Katie reads articles and shares it with her friends, the key concepts from the article are recorded in user models maintained per user. Result ranking, clustering, and summarization of clusters are continuously adapted based on the user model.

Katie can also explicitly restrict her results to be from particular regions or categories. These customization preferences are recorded and subsequent results are tuned to these preferences.

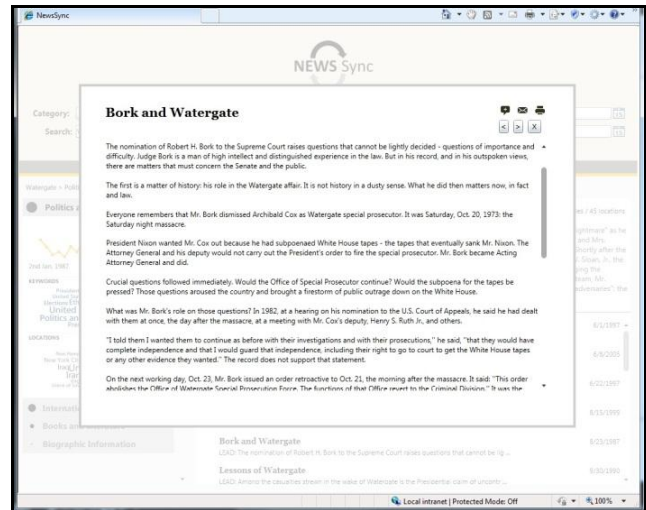


Figure 7. News Sync article obtained after clicking an article in the summary view (Figure 6).

CORPUS CHARACTERISTICS

The system was developed as part of HCIR Challenge 2010, a shared-data initiative run over the summer of 2010, associated with the Human Computer Interaction and Information Retrieval workshop. As part of the challenge, the organizers released a dataset consisting of over 1.8 million articles published by the New York Times from 1987 to 2007. The corpus only had textual content and did not have any multimedia content, including images and videos. The dataset, however, included rich, manually annotated meta-data. All news articles were tagged with the section title and page numbers or URL where they appeared in the print and online versions of New York Times. The named entities in these articles, such as person names, locations, and organizations, were identified and normalized. Further, articles were tagged with other general descriptors that could be used as keywords describing the main themes of the article.

The lack of multimedia content presented some challenges during the design of a news interface. The focus of the interface design was directed away from integrating multimedia content (such as photos and videos) with news, and was refocused towards enhancing user experience by allowing news exploration. The result section was designed to project the navigational aspects of news, allowing users to zoom in and out of news clusters based on time, location, entities, and other key concepts. A version of *News Sync* was demonstrated at HCIR 2010 (Vydiswaran et al., 2010).

EVALUATION

We piloted an initial prototype of *News Sync* internally within our organization. In the first phase, we released it to a small user-base to understand how users interact with the system, using implicit and explicit feedback. Usability studies and a follow-up survey helped us understand popular features and usage patterns.

The next phase was a wider deployment of two competing interfaces with the same look-and-feel and result set, where the exploration features were enabled in one and disabled in the other. This study helped us evaluate the impact of news exploration on user satisfaction and task completion.

Usability Study

In the first evaluation phase, we conducted a boxed usability study. Users were invited for the study of a news site. They were first asked if they had heard of an archival search system and if they read news online. Almost all users displayed some understanding of online news systems, even if they had not used an archival search software before. This helped set the *News Sync* system in perspective. None of the users had seen the prototype before and no user manual was given on how to interact with the system. Users were asked to complete a set of tasks one-by-one, as they explored the system interface and various options that were available. They were encouraged to speak aloud as they navigated the system and their observations and comments were recorded. We specified a mix of easy and difficult tasks covering the different exploration-based news scenarios, as given below:

1. **Clinton Impeachment:** You are discussing the impeachment trial of President Clinton with your friends and want to find arguments in favor or against the impeachment. Find arguments for or against President Clinton's impeachment.
2. **Princess Diana:** You are interested in some details of Princess Diana's death. Find the date and cause of her death. Can you also find some information about the investigation that followed?
3. **Catching up on Sports:** Assume today's date is 8th February 2006, and you were out on vacation since 25th January 2006. Find out what happened in the world of sports during this time.
4. **Diaspora Digest:** If you are from another country (or city, other than your current place of residence), find key events that took place in that country/city in 2005.
5. **Pizza Prices:** You are worried that food prices have skyrocketed over the last twenty years. As an example to support that argument, you want to find how pizza prices have varied from 1987 to 2006 in New York.
6. **Communist Legislator:** You hear that a member of the Communist party once held a legislative or executive post in New York state. Can you identify this person and the post (s)he held?

As is evident above, some tasks were specific and expected users to find exact answers (e.g. date of "Princess Diana's death"), while others were open ended (e.g. "Diaspora Digest" of any specific location – city or country). The first four tasks were more adaptive to exploration, while the last two had specific requirements. The task on "Pizza prices" was an example of a query on topics with long history and

needed analysis of many documents to arrive at an answer, while the task on "Communist legislator" was an example of a query looking to retrieve specific information (such as entity names), similar to a question-answering use case.

The usability study was done with 6 users. Users were asked to complete as many tasks as possible in one hour. The primary observation from the study was that the difficulty of the task significantly impacted user satisfaction. Most users were satisfied with the "Catching up on Sports" and "Clinton impeachment" tasks, as they could find expected answers fairly quickly. On the other hand, users did poorly on "Pizza prices" and "Communist legislator" tasks, because these involved analyzing a lot of documents and users were frustrated with the task. We observed that when the task is loosely defined (tasks such as "Diaspora Digest" and "Catching up on Sports") for a particular time period, they could easily find what they were expecting. Adding UI components such as the sparkline was also helpful in directing users to specific answers quicker (such as date of "Princess Diana's death"). Cluster labels also significantly reduced search time. For example, in the "Clinton impeachment" task, when users queried for reasons using the (obvious) keywords "Clinton impeachment", they got a cluster on "Ethics" that seemed highly relevant.

In general, users found the UI elements easy to work with. After interacting with the system for a couple of minutes, users could easily navigate between list view, article view, and filtered list of results. Users did not find the extractive abstract very useful. The usability study exposed the- need for keyword highlighting and some changes to UI elements (such as avoiding the overlay of article view over result list) to ease browsing. Also, users were unaware of advanced search options (such as using logical operators), and a couple of users requested this feature.

Effect of exploratory interface

In the next set of evaluations, we explored the effects of providing an exploratory interface. Emails were sent to select mailing lists within our institution, and interested participants were directed to one of two interfaces based on their last name. One version of the interface was the same as the one used in the usability study. A second, control interface was created where the results were shown as a list of retrieved documents, similar to what a search engine would provide. Users were able to look at articles by clicking on the result documents, but all other exploratory features were hidden. The two interfaces appeared similar and the users were not told which version they were using – the control version or the experimental version. Users were asked to complete at least two of the six tasks listed above. After completing them using the assigned interface, they could also try the other interface if they wanted. They could also search for any news topic from the twenty year period (1987-2007) for which data was available. A feedback frame was added to both interfaces and users were

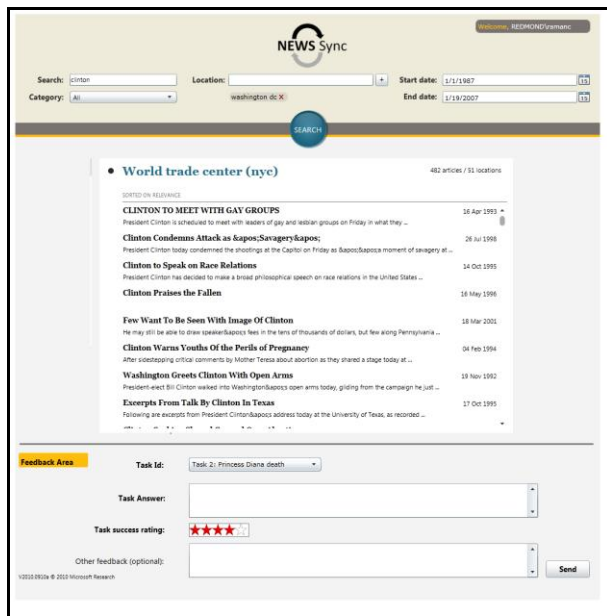


Figure 8. Control interface, with exploration disabled and a feedback frame added.

encouraged to give task-based feedback on different aspects of the system, such as clustering, speed of response, etc. Screenshots of the two interfaces are shown in Figure 8 and Figure 9. Users were asked to rate their satisfaction level using a five-point Likert-type rating scheme (see Table 1). Note that the ratings scale was not defined based on interface usability, but on satisfactory task completion.

In this round of evaluation, 303 users participated, issuing a total of 955 queries, and spending on-an-average 11.37 minutes interacting with the interface. The system received 126 feedback comments – 77 from those using the baseline version and 49 from those using the exploratory version.

Out of seven tasks (six defined tasks and one category for all other tasks/queries that users came up with), the users rated the first three tasks (viz., “Clinton impeachment”, “Princess Diana’s death”, and “Catching up on Sports”) as most successful (see Table 2). The tasks on “Pizza prices” and “Communist legislator” were known to be difficult, and these received poorer scores. We were somewhat surprised that the “Diaspora Digest” task was considered the least successful, because in the controlled usability study, users addressing this task could easily find information they were looking for. On analyzing the ratings and user comments

Rating	Description
1 star	Did not find any answer
2 stars	Found some answers, but missed a lot.
3 stars	Found enough, but could be better.
4 stars	Found most answers.
5 stars	Found all answers I was looking for.

Table 1. Feedback rating scale and its meaning.

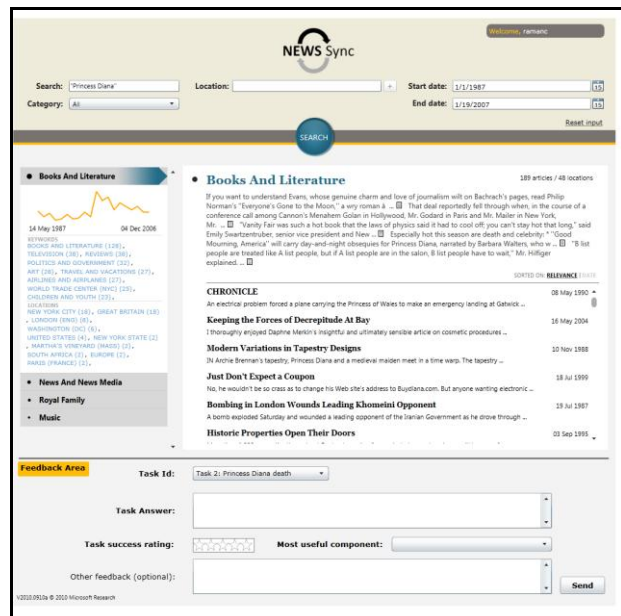


Figure 9. Original interface, with exploration enabled and a feedback frame added.

for this task, we found that users were unclear on what query to input. Some queried for very specific or unpopular places like “Lincroft, NJ”, “Haifa”, or “Guyana” and were not satisfied with the number of articles they found. We conclude that more instructions were needed to explain the “Diaspora Digest” task, and adding additional sources for local news may address the issue of insufficient coverage of location-specific news in the corpus we used in our study.

When users were allowed to search for any news of their interest, they queried for specific categories (e.g. Editorials), favorite sports teams (e.g. “Pittsburgh Steelers”), and some international conflicts (e.g. “Kashmir”) to take advantage of the design feature that news articles were clustered in orthogonal dimensions.

We asked the users to give feedback ratings based on their satisfaction level for the tasks. Users found that different UI components helped them answer different tasks better. When asked to choose the most useful feature, users chose result clustering and keyword based filtering features for the “Clinton impeachment” task and a couple of other tasks, while news trend information encoded in the sparkline was chosen as the most useful feature to answer the task on “Princess Diana’s death”. Table 2 lists the most useful feature users chose for each task. Users did not distinguish between the two interfaces in their overall ratings and both versions got statistically similar scores. Our conjecture is that users gave ratings based on the difficulty of the task and not on the merits of the interface per se. But, allowing users to give general comments about the system helped us capture suggestions and reactions as soon as they finished a task. In general, users believed the system “flow was smooth”, and the interface “well designed” and “NICE!”. Those who worked with both versions found the version

Task	N	Rating	Most useful UI feature
Clinton impeachment	32	3.73	Clustering of results, Filtering with keyword tags
Princess Diana's death	25	3.74	Filtering with time sparkline
Catching up on Sports	16	3.62	Clustering of results
Diaspora Digest	18	2.32	Top input area and search button
Pizza prices	12	2.49	Clustering of results
Communist legislator	10	2.40	Filtering with location tags
Others	13	3.08	Filtering with keyword tags

Table 2. Average rating and most useful features found by users per task. N denotes number of votes.

with exploratory features “Wow! waaaaay (sic) better than [the control version]”. One user commented, “I like all the new touches!! The keyword highlighting is extremely useful in results. I really like the timeline display. Keyword subdivisions are good”. Another user said, “I was impressed that I could actually find pizza price articles. I expected the task to be too difficult”. Some users were impressed with the results even in presence of spelling errors in the query, while others believed that librarians and archive professionals would find this system very useful.

Many users using the control version requested a feature to change the sorting order from relevance to date and to highlight keywords in the text. Both these features were hidden in the control version and were available to the users of the other interface, and were in fact appreciated by users who tried both versions. The users suggested improvements like adding ability to (a) copy and paste text from articles, and (b) take notes. Users wanted the articles to be formatted better to allow finding relevant content easily. Some users wanted more control on search and requested advanced search features like quoted search and logical operands. However, they were unaware that these options were in fact available in both interface versions. Further, some users did not realize that they could look at the peaks in the sparkline and click on them to filter results. Some UI walk-through or highlighting interactive aspects of UI features may have helped alleviate this issue.

RELATED WORK

In the past, researchers have looked into generating a personalized webpage of relevant news based on the user's topics of interest. Kamba et al. (1995) conducted one of the early studies on presenting an interactive newspaper on the Web. They proposed a system that builds web pages dynamically as the user browses the newspaper. Anderson & Horvitz (2002) developed a personalized web page as a montage of links of frequently viewed pages that changes dynamically with the time at which the page is viewed. The system learns which pages are viewed regularly at certain time periods and presents content based on the user's interests and browsing pattern. For example, a user might be shown weather forecasts and key news in the morning; the stock price ticker and work-related articles during the day; and traffic congestion pattern in the evening.

There has also been work in providing personalized newsfeeds. Gabrilovitch et al. (2004) analyzed inter-/intra-document differences and similarities to recognize novel content in articles and how information evolved over time. They developed measures to rank news by novelty, and pick the best (most novel) update to send to the user as a newsfeed. (Columbia's Newsblaster) is a news clustering and summarization system to track and summarize news on a day-to-day basis. Although similar backend technological tools were used in *News Sync*, we studied the problem of enhancing news exploration with better visualization and user interaction. Other researchers, such as Tintarev & Masthoff (2006), studied different measures of similarity of news headlines to improve news recommendation. Bogers & Bosch (2007) studied a few retrieval techniques for news recommendation. Although we used the state-of-the-art, language modeling based retrieval function, the choice of the retrieval function is orthogonal to the work presented in this paper, and alternate retrieval techniques may be used.

A lot of research has also happened in the realm of interface design. For example, Shneiderman (1994) suggests the use of dynamic queries to update search results as users adjust sliders and other UI elements. Teitler et al. (2008) suggest NewsStand, which uses geographic information in news articles to overlay news on a map. This presents users with a geographic perspective of where the news comes from and helps them cluster and explore news based on location. Wang et al. (2006) did some preliminary work on selecting and presenting news in the context of an RSS news reader. Marshall (2007) suggests the use of an RSS-based Times News Reader application to improve the readership of traditional newspapers by presenting news in a personalized and browsable interface, leading to a better news reading experience. Our system, *News Sync*, extends this experience by encouraging news readers to get historic perspective of news in a way that helps exploration of news and, in turn, improve user stickiness and loyalty to the media source.

CONTRIBUTIONS AND CONCLUSIONS

In this paper, we proposed a use-case driven approach to provide a captivating, sticky news consumption experience, using techniques from search, language processing, visualization and learning. We presented key requirements and design of an exploration-based interface, based on three

news exploration scenarios. We described a news exploration prototype called *News Sync* that we developed and deployed internally.

In *News Sync*, we provide controls to users to explore news by specifying topics, a time range, and/or locations of interest. We react immediately to user inputs to show not just the relevant articles, but additional information including clusters and summaries, tag clouds of locations and key concepts, and a sparkline to show temporal trends. We also adapt results based on user preferences and a model of the user acquired over time, to ensure that the user gets maximally relevant content.

Results from a usability study and a small deployment study suggest that users definitely liked visualizing and filtering news on the sparklines and tag clouds, and the ability to directly explore the news archive. Users seem to relate more to the three use-cases we proposed. Task success ratings seem to be related to the quality of the results returned rather than the interface type; for example, the paucity of results about small towns was seen as a task failure than the reality of news coverage by a major newspaper. Carefully constructed examples (with answers certain to be in the archive) may have obviated this problem, but could be testing artificial situations.

Although the domain in this prototype was news, such an exploratory interface would be of significant relevance to other domains such as legal search, patent search, and in intelligence gathering tasks, where there is structured textual information, especially where there is a clear timeline. We propose to try this tool on some other domain. Further, in this work, we have relied on news only from a single source, namely the New York Times. We hope to ultimately extend this to multiple sources, deal with different points of view and sentiments, and work with live news streams and with associated multimedia information.

ACKNOWLEDGMENTS

This work was done at Microsoft Research in 2010, while the first two authors were research interns. The authors thank Sarah Wolfson for her inputs in interface design, and all Microsoft employees who tested our prototype system and gave critical feedback to help with the study.

REFERENCES

Anderson, C. R. & Horvitz, E. (2002). Web Montage: A Dynamic Personalized Start Page. *Proceedings of the 11th International Conference on World Wide Web (WWW)* (pp. 704-712). ACM.

Billsus, D. & Pazzani, M. J. (2000). User Modeling for Adaptive News Access. *User Modeling and User-Adapted Interaction* 10(2-3), 147-180.

Bogers, T. & van den Bosch, A. (2007). Comparing and Evaluating Information Retrieval Algorithms for News Recommendation. *Recommender Systems* (pp. 141-144).

Columbia's Newsblaster <http://newsblaster.cs.columbia.edu>

Gabrilovich, E., Dumais, S., & Horvitz, E. (2004). NewsJunkie: Providing Personalized Newsfeed via Analysis of Information Novelty. *Proceedings of the 13th International Conference on World Wide Web (WWW)* (pp. 482-490). ACM.

Gather (2010). The Changing Face of News Media. <http://www.emarketer.com/> Retrieved May 25th, 2010.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). : Book section 14.3.12 on Hierarchical clustering. *The Elements of Statistical Learning*, 2nd ed. (pp. 520-528). Springer.

Kamba, T., Bharat, K., & Albers, M. C. (1995). The Krakatoa Chronicle – An Interactive, Personalized, Newspaper on the Web. *Proceedings of the 4th International Conference on World Wide Web (WWW)* (pp. 159-170). ACM.

Marshall, C. C. (2007). The Gray Lady Gets a New Dress: A Field Study of the Times News Reader. *Proceedings of the 7th ACM / IEEE-CS Joint Conference on Digital Libraries (JCDL)* (pp. 259-268). ACM.

Nenkova, A. & Vanderwende, L. (2005). The Impact of Frequency on Summarization. *Technical Report, MSR-TR-2005-101*. Microsoft Research.

Nielsen/Net Ratings. MegaPanel Data. Retrieved from <http://www.naa.org/TrendsandNumbers/Newspaper-Websites-Nielsen.aspx>

Pew Research Center for the People and the Press (2010). Americans Spending More Time Following the News – Ideological News Sources: Who Watches and Why. *Pew Research Center Biennial News Consumption Survey*.

Shneiderman, B. (1994). Dynamic Queries for Visual Information Seeking. *IEEE Software* 11(6), 70-77. IEEE Computer Society Press.

Teitler, B. E., Lieberman, M. D., Panozzo, D., Sankaranarayanan, J., Samet, H., & Sperling, J. (2008). NewsStand: A New View on News. *Proceedings of the 16th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS)* (pp. 18.1-18.10). ACM.

Tintarev, N. & Masthoff, J. (2006). Similarity of news recommender systems. *Proceedings of AH'06 Workshop on Recommender Systems and Intelligent User Interfaces*.

Vydiswaran, V.G.V., van den Eijkhof, J., Chandrasekar, R., Paradiso, A., & St. George, J. (2010). News Sync: Three Reasons to Visualize News Better. *Proceedings of the Fourth Workshop on Human-Computer Interaction and Information Retrieval (HCIR)* (pp. 45-48).

Wang, T., Yu, N., Li, Z., & Li, M. (2006). nReader: Reading News Quickly, Deeply, and Vividly. *Proceedings of CHI '06 Extended Abstracts on Human Factors in Computing Systems* (pp. 1385-1390). ACM.