

Creating Temporally Dynamic Web Search Snippets

Krysta M. Svore, Jaime Teevan, Susan T. Dumais
Microsoft Research
Redmond, WA, USA
{ksvore, teevan, sdumais}@microsoft.com

Anagha Kulkarni
Carnegie Mellon University
Pittsburgh, PA, USA
anaghak@cs.cmu.edu

ABSTRACT

Content on the Internet is always changing. We explore the value of biasing search result snippets towards new webpage content. We present results from a user study comparing traditional query-focused snippets with snippets that emphasize new page content for two query types: general and trending. Our results indicate that searchers prefer the inclusion of temporal information for trending queries but not for general queries, and that this is particularly valuable for pages that have not been recently crawled.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information search and retrieval – query formulation, search process.

Keywords

Web search, query dynamics, search result snippets.

1. INTRODUCTION

Web content change often drives people to visit webpages to see what is new. For example, people visit news sites, blogs, and social media sites to stay up-to-date. We believe the search experience can be improved if people can quickly understand how a page has changed. Search result snippets typically provide a query-biased summary of the results. For example, Figure 1(a) shows a snippet for the query *tom bosley* that provides background information about the actor. However, people often use search engines to find results that have changed [1]. Someone who searched for *tom bosley* shortly after he died probably already knew who he was and instead wanted to learn more about his death. In cases like this, the availability of time-sensitive information in the snippet (e.g., Figure 1(b)) may help the user make better relevance decisions.

We explore the effectiveness of including new webpage content in search result snippets. We develop a technique for identifying important new content on a webpage, and generate several types of snippets that take advantage of this information. We conducted a study in which people judged these temporal snippets versus baseline snippets for 465 unique query-URL pairs for two types of queries: general and trending. The results demonstrate that for trending queries searchers prefer temporal content over static content. By leveraging temporal webpage information in snippets, we can better indicate the relevance of the page to the query and enhance the searcher’s experience.

Researchers have recently studied temporal features of snippets. Alonso et al. [2] identified temporal expressions in webpages, and developed snippet generation algorithms that included explicit temporal expressions in sentence selection. Subjective evaluations of the usefulness of temporal expressions were obtained for 10 queries. Campos et al. [3] examined the extent to which queries and search snippets contained explicit temporal expressions. They found that snippets were a rich source of temporal information and could be used to identify whether queries had a temporal intent. Our research differs in that we focus on changes in page

content rather than explicit mentions of time, and we present a more extensive user study with comparisons to a baseline non-temporal snippet.

Researchers have examined how changes in query frequency, document content and user interaction with documents vary over time. Kulkarni et al. [5] developed features to describe changes in query popularity and showed that these features, accompanied by changes in result content, can be a good indicator of a change in query intent. Diaz [4] developed algorithms for identifying queries that are related to breaking news and for blending relevant news results into core search results. We use changes in page content as an indication of the temporal relevance of a sentence when selecting sentences to include in a snippet. We examine how web content dynamics can be used to help users make more informed decisions about the relevance of results.

2. GENERATING TEMPORAL SNIPPETS

The snippets we study are query-focused for a given URL and query. We consider two snippet lengths: long (4 sentences) and short (2 sentences), and two variations of temporal snippets: one that blends new content with the static baseline snippet, and one that only contains new content. To create the snippets for a query-URL pair, we crawl the current version of the page (N for new), and the version of the page that the search engine crawled (C for cached). For each sentence on the webpage we determine its position in the document and the number of total and unique query terms it contains. The *short baseline snippet* is formed by choosing the two earliest sentences from page N that contain the most number of unique query terms. For the *long baseline snippet*, we choose the four earliest sentences with the most number of unique query terms. We do not use the search engine snippets as the baseline since they are of variable length and may be personalized.

For the temporal and blend snippets, we identify the sentences that are the most *changed* from the cached page C . For each sentence on page N , we determine its Dice coefficient with every sentence on page C . The Dice score of the most similar cached sentence is assigned to each sentence on page N , and sentences are ranked by their Dice scores. Ties are broken by the number of

[Tom Bosley: Biography from Answers.com](#)

Tom Bosley won a Tony Award in 1958 for his lead role as New York mayor Fiorello LaGuardia in the Broadway musical *Fiorello!*. But **Tom Bosley** is better remembered for his role on the long-running TV show *Happy Days* (1974-1984).
www.answers.com/topic/tom-bosley

(a) Baseline Snippet

[Tom Bosley: Biography from Answers.com](#)

Bosley died at 4:00 a.m. of heart failure on October 19, 2010, at a hospital near his home in Palm Springs, California. His agent, Sheryl Abrams, said **Bosley** had been battling lung cancer.
www.answers.com/topic/tom-bosley

(b) Temporal Snippet

Figure 1: Snippets for the same URL for the query *tom bosley*. One (a) is query focused; the other (b) includes new content.

Table 1. Snippet preference for general and trending queries. (Maximum value for Temporal v. Baseline is shaded.)

Snippet Comparison		General or Trending Queries					
Snippet Length	Temporal Content	General Queries			Trending Queries		
		Temp.	Base.	None	Temp.	Base.	None
Short	New Only	17%	23%	60%	18%	20%	62%
Short	Blend	17%	18%	65%	23%	16%	61%
Long	Blend	11%	14%	75%	12%	12%	76%

unique query terms, followed by the sentence position. *Short temporal* (“new only”) snippets are formed from the two top-ranked temporal sentences. The *short blend snippet* consists of the top-ranked baseline sentence and the top-ranked temporal sentence. The *long blend snippet* is formed from the two top-ranked baseline sentences and the two top-ranked temporal sentences. For all snippets, the selected sentences are shown in order of their original positions on page *N*. Figure 1 shows the short baseline and temporal snippets for a result for the query *tom bosely*. We do not include long temporal snippets since many pages did not contain 4 temporal sentences.

3. STUDYING TEMPORAL SNIPPETS

Sixty-six people, responding to an email request sent to Microsoft employees, participated in a remote study comparing the different search result snippets. Participants were shown a query and asked to indicate a preference for one of two snippets (a baseline snippet and a temporal snippet) for the same URL. We studied three cases for each query-URL pair: 1) Short baseline vs. Short temporal; 2) Short baseline vs. Short blend; 3) Long baseline vs. Long blend. Participants judged all three snippet combinations for 5 distinct URLs for each query; however, combinations with no differences between temporal and baseline snippets were not judged. Placement of each snippet type was randomized to appear on the top or bottom.

Queries were drawn from two sources: 1) a frequency-weighted sample of *general* Bing search engine queries (292 query-URL pairs), and 2) a sample of *trending* Bing queries that had a spike in frequency the prior day (173 query-URL pairs). One challenge in evaluating the temporal nature of a query is that relevance judgments must occur at the time changes happen. Another challenge is that participants not interested in the query topic may not have the background knowledge to understand what is new or interesting about the query. For this reason, judgments collected in an experimental setting are likely to be biased toward background information versus timely information, as compared with real-world situations. To mitigate this, study participants judged snippets the day after events happened (for trending queries) and selected queries to judge from a list of candidate queries, enabling them to select queries that they knew about.

Table 1 shows the snippet preferences for the general and trending queries. For the general queries, baseline snippets were preferred to temporal snippets, between 4% and 32% more often, depending on the condition. This preference is expected since sentences in the baseline snippet are chosen based on the presence of the query terms, and thus are consistently more query-focused than the temporal snippets. For example, in Figure 1, the full query *tom bosely* appears twice in the baseline snippet, while only *bosely* appears in the temporal snippet. Previous work [6] suggests users prefer snippets with many query terms. The baseline snippets may also perform well since they typically provide good background context for the result. This is evidenced by the fact that preferences are less pronounced when there is additional context provided in the blended temporal snippet. Likewise, long snippets result in fewer preferences (75% of the pairs have no preference),

Table 2. Trending query snippet preference based on age of the cached page. (Max. value for Temporal v. Baseline is shaded.)

Snippet Comparison		Fresh or Old Cached Page					
Snippet Length	Temporal Content	Fresh Cache			Old Cache		
		Temp.	Base.	None	Temp.	Base.	None
Short	New Only	14%	21%	65%	20%	19%	61%
Short	Blend	19%	22%	59%	24%	13%	63%
Long	Blend	10%	14%	76%	13%	10%	77%

perhaps since they enable more background content to be shown. In contrast to the general queries case, for the trending queries, participants preferred snippets with temporal content to the baseline snippets, despite the fact that our judges did not necessarily know that the query had a temporal component. As long as some background context was preserved in the temporal snippets (i.e., the two blended temporal conditions), there was a preference for the temporal snippet. Participants also had fewer preferences in the case of longer snippets. But for the case of the shorter snippets, participants preferred the temporal blended snippet 59% of the time there was a preference.

These results suggest that for trending queries, displaying new content in the snippet can be quite valuable. On manual inspection, we noted that several of the baseline snippets for the trending queries also contained very timely content. In all of these cases, the search engine had recently crawled the page, so the cached version contained very fresh content. To better understand the impact of the time between cache and crawl, we further broke down the data for our trending queries based on the age of the cached page. For 58 of the 173 trending queries, the cached page had been crawled very recently, within one day or less (called *fresh*). For the other 115 queries, the cached webpage was over one day old (called *old*). As can be seen in Table 2, when the cached page was fresh and thus likely similar to the live page, the baseline snippet was preferred over the temporal snippet generated using the difference between the two pages. In contrast, when the cached page was older people were much more likely to prefer the temporal snippet, especially for short blended snippets.

4. CONCLUSION

We performed a user study to determine the impact of including new content in search result snippets for general and trending queries. Our results indicate that users find new content useful for trending queries and when the cached page content is older. A snippet blending static and temporal information also appears beneficial, particularly for trending queries. In future work we hope to study re-finding behavior, which accounts for a significant portion of web search activity. Temporal snippets may be particularly valuable when the user has seen the snippet before.

5. REFERENCES

- [1] Adar, E., Teevan, J. & Dumais, S. T. Resonance on the web: Web dynamics and revisitation patterns. *CHI 2009*, 1381-1390.
- [2] Alonso, O., Baeza-Yates, R. & Gertz, M. Effectiveness of temporal snippets. *WWW 2009: Workshop on Web Search Results Summarization and Presentation*.
- [3] Campos, R., Dias, G. & Jorge, A.M. What is the temporal value of web snippets? *WWW 2011: Temporal Web Analytics Workshop*.
- [4] Diaz, F. Integration of news content into web results. *WSDM 2009*, 182-191.
- [5] Kulkarni, A., Teevan, J., Svore, K. M. & Dumais, S. T. Understanding temporal query dynamics. *WSDM 2011*, 167-176.
- [6] Tombros, A. & Sanderson, M. Advantages of query-biased summaries in information retrieval. *SIGIR 1998*, 2-1.