# A New Visual Search Interface for Web Browsing

Songhua Xu[♮,♯,‡,*]

♮: College of Computer Science and Technology, Zhejiang University, Hangzhou, Zhejiang, 310027, P.R. China

Tao Jin[‡]

♯: Department of Computer Science, Yale University, New Haven, Connecticut, 06520-8285, USA

Francis C.M. Lau[‡]

‡: Department of Computer Science, The University of Hong Kong, Pokfulam Road, Hong Kong, P.R. China

## ABSTRACT

We introduce a new visual search interface for search engines. The interface is a user-friendly and informative graphical front-end for organizing and presenting search results in the form of topic groups. Such a semantics-oriented search result presentation is in contrast with conventional search interfaces which present search results according to the physical structures of the information. Given a user query, our interface first retrieves relevant online materials via a third-party search engine. And then we analyze the semantics of search results to detect latent topics in the result set. Once the topics are detected, we map the search result pages into topic clusters. According to the topic clustering result, we divide the available screen space for our visual interface into multiple topic displaying regions, one for each topic. For each topic's displaying region, we summarize the information contained in the search results under the corresponding topic so that only key messages will be displayed. With this new visual search interface, users are conveyed the key information in the search results expediently. With the key information, users can navigate to the final, desired results with less effort and time than conventional searching. Supplementary materials for this paper are available at **http://www.cs.hku.hk/∼songhua/visualsearch/**.

## Categories and Subject Descriptors

H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*graphical user interfaces, screen design*; H.4.3 [ **Information Systems Applications**]: Communications Applications—*information browsers*; H.3.1 [**Information Storage and Retrieval**]: Content Analysis and Indexing; H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval—*clustering*

*Contact him at A DOT B AT C DOT com in which A = "songhua", B = "xu", and C = "gmail".


*WSDM'09,* February 9-12, 2009, Barcelona, Spain.


## General Terms

Algorithms, Design, Experimentation, Human Factors, Measurement, Performance

## Keywords

Visual search interface, web search, online browsing and navigation, document summarization

## 1. INTRODUCTION

In this information era, we are constantly inundated by massive online information. To help us steer in the information ocean, search engines are indispensable. Indeed, search engines are unarguably one of the most useful tools that computer scientists have offered to the modern society. They have changed our way of accessing information and have reshaped our methods for organizing information. Hence, there have been a great deal of research efforts devoted to search engine design, much of which however focus on the retrieval part of the process.

In this paper, we introduce a novel information searching service which would create a convenient digest with key messages from the raw search results. The digest facilitates further flexible selection by the user. The service is presented via a new visual search interface. (Figures 1 and 2 together show a simple illustrative example). The key idea in our proposal is to group the search results into topic groups at the top of a hierarchy and present selected search results in summary form at the high levels to the user who may then choose to navigate down to the lower levels where the actual webpages or documents can be found. With this new search interface, end users can quickly grasp the topics and contents in the search result set without being exposed to too many details too early. We believe presenting initially a high-level and an overview of the search results would make the web search experience more efficient and effective.

## 2. OVERVIEW

We first describe typical search experiences from using our visual interface in Sec. 2.1. We then discuss how our visual interface functions at a high level in Sec. 2.2. We summarize the contributions of the proposed design in Sec. 2.3.

### 2.1 Search Using Our Visual Interface

A typical web search scenario using our visual interface is as follows. After submitting a query to the search interface, the interface will bring up a summarized view of all

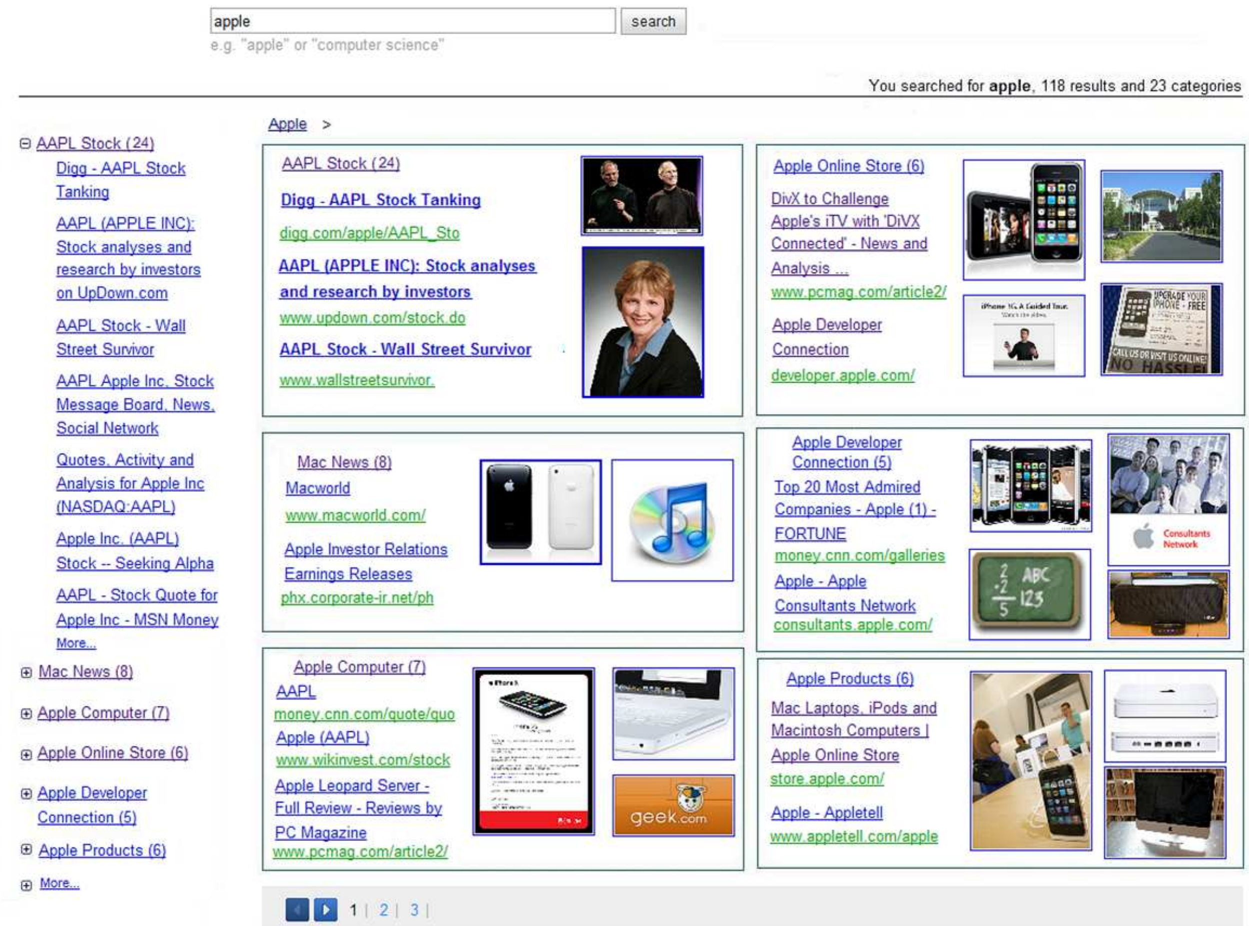


**Figure 1: A search example using our visual search interface. The search keyword used here is "Apple". Here we show the initial search result page. After the user clicks on the category "Apple Stock", he is directed to a new page providing summary information of search results on apple stock (see Figure 2).**

the search results. Unlike traditional search interfaces which present their long "list" of search results by spreading them over many pages, our visual interface organizes the search results in a top-down style which is analogous to the presentation of a satellite map in an electronic map program (e.g., Google Earth) where information is hierarchically organized and progressively presented to the user.

Initially the user sees a high-level overview of all the topics present in the search results. Each topic is mapped to a certain screen region. The size and location of the region depend on the relevance of the topic to the query such that the more relevant or user-interested topics are displayed using bigger regions in the upper part of the window and less relevant ones are displayed in the bottom part of the window. Upon the user clicking a certain region/topic, a zoom-in view comes up, which presents more detailed information on the topic. Depending on how much information there is for the topic, the user will either see in this zoom-in view another overview (of the sub-topics under this topic), or he will see the abstracts of the search results directly. In the current prototype, we allow at most four levels in the hierarchy; that is, a user will click at most three times before reaching the lowest level of the topic hierarchy where abstracts of the corresponding search results will be displayed.

## 2.2 Main Idea of Our Visual Interface

Given a query, we first use a commercial search engine (Google) to fetch the top $N$ relevant webpages, which is user tunable. With this set of search results, we apply semantics-based document clustering to detect topics present in these webpages and to classify all the search result pages into the corresponding topic groups. For each topic group, we create a panel for it in the window of our visual interface. In each panel, we display the automatically detected key words or sentences along with the key images extracted from these webpages if any. After browsing through the panels, a user can then click on the most interested panel, which will bring in zoom-in view of the topic group. The number of documents and/or images displayed in each topic displaying region is dynamically determined according to the size of the region.

The rest of the papers is organized as follows. Sec. 3 discusses how to detect topics latent in the search results via document clustering. Sec. 4 explains how to allocate displaying regions for individual topic clusters. Sec. 5 discusses the issue of selecting the most essential texts and images

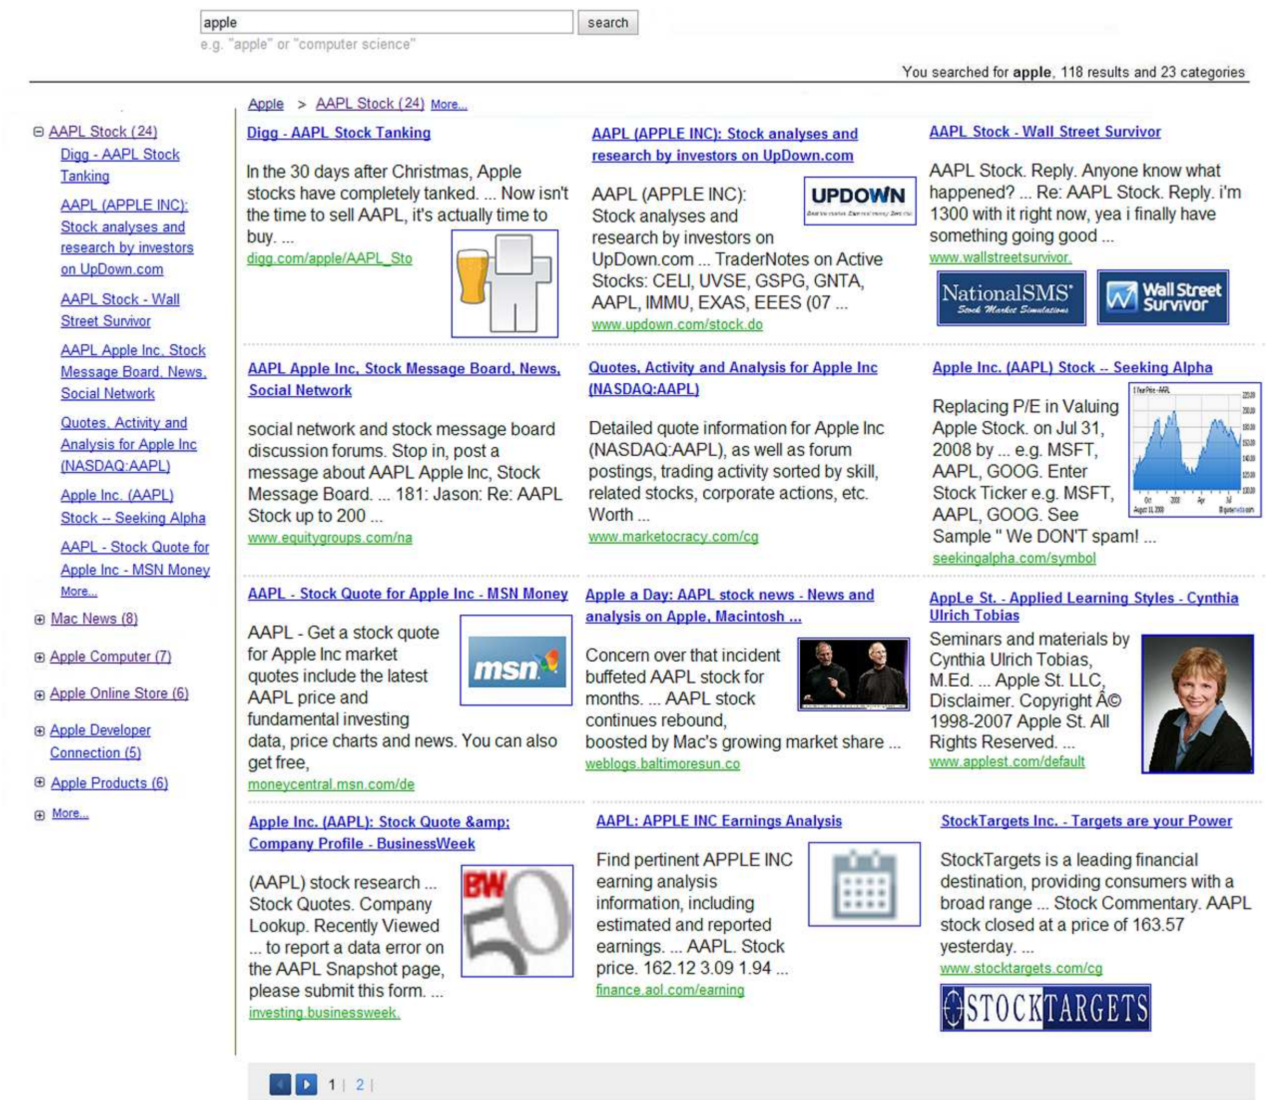


**Figure 2: (Continued from Figure 1) After the user clicks on the category "Apple Stock" in the initial page, he is directed to a new page providing summary information of search results on apple stock.**

from the search results. Sec. 6 presents some examples of searching using our visual interface. Sec. 7 presents a user study we conducted to evaluate the effectiveness of our visual interface for web browsing. Sec. 8 surveys the most related work. Finally, we conclude the paper in Sec. 9 and point out future research directions.

## 2.3 Contributions

- Our visual interface is content or semantics oriented rather than physical link or document oriented. In traditional search interfaces, each document or website is presented as an isolated record in the search result page, regardless of whether it might be similar or related to other documents or websites or not. In contrast, in our new search interface, an item in the displaying region corresponds to a topic or sub-topic. Redundant or nearly redundant results will no longer be repeatedly represented. Also, more important topics will be displayed in higher levels of the search result view using larger space and with more details, while less important information will be deferred to the lower levels of the search result view using less screen space and with less details. Such a human-centric way of search result presentation should facilitate the user's navigation through the search results.

- In addition to the individual search result documents or images, the whole search result set is organized by topics in a visually efficient manner. The relationship between different search result documents and images are clearly revealed in this presentation, which significantly reduces the user's processing overhead in getting to the desirable contents and topics in the search result set. Such method of visual search result browsing can more effectively convey the most essential information in the search results to the user than conventional search interfaces with text-based outputs.

# 3. SEARCH RESULT CLUSTERING

## 3.1 Semantically Meaningful Document Clustering

When the initial search result set is obtained from the search engine, we detect the topics that are latent in the result set and then classify the search results into clusters by their topics. In our design, we carry out document clustering based on nonnegative matrix factorization [26]. Nonnegative matrix factorization [17, 18] is a new matrix analysis technique which has led to a number of successful applications in text mining [24]. Document clustering algorithms based on nonnegative matrix factorization have superior ability in identifying topics present in a document corpus, and document clusters formed according to nonnegative matrix factorization correspond very well to semantically meaningful topics in the document content space.

## 3.2 Hierarchical Document Clustering

The original document clustering algorithm in [26], however, cannot automatically determine the number of clusters in the document corpus, which needs to be specified by the user apriori. Such a manual specification is not desirable in our application as it would mean too much user intervention, and not practical at all since a user cannot tell how many topics there are in the search results before seeing the entire result set. Also, as mentioned earlier in Sec. 2.1, each topic of the search result set will be displayed using a separate region. Thus there cannot be too many topics to be exposed to the user initially. Therefore we go for a hierarchical topic presentation for our application. We adopt the open source hierarchical document clustering software package at http://demo.carrot2.org in our current implementation to produce the number of clusters on each level of the topic hierarchy.

# 4. DISPLAYING REGION ALLOCATION FOR TOPIC CLUSTERS

Our heuristic in allocating displaying regions for topic clusters is that the more important a topic cluster is, the larger the space in which it shall be displayed. To measure the importance of a topic cluster in the search result set, we utilize the webpage ranks provided by Google. Basically, for each topic cluster of search results, we compute the average webpage rank held by all the search results in the topic cluster in Google's webpage ranking. The smaller the average rank, the more important the topic cluster is since it is judged by Google as being more relevant to the query. Assuming there are $n$ topic clusters in the search result set, $\mathbf{dc}_1, \cdots, \mathbf{dc}_n$ with their average webpage ranks in Google being $r_1, \cdots, r_n$ respectively. Assuming the whole screen space occupied by our visual search interface is $S$, then ideally the document topic cluster $\mathbf{dc}_i$ shall be allocated a region of size $s_i$ where:

$$s_i \triangleq \frac{\frac{1}{r_i}}{\sum_{i=l}^{n} \frac{1}{r_l}} S. \tag{1}$$

For simplicity and visual elegance, we allocate a rectangular region to a document topic cluster. Our layout allocation problem is reduced to a rectangle-in-rectangle packing problem. To determine the layout of all the topic displaying regions, we adopt the automatic yellow-page pagination and layout algorithm proposed by Johari et al. [12].

# 5. SELECTIVE PRESENTATION OF SEARCH RESULTS

Now we have determined the number of topics suitable for presentation at each level of the user view. The next question is how to select the most useful search results and automatically create a summarized version of them to present to the users. At the moment, we only consider presenting selected texts or images embedded in the search result documents. The details follow.

We offer three modes in the system: text only, image only, and text and image combined. In the first two modes, only text or image information will be extracted for presentation. In the last mode, both of them will be extracted. Our decision to select both images and texts as highlights of a document to be presented is supported by prior study conducted by Woodruff et al. [25], which shows that using mixed text summaries and thumbnails achieves a better performance than using either text summaries or thumbnail images in informing users of search results. The amount of information extracted for each mode is processed as follows:

**Text only.** We use the iterative graph-based extractive summarization algorithm proposed by Mihalcea and Tarau [20] to extract key sentences. Given the screen area $size(d_{x_j})$ allocated for displaying information in a representative search result document $d_{x_j}$, we use the algorithm to generate a few key sentences so that when these sentences are displayed, the screen space they occupy is closest to $size(d_{x_j})$. If the displaying space is larger than $size(d_{x_j})$, we will trim off the ending words in the extracted key sentences.

**Image only.** Given a representative search result document, if it contains images, we first differentiate those content images from advertisements or website navigation images. Sophisticated image classification techniques are available for this task. In our current system implementation, we use a simple heuristic: if an image is larger than $200 \times 200$ pixels and is located in the center of the document, i.e., it is not at the corner of the webpage, and it is not floating around (which is typical of many advertisements), we consider it an actual content image of the document. If there are more than a certain number of content images detected in a document, we consider the document as being image intensive. Right now we set the threshold to five images. For a document that contains less than five images, we will pick the first image for presentation. For a document that contains more than five images, we will select the first $a \times b$ content images for presentation. In the latter case, given the screen space for displaying images in the document, we constrain each image selected to be at least 50 pixel wide and long. This gives the maximum number of rows and columns of images, i.e., $a$ and $b$. If there are less than $a \times b$ content images in the document, we will find a maximum number $a' \times b'$ which is smaller than or equal to the total number of content images in the document. Once

Table 1: Performance statistics on web browsing using our visual search interface

| | Number of | | | Time | | | | |
|---|---|---|---|---|---|---|---|---|
| Keyword | Webpages | Images | Clusters | Clustering | Text Selection | Image Selection | Layout | Total |
| apple | 98 | 882 | 20 | 4.47s/62.6% | 1.34s/18.76% | 1.21s/16.95% | 1.12s/15.69% | 7.14s |
| computer science | 100 | 921 | 21 | 3.34s/48.06% | 1.27s/18.27% | 1.33s/19.14% | 1.01s/14.53% | 6.95s |
| greenhouse effect | 96 | 768 | 20 | 2.21s/39.39% | 1.22s/21.74% | 1.17s/20.86% | 1.01s/18.00% | 5.61s |
| Android | 98 | 834 | 23 | 4.32s/55.67% | 1.23s/15.85% | 1.25s/16.11% | 1.05s/13.53% | 7.76s |
| iphone | 94 | 891 | 18 | 4.15s/51.68% | 1.30s/16.19% | 1.39s /17.31% | 1.19s/14.82% | 8.03s |
| financial crisis | 96 | 698 | 22 | 3.02s/46.68% | 1.20s/18.55% | 1.22s/18.86% | 1.03s/15.92% | 6.47s |
| olympics | 99 | 812 | 22 | 4.02s/53.96% | 1.19s/15.97% | 1.15s/15.44% | 1.09s/14.63% | 7.45s |
| Hong Kong | 98 | 884 | 18 | 3.08s/46.67% | 1.21s/18.33% | 1.24s/18.79% | 1.07s/16.21% | 6.6s |
| Crater Lake | 97 | 890 | 18 | 4.07s/53.27% | 1.25s/16.36% | 1.19s/15.58% | 1.13s/14.79% | 7.64s |
| Halloween | 98 | 859 | 23 | 3.08s/48.12% | 1.13s/17.66% | 1.15s/17.97% | 1.04s/16.25% | 6.4s |

these images are selected, we will tile them into a big rectangle for display using an open-source image mosaic package (http://jimage-mosaic.sourceforge.net/).

**Mixed text and image.** In this mode, both text and image information will be selectively presented to the end user. Right now we restrict the right part of the document's displaying window to be allocated to the presentation of image(s) and the left part to texts. The proportion of the image displaying versus text displaying region is derived from the proportion of the space occupied by images versus that occupied by texts in the original document.

## 6. WEB SEARCH USING OUR VISUAL INTERFACE

Figures 1 and 2 together show an example of searching using our prototype visual search interface. The user submitted the search keyword "apple". He is then presented with an overview of all the topics contained in the search results, which is given in a tree structure on the left margin of Figure 1. The number of search result articles for a certain topic or sub-topic is listed inside the brackets. In the main screen, six most important topic categories are identified by our algorithm and each is allocated a certain screen displaying region having a light green boundary. Inside each region, the topic title, the titles of a few most representative articles in the topic category and their URLs as well as a few most revealing thumbnail images are presented. The user can click on the entry titled "more ..." to request for summary information of more topic categories. He can also click on an article title or URL to open the article. He can also click on the topic category title, either in the topic tree or at the top of the topic displaying region, to enter the next level of the topic. In this search example, after a quick browsing, the user wants to know more about apple stock; so he clicks the entry entitled "AAPL Stock (24)". This takes him to the next level where more summary information about articles on apple stocks is provided (Figure 2).

For this search example, there is no deeper level of subtopic under the apple stock topic category. So all the entries will be the search result articles themselves. Due to screen space limit, only summary information for 12 out of all the 24 search results is included. Again, the user can click on an article title or URL to open the article. And if he clicks on the entry entitled "more ..." either in the topic tree or in the summary information displaying area, summary information of more search result articles will be displayed.

With the help of the tree navigator on the left side of the screen, the user can freely and easily switch to a new topic or sub-topic. Figure 3 shows another search example of using our visual search interface where the search keyword is "computer science". Due to space limit, only the initial search result page is shown.

Lastly, we show some statistics on the performance of web browsing using our visual search interface in Table 1. Ten query examples are shown here. From these measurement data we can see that users can get the magazine style summarized presentation of their search results using our visual search interface within 8 seconds generally. The experiments were carried out on a desktop computer equipped with a 2.83 GHz Intel(R) Core(TM)2 Duo CPU and 2.0 GB memory. The operation system running on the computer is Microsoft Windows XP Professional. In real deployment, much of the processing functions can be performed on the server side, and obviously a distributed implementation using multiple servers would shorten the times to be commensurate with that of using any other ordinary search engine.

## 7. A USER STUDY

To further evaluate the effectiveness of our approach, we invited 30 graduate students to participate in a user-study to examine the benefit of using our new visual search interface in a controlled setting. We choose Google as the benchmark since it is the most popular and commercially successful one. In the current study, we use 30 explorative questions, which are listed at Table 2. In compiling these questions, we select those that require explorative web search to draw the answers since our new visual search interface is designed to facilitate web browsing and navigation.

In our user study, each participant is asked to find answers from the Internet for 15 randomly chosen questions out of the 30 explorative questions using our newly proposed visual search interface and the remaining 15 questions using Google. Note that if a user is randomly chosen to answer a question using Google search, he would not be asked to answer the same question again using our search interface; and vice verse. In our current study, a user is always first asked to answer 15 random questions using our visual search interface and then to answer the rest 15 questions using Google. We do not implement more sophisticated counterbalancing strategy regarding the order of presentation for the two search interfaces. This is because these two interfaces differ in many aspects, in particular our interface newly exhibits the concept of topic clusters, provides summaries for each topic cluster in both the image and text forms, and has a different layout arrangement than that of Google. In addition, all the 30 subjects we recruited for the user study are familiar with Google and have used Google a number of times in the past for web search. Hence we assume conducting the user study first using our interface and then using Google would not noticeably change the user behavior for them to search web using Google. This user study however does carry a little bias against our new visual interface because the users are not as familiar with our interface as with Google. Throughout the user study, when a participant is using our interface for web search, our interface always functions in the mixed image and text mode, i.e. both the text and image abstracts of a topic, or sub topic or a search result article are provided. In our study, we also do not constrain the number of queries each participant shall issue to answer each equation, either using our interface or Google. Instead, it is solely left to the participant's own judgement when they shall stop querying to acquire answers for a particular question during the user study. However, from the time and the number of mouse clicks it took for the subject to answer a question, which we do record during the user study, a rough estimation on the efforts the subject spent on addressing the question is possible.

We record the time each user spends on answering each question (completion time), the number of mouse clicks to answer each question (no. of clicks), how satisfied the participant feels about the search engine or interface in assisting him to answer a question (satisfaction feedback score). The score scale is from 0 to 10 where a larger number means more satisfaction. At the end of the experiment, we also manually evaluate the quality of each user answer (answer quality). In this evaluation, we first look thoroughly for answers to these questions through Internet search and then come up with a grading guideline specifying what types of answers will be given what ranges of score. And then we read all the user answers one by one. After that, we first decide which range of the quality score should be given according to the grading guideline. Within that range, we subjectively determine which score should be given. The score range is from 0 to 1, and the larger a quality score, the better the quality of the answer.

We compare the user performance data in answering the 30 questions using Google and our search interface respectively, as shown in Figures 4–7. The comparison is based on the above four aspects, i.e., completion time, no. of clicks, satisfaction feedback score and answer quality. In each figure, we compare the mean of the two types of performance data for each question. From these user study data, we conclude using our new visual search interface, users can answer explorative questions with less time, less mouse clicking, and can achieve higher question answering quality; and the users in the study are generally more satisfied with using our interface than Google in searching for the answers to these questions. Given that there are many differences between our new search interface and that of Google, it is difficult to pin down to which aspect of the new interface design the improved user search experience shall be attributed. However, all the design features of our new interface are introduced to allow users to navigate through the search result set in a topic organized fashion meanwhile being visually informed of the highlighted texts and images in the topic. Thus it is safe to assume that the topic oriented search result presentation along with visual topic summary can indeed facilitate users to browse the web and conduct explorative search.

# 8. RELATED WORK

## 8.1 Cluster-Based Document Collection Browsing

One of the closest related work to our study here is the scatter/gatther tool developed by Xerox, which offers a cluster-based document browsing method. Their system divides a document corpus into groups, and allows users to iteratively examine the resultant document groups or sub-groups for content navigation. Their early work had proposed using scatter/gatther as a tool for browsing large to very large document collections [9, 8], wherein the focus was placed on the efficiency of clustering: a linear time algorithm [9] and a constant interaction-time algorithm [8] were proposed respectively for clustering documents in the scatter process. They demonstrated that scatter/gatter can be employed to browse the topic structure of a very large text collection [22]. This becomes a new way to examine search results. Compared with their method, our work concentrates on the design of a new visual search interface that facilitates intuitive examination of search results in a topic-oriented fashion. Thus much of our effort is dedicated to achieving a most effective presentation of search results using the limited screen space and to support efficient navigation and location of user-desired search results.

## 8.2 Visual Search Interfaces and Visual Data Exploration

One of the earliest web visualization systems was proposed by McCrickard and Kehoe [19] nearly a decade ago for interactively visualizing web search results. Beale et al. [3] used 3D spatial structures to visualize the relationship between a query and its search results for helping users to navigate through search results. However, their visualization per se does not reveal much information in the search results. Whereas in our system design, we bring forward the exposure of key information in the search results to the search result browsing step. Robertson [23] introduced the data mountain metaphor to leverage users' spatial memory for document management. Amento et al. [1] introduced a system called "TopicShop" which can build site profiles for relevant websites. An empirical evaluation of their method was also reported. Compared with their system, our system is oriented towards individual search results, not websites as targeted in their system.

Benford et al. [5] surveyed the work finished by the year 2000 on visualizing the World Wide Web using 3D-graphics, which include visualization efforts on web structure, browsing history, searches, evolution of the web, as well as the dynamic behaviors of multiple users. Au et al. [2] introduced a system to visually display the distribution of search results by grouping these results into topics according to their keywords. But their system does not allow progressively examining information in the search results as explored and supported in our system. In comparison, the principal advantage of our system is that users can simultaneously be informed of the key information contained in the search results while exploring the distribution of topics in the result set.

**Q1** What are the most well sold Apple computer models?

**Q2** What are the most representative Hello Kitty (a Japanese cartoon character) images?

**Q3** What is the song by which Sarah Brightman became famous?

**Q4** What are the highlights in 2008 Beijing Olympics opening ceremony?

**Q5** What are the golden meddles that Jingjing Guo (a Chinese athlete) won in history?

**Q6** What is the operation system used for the MIPS machines?

**Q7** Find five entry-level piano practising notes.

**Q8** What are the popular team management textbooks?

**Q9** What is the anecdote associated with Lan Tingxu (a famous Chinese calligraphy artwork)?

**Q10** How did Xizhi Wang (a famous ancient artist) become famous?

**Q11** What are the strategies to keep you focused on your work?

**Q12** What is this summer's traveling recommendation in East Asia?

**Q13** What are the popular things people do in their holidays this year at Hong Kong?

**Q14** What is the hotel recommendation on the island Bali?

**Q15** What are the most popularly recommended stocks at HK stock market this week?

**Q16** How many times has Bill Gates spoken in China?

**Q17** What is special with water cube's construction for Beijing Olympics Games?

**Q18** What are the most urgent environment issues?

**Q19** What is greenhouse effect?

**Q20** What are the popular movie reviews with the film Ice Age?

**Q21** What are the recent photos with wide tigers in China's countryside?

**Q22** What are the popular movie reviews with the film Fongfu Panda?

**Q23** How to write an abstract?

**Q24** What are the popular textbooks in machine learning?

**Q25** Find a few photos about the firework in an olympics' opening ceremony.

**Q26** What are popular textbooks for oral English practice?

**Q27** What are the popular java tutorial materials?

**Q28** What is the effect of chocolate on a person's emotion?

**Q29** What is the rule for the game "the lord of the ring" (a game adapted from the film The Lord of the Ring)?

**Q30** What are the recent prices for the popular Canon digital camera models?

**Table 2: The 30 questions adopted in our user study.**

Becks et al. [4] proposed a text-access interface based on a document map metaphor, which offers a graphical overview of document knowledge for knowledge management. Cockburn and McKenzie [7] evaluated the effectiveness of spatial memory in 2D and 3D physical and virtual environments. During their evaluation, they focused on the added value of using the third dimension in the virtual document space for spatial memory. Lam and Baudisch [16] introduced a technique called "summary thumbnails" for browsing webpages on small screen devices like PDAs. In their work, texts from the original webpages are reduced to fit the small screen while the original page layout is preserved. Also related is the study conducted by Kaasten et al. [13] on the minimum sizes of thumbnail images and lengths of title and URL texts for people to recognize previously visited webpages. Since our work is mostly concerned with effectively utilizing the limited screen space to inform the users of search results, during our visual interface design, we take into account their recommendation on the minimum sizes of thumbnails, titles and URLs for easy user recognition of webpages and websites.

## 9. CONCLUSION

We have introduced a visual search interface to facilitate efficient navigation of search results. Using this new interface for web searching, a user can quickly grasp the latent topics in the search results as well as quickly locate the documents he is interested in. This new interface presents search results to the end users in a semantics oriented rather than physical structure oriented fashion as is done in current search engine interfaces. We believe this new way of search result presentation can improve user experiences when searching the Internet.

In the near future, we plan to conduct a comprehensive user study to quantitatively and objectively measure the effectiveness of our new search interface in contrast with the traditional text-based web search interfaces and other existing and emerging commercial visual search interfaces. We also plan to strengthen and enhance the algorithmic components currently used in our visual search interface, including tuning the algorithms we adopted from existing literature and publicly available open-source software packages we have made use of in the implementation.

## 10. ACKNOWLEDGEMENTS

We thank Hao Jiang and Wenxia Yang for helping us in some of the experiments. The first author would like to thank David Gelernter for many inspiring discussions on intelligent web, from which the idea of this paper stems. This work has a patent pending.

## 11. REFERENCES

[1] B. Amento, W. Hill, L. Terveen, D. Hix, and P. Ju. An empirical evaluation of user interfaces for topic management of web sites. In *CHI '99: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 552–559, New York, NY, USA, 1999. ACM.

[2] P. Au, M. Carey, S. Sewraz, Y. Guo, and S. M. Ruger. New paradigms in information visualization. In *SIGIR '00: Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and*

*Development in Information Retrieval*, pages 307–309, New York, NY, USA, 2000. ACM.

[3] R. Beale, R. McNab, and I. Witten. Visualising sequences of queries: a new tool for information retrieval. *Proceedings of International Conference on Information Visualisation*, pages 57–62, 1997.

[4] A. Becks, C. Seeling, and R. Minkenberg. Benefits of document maps for text access in knowledge management: A comparative study. In *SAC '02: Proceedings of ACM symposium on Applied computing*, pages 621–626, New York, NY, USA, 2002. ACM.

[5] S. Benford, I. Taylor, D. Brailsford, B. Koleva, M. Craven, M. Fraser, G. Reynard, and C. Greenhalgh. Three dimensional visualization of the world wide web. *ACM Computing Surveys*, 31:1–16, 1999.

[6] K. Borner, Y. Feng, and T. McMahon. Collaborative visual interfaces to digital libraries. In *JCDL '02: Proceedings of the 2nd ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 279–280, New York, NY, USA, 2002. ACM.

[7] A. Cockburn and B. McKenzie. Evaluating the effectiveness of spatial memory in 2d and 3d physical and virtual environments. In *CHI '02: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 203–210, New York, NY, USA, 2002. ACM.

[8] D. R. Cutting, D. R. Karger, and J. O. Pedersen. Constant interaction-time scatter/gather browsing of very large document collections. In *SIGIR '93: Proceedings of the 16th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 126–134, New York, NY, USA, 1993. ACM.

[9] D. R. Cutting, D. R. Karger, J. O. Pedersen, and J. W. Tukey. Scatter/gather: a cluster-based approach to browsing large document collections. In *SIGIR '92: Proceedings of the 15th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 318–329, New York, NY, USA, 1992. ACM.

[10] M. A. Hearst, D. R. Karger, and J. O. Pedersen. Scatter/gather as a tool for the navigation of retrieval results. In *Proceedings of AAAI Fall Symposium on AI Applications in Knowledge Navigation and Retrieval*, 1995.

[11] M. A. Hearst and J. O. Pedersen. Reexamining the cluster hypothesis: scatter/gather on retrieval results. In *SIGIR '96: Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 76–84, New York, NY, USA, 1996. ACM.

[12] R. Johari, J. Marks, A. Partovi, and S. Shieber. Automatic yellow-pages pagination and layout. *Journal of Heuristics*, 2(4):321–342, 1997.

[13] S. Kaasten, S. Greenberg, and C. Edwards. How people recognize previously seen web pages from titles, urls and thumbnails. In *People and Computers XVI (Proceedings of Human Computer Interaction 2002)*, pages 247–265, 2002.

[14] D. A. Keim. Information visualization and visual data mining. *IEEE Transactions on Visualization and Computer Graphics*, 08(1):1–8, 2002.

[15] W. Kules, M. Wilson, M. Schraefel, and B. Shneiderman. From keyword search to exploration: How result visualization aids discovery on the web. Technical Report 1516920080208, School of Electronics and Computer Science, University of Southampton, 2008.

[16] H. Lam and P. Baudisch. Summary thumbnails: readable overviews for small screen web browsers. In *CHI '05: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 681–690, New York, NY, USA, 2005. ACM.

[17] D. D. Lee and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

[18] D. D. Lee and H. S. Seung. Algorithms for non-negative matrix factorization. In *NIPS '00: Proceedings of Advances in Neural Information Processing Systems*, pages 556–562, 2000.

[19] S. McCrickard and C. Kehoe. Visualizing search results using sqwid. In *Proceedings of the 6th International World Wide Web Conference*, 1997.

[20] R. Mihalcea and P. Tarau. A language independent algorithm for single and multiple document summarization. In *IJCNLP '05: Proceedings of the International Joint Conference on Natural Language Processing*, pages 19–24, October 2005.

[21] D. A. Nation, C. Plaisant, G. Marchionini, and A. Komlodi. Visualizing web sites using a hierarchical table of contents browser: Webtoc. In *Proceedings of the 3rd Conference on Human Factors and the Web*, pages 67–75, 1997.

[22] P. Pirolli, P. Schank, M. Hearst, and C. Diehl. Scatter/gather browsing communicates the topic structure of a very large text collection. In *CHI '96: Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 213–220, New York, NY, USA, 1996. ACM.

[23] G. Robertson, M. Czerwinski, K. Larson, D. C. Robbins, D. Thiel, and M. van Dantzich. Data mountain: using spatial memory for document management. In *UIST '98: Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology*, pages 153–162, New York, NY, USA, 1998. ACM.

[24] F. Shahnaz, M. W. Berry, V. P. Pauca, and R. J. Plemmons. Document clustering using nonnegative matrix factorization. *Information Processing and Management: An International Journal*, 42(2):373–386, 2006.

[25] A. Woodruff, R. Rosenholtz, J. B. Morrison, A. Faulring, and P. Pirolli. A comparison of the use of text summaries, plain thumbnails, and enhanced thumbnails for web search tasks. *Journal of the American Society for Information Science and Technology*, 53(2):172–185, 2002.

[26] W. Xu, X. Liu, and Y. Gong. Document clustering based on non-negative matrix factorization. In *SIGIR '03: Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Informaion Retrieval*, pages 267–273, New York, NY, USA, 2003. ACM.

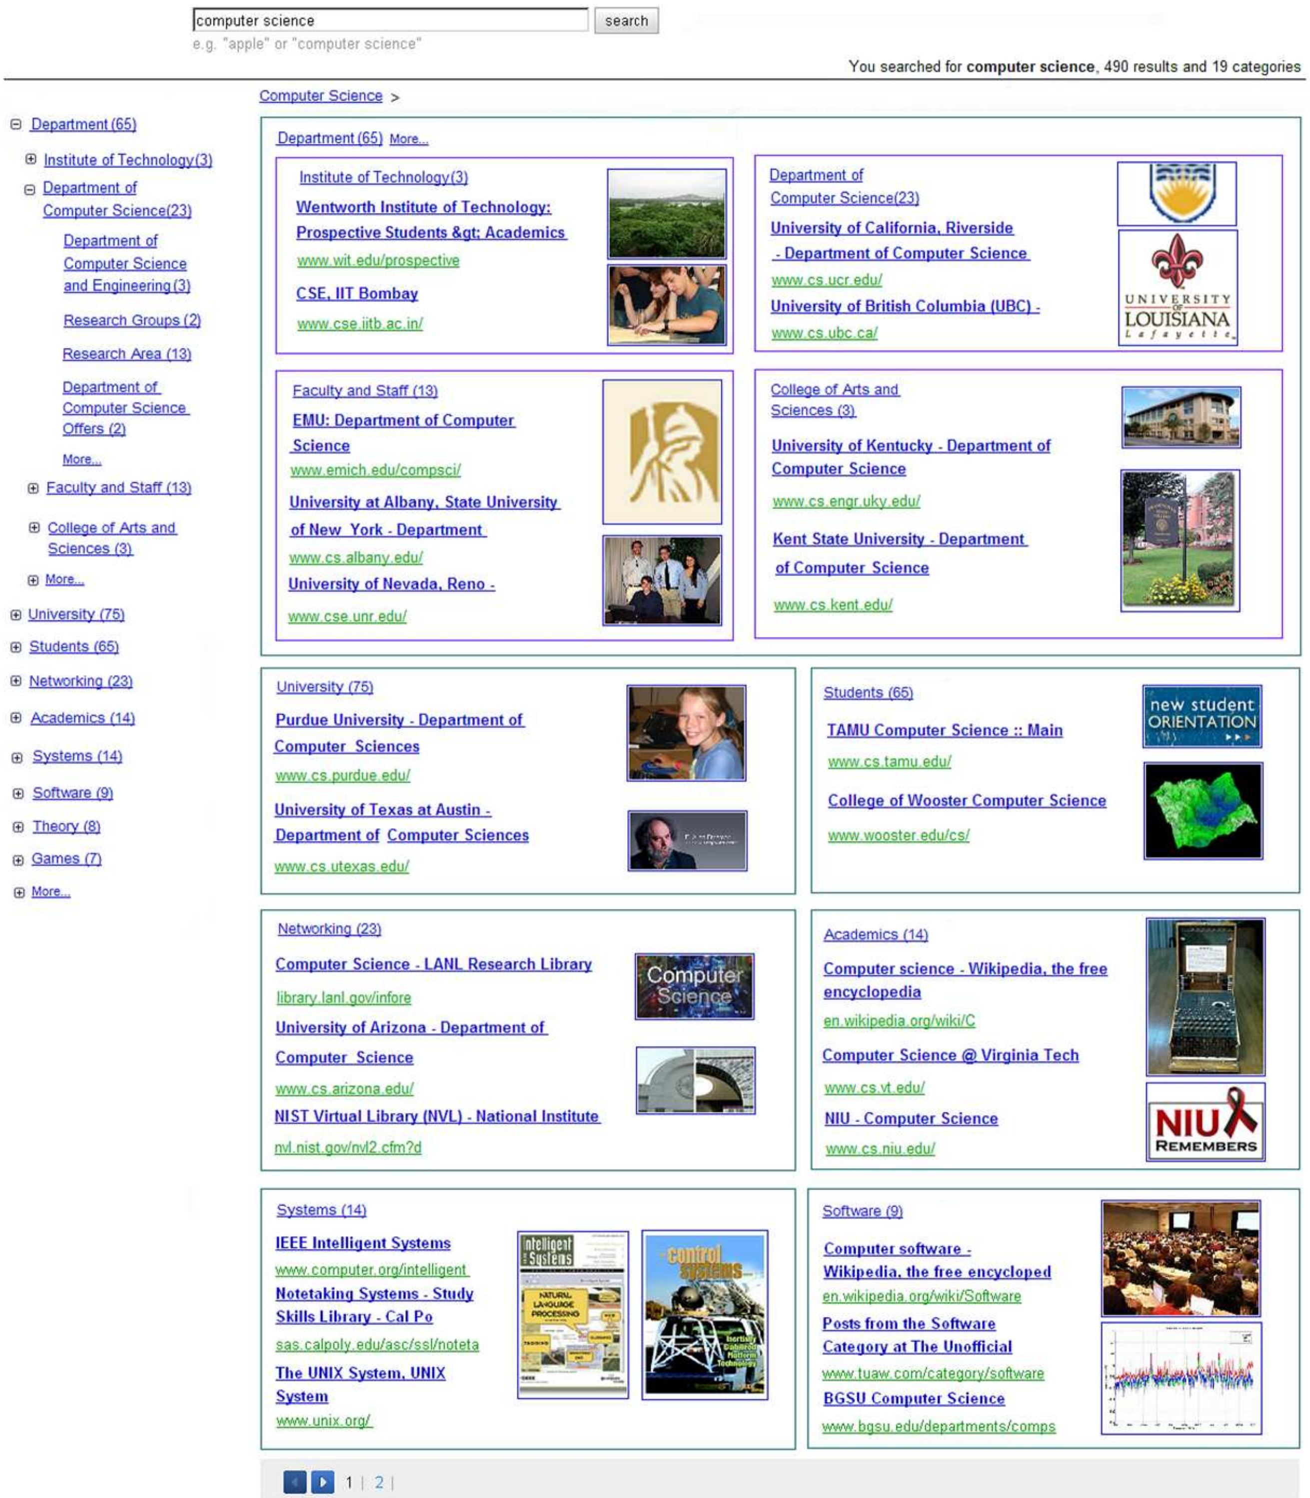


**Figure 3: A search example using our visual search interface. The search keyword used here is "Computer Science". Here we show the initial search result page.**

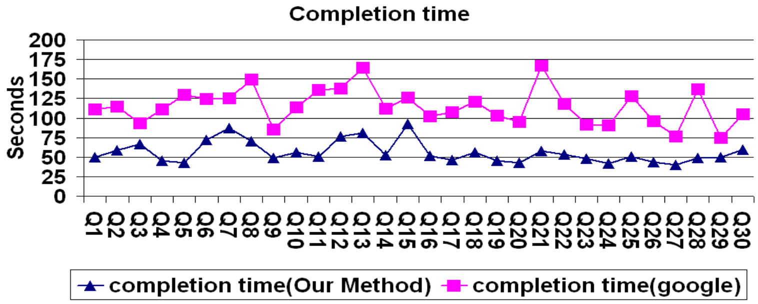


**Figure 4: Comparison of user completion time for answering the 30 questions using Google versus our search interface respectively. Here we compare the mean of completion time using the two methods.**

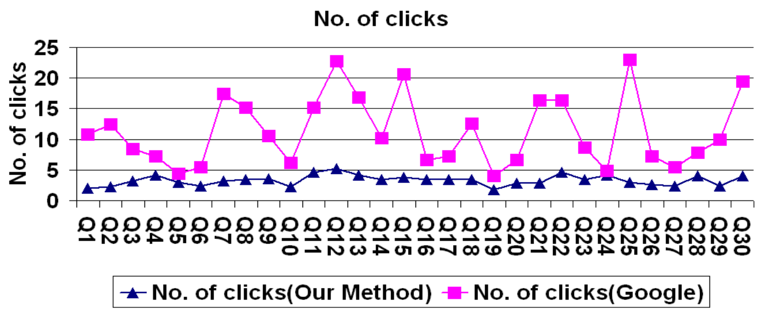


**Figure 5: Comparison of user mouse click numbers for answering the 30 questions using Google versus our search interface respectively. Here we compare the mean of user mouse click numbers using the two methods.**

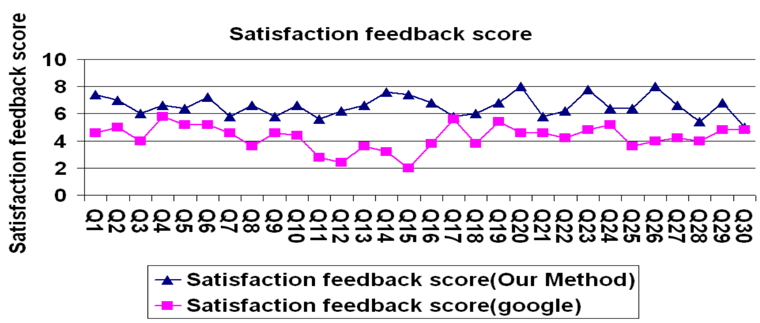


**Figure 6: Comparison of user satisfaction feedback scores of using Google versus our interface in answering the 30 questions respectively. Here we compare the mean of user satisfaction feedback scores using the two methods.**

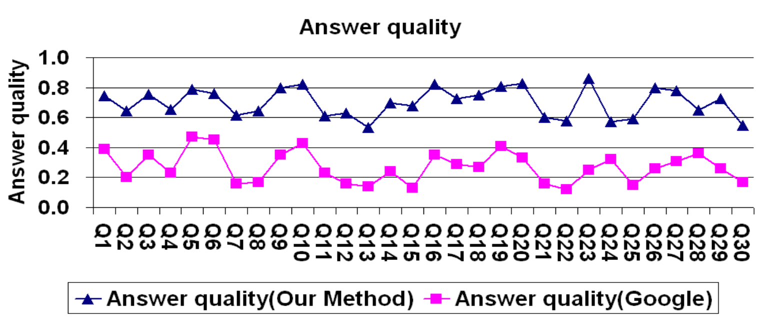


**Figure 7: Comparison of user question answering quality of using Google versus our interface in answering the 30 questions respectively. Here we compare the mean of user question answering quality using the two methods.**