

Queries? Links? Is there a difference?

Gene Golovchinsky[†]
FX Palo Alto Laboratory
3400 Hillview Avenue, Bldg. 4
Palo Alto, CA 94304 USA
+1 415 813-7361
gene@pal.xerox.com

[†]This research was conducted at the University of Toronto as part of the author's Ph.D. dissertation.

ABSTRACT

Hypertext interfaces are considered appropriate for information exploration tasks. The prohibitively expensive link creation effort, however, prevents traditional hypertext interfaces from being used with large coherent collections of text. Such collections typically require query-based interfaces. This paper examines a hybrid approach: the system described here creates anchors dynamically based on users' queries, and uses anchor selection as a query expansion mechanism. An experiment was conducted to compare browsing behavior in query- and link-based interfaces. Results suggest that query-mediated links are as effective as explicit queries, and that strategies adopted by users affect performance. This work has implications for the design of information exploration interfaces; the dynamic link algorithms described here are being incorporated into a Web server.

Keywords

Hypertext, dynamic links, browsing, newspaper metaphor, information exploration, information retrieval

INTRODUCTION

Hypertext

Hypertext interfaces have evolved to support information exploration tasks by providing users with a flexible interaction style that facilitates browsing through text [14]. The recent increase in the popularity of the Web has moved hypertext interfaces from the research domain into the lay public's almost implicit consciousness. The responsive interaction style with low cognitive overhead (compared with traditional query interfaces) may partially account for this increase.

There are some well-known problems with hypertext interfaces [6]. Users may not know how to get to the desired information, or how to return to previously-visited nodes. Some of these problems arise from the non-linear structure of hypertext; others may be related to hyperbase

size. As the number of nodes grows, the number of potentially relevant links emanating from any single node increases. At some point, the sheer number of links may overwhelm the reader. Furthermore, authors' reasons for including specific links may not always be clear to readers.

From an author's perspective, creating a large-scale hypertext is an arduous task. As the number of nodes increases, the number of possible (and desirable) links may increase geometrically. This link proliferation problem has limited most coherent hypertexts to a few thousand nodes; a 10,000-node hypertext is considered huge.¹ Automatic text-to-hypertext conversion has met with variable success [5], [8]. One problem with automatic creation of hypertexts is lack of feedback about which of the many potential links will be useful to the reader. Dynamic linking techniques based on user feedback have been implemented in some small systems (e.g., [15]), but these techniques have not been shown to scale to large collections.

Thus, hypertext has shown the promise of fluid, intuitive interaction but has suffered from limitations of scale. Some of these limitations have been resolved by incorporating query facilities into hypertext interfaces (e.g., [7], [2], [12]). These query interfaces, however, tended to be separate from hypertext browsing, rather than being incorporated into the same process.

Queries

Information retrieval techniques based on full-text indices allow fast, efficient searching of extremely large collections of text. Multiple gigabytes can be handled with ease, and techniques are being developed to handle even larger collections [11]. Efficient use of these systems has depended in part on expertise in query construction. Although some performance gains may be realized from finely tuned queries, there is substantial evidence (e.g., [3]) that casual users cannot master the required query syntax (Boolean or otherwise). Natural-language queries, on the

¹The Web is an exception because it is not coherent: each author creates a small coherent space that is arbitrarily connected to other such spaces. The number of nodes in each such space rarely exceeds a few hundred. In any case, the large web sites tend to use links for structural rather than for semantic purposes.

other hand, have been shown to be effective [16]. Belkin and Croft, for example, reported increased recall and precision of passage-based queries compared with Boolean queries [1].

Query-mediated hypertext

Syntactic problems associated with query formulation may be removed by accepting natural-language text and converting it to a query automatically (e.g., [16]). Interaction may be streamlined further if users can select passages from displayed text to specify their information-seeking intent. Thus, a query may be reduced to a drag-selection of some passage that would cause other, similar, passages to be displayed. Users could then hop from one document to another in a manner similar to link selection. The difference between this style of interaction and traditional hypertext is that the links are transient — no link is present until the reader creates it, and the link (potentially) disappears after the session ends.

It is also possible, however, to make the links more explicit. Given some representation of a user's information need, the system may retrieve relevant articles, and identify potential links within these articles. Selecting the anchor of one of these links would cause the system to use that anchor's context to expand or to modify the query. The results of the new query would then replace the source article, producing a hypertext-like connection between the source and the destination.

Thus it should be possible to create hypertext-like interfaces to collections of tens or hundreds of thousands of nodes, without incurring the costs of manual link construction. The following section describes in detail the interface of a query-mediated hypertext system; this description is followed by some experimental evidence for the effectiveness of this approach.

VOIR

VOIR (Visualization Of Information Retrieval) is an information exploration interface that implements several types of queries with a newspaper-like display and visualization tools to facilitate interactive browsing through large collections of text. It uses Inquery² [17] as the search engine, although the algorithms may be adapted easily to other search engines that produce ranked output. Nodes are stored in the database as SGML documents, although the tags are used only for formatting purposes. The system has been tested with a 250MB database (about 74,500 nodes). Reasonable performance is expected for collections up to ten times larger. The following subsections describe the

system's user interface (the newspaper metaphor) and its linking algorithms.

Newspaper interface metaphor

Hypertext systems typically use one-to-one links, although a significant minority supports one-to-many linking. When queries are used to mediate links, every link connects a node to an ordered set of other nodes. Thus, although it may be possible to open each node in a separate window and to let the user manage their layout, the large number of windows created in this manner would quickly overwhelm the user.

One solution to this problem is to use the newspaper metaphor to organize search results [13], [10]. Newspapers are designed to organize large numbers of loosely-related units of internally coherent text. Articles may also be grouped by topic, related articles appearing together. Finally, the amount of space allocated to an article and its position on the page provide the reader with cues to the article's importance.

The VOIR system uses the newspaper metaphor to arrange retrieved nodes in a multi-column page layout. It assigns more space to the more relevant articles. The text columns are fixed in size, but may be resized to reflect the degree of relevance of each article to the query [13]. Articles are grouped eight to a page, and enough pages are created to accommodate the entire set of retrieved articles. The user may flip from page to page, and may back up to any previous set of pages corresponding to a prior query. Figure 1 shows a screen shot of a VOIR interface.

Some care has been taken to make the interface appear as similar to a printed newspaper as possible. Heavy borders and standard widgets were avoided to foster the illusion of a printed page. A concession was made for scroll bars, but white space was used to separate columns when the entire article was visible. Different fonts were used for headlines and for text. These features were designed to encourage the acceptance of the newspaper metaphor. It is not clear how this design affected subjects of the experiment described below. Although it may be interesting to compare performance of this interface with a more traditional equivalent, differences between interfaces are unlikely to have significant impact on intermediate or expert users over long periods of time. It is possible, however, that novices may find such interfaces more familiar.

Whereas the newspaper metaphor addresses primarily the output aspects of the user interface, the input is handled through the selection of hypertext links. In addition, visualization techniques are used to help users orient themselves within the retrieved information. Link creation and selection, and visualization of search results are discussed in the sections that follow.

Topic specification

Different users frequently require different information from the same documents. Furthermore, the same user may

² Copyright (c) 1990-1994 by the Applied Computing Systems Institute of Massachusetts, Inc. (ACSIOM). All rights reserved. The INQUERY SYSTEM was provided by the Center for Intelligent Information Retrieval (CIIR), University of Massachusetts Computer Science Department, Amherst, Massachusetts. For more information, contact ACSIOM at 413-545-6311.

The Wall Street Journal® Hypertext

"Mulroney Sees Big Changes From Canadians' Demands"

New query
Done

Related

[Mulroney Sees Big Changes From Canadians' Demands](#) 12/24/90

By John Uquihart
Staff Reporter of The Wall Street Journal

OTTAWA -- Prime Minister Brian Mulroney said Canada's federal system must undergo "substantial change" to meet the demands of Quebecers and other Canadians.

He said at a news conference that "the alternative to a fundamentally reformed Canada is probably no Canada."

Mr. Mulroney didn't indicate what specific changes might be made in the federal system. But he said he is confident that "we will be able to devise a more modern, flexible brand of federalism that will accommodate the desires and aspirations of all Canadians."

The largely French-speaking province of Quebec is preparing demands for more autonomy, backed by polls showing record support for Quebec's gaining independence. In western Canada, protest movements are demanding constitutional changes to give western provinces greater influence in Canada's Parliament.

Mr. Mulroney said Quebec separatists would like a vote on Quebec independence

Related

[Canada's Prime Minister Proposes Shift Of Some Federal Powers to the Provinces](#)

By John Uquihart
Staff Reporter of The Wall Street Journal

02/13/91

Canadian Prime Minister Brian Mulroney said he is prepared to rewrite the Canadian constitution to prevent the nation's breakup.

"Change is possible, change is needed and change is coming," he said in a constitutional policy speech in Toronto.

Much of the proposed change is expected to involve a transfer of

Related

[Canada Parties Clear Plan to Rewrite Constitution to Meet Provincial Demands](#)

By John Uquihart and G. Pierre Goad
Staff Reporters of The Wall Street Journal

03/02/92

OTTAWA -- Canada's major political parties approved a plan to rewrite the constitution to deal with the longstanding demands of Quebec, other provinces and native Canadians.

The plan would give Quebec some of the powers it seeks, but whether it can win the necessary broad support in the rest of Canada was uncertain. Most of the constitutional changes proposed in the plan require the approval of the federal government and at least seven of Canada's provinces; some require the consent of all 10 provinces.

Quebec Premier Robert Bourassa plans to comment in detail on the federal blueprint tomorrow, a spokeswoman said.

At first glance the federal plan falls short of the demands by the ruling Quebec Liberal Party. Opposition leader Jacques Parizeau told reporters in Quebec City that the federal plan is

Related

[Canada Talks Stall After Fifth Day Over Quebec's Bid for Special Status](#)

By John Uquihart and G. Pierre Goad
Staff Reporters of The Wall Street Journal

06/08/90

OTTAWA -- Negotiations to resolve Canada's constitutional impasse stalled last night after five days of talks, and several provincial premiers warned that the situation is critical.

Canadian Prime Minister Brian Mulroney and the nation's 10 provincial premiers have been meeting here since Sunday night to try to resolve the constitutional dispute. They will meet again this morning. The premiers are stalemated over recognition of French-speaking Quebec as a "distinct society" within Canada.

"We say enough is enough," said Quebec Premier Robert Bourassa. Quebec "can't and won't accept amendments" to the Meech Lake accord, a package of proposed constitutional amendments. One Meech Lake clause recognizes Quebec as a "distinct society."

The provision is highly prized by Quebec as recognition of its cultural uniqueness even though constitutional lawyers say it won't add

Related

[Mulroney Detects Progress in Talks Over Constitution](#)

By John Uquihart and G. Pierre Goad
Staff Reporters of The Wall Street Journal

06/06/90

OTTAWA -- Canadian Prime Minister Brian Mulroney reported a "small degree of progress" on the third day of negotiations on Canada's constitutional impasse.

Mr. Mulroney declined to disclose details of the negotiations, which involve himself and 10 provincial government leaders. But government officials said that much of yesterday's negotiations involved differences over a proposed constitutional amendment that would set new rules for changing the federal Senate.

The proposed amendment is part of the constitutional accord that was negotiated three

Related

[International: Quebec Ruling Party Expected to Unveil Its Proposal to Keep Province in Canada](#)

By G. Pierre Goad
Staff Reporter of The Wall Street Journal

01/29/91

MONTREAL -- The ruling Quebec Liberal Party is to unveil today a constitutional plan that is expected to demand greater powers for the French-speaking province and a drastic decentralization of the

Related

[International: Quebec Leaders Clear a Hurdle in Impasse Involving Constitution](#)

Tentative Accord Is Reached On Status for Quebec, But Opposition Remains

By G. Pierre Goad and John Uquihart
Staff Reporters of The Wall Street Journal

06/11/90

OTTAWA -- Canada's leaders patched together a tentative

Related

[International: Quebec Ruling Party Expected to Unveil Its Proposal to Keep Province in Canada](#)

By G. Pierre Goad
Staff Reporter of The Wall Street Journal

01/29/91

MONTREAL -- The ruling Quebec Liberal Party is to unveil today a constitutional plan that is expected to demand greater powers for the French-speaking province and a drastic decentralization of the

Related

[Mulroney Sees Big Changes From Canadians' Demands](#)

By G. Pierre Goad and John Uquihart
Staff Reporters of The Wall Street Journal

06/05/90

OTTAWA -- Canadian Prime Minister Brian Mulroney and the ten provincial premiers will resume talks this morning in a bid to resolve Canada's constitutional impasse after making only limited progress yesterday.

Mr. Mulroney said last night that a federal government proposal

Related

[Canada's Prime Minister Proposes Shift Of Some Federal Powers to the Provinces](#)

By John Uquihart
Staff Reporter of The Wall Street Journal

02/13/91

Canadian Prime Minister Brian Mulroney said he is prepared to rewrite the Canadian constitution to prevent the nation's breakup.

"Change is possible, change is needed and change is coming," he said in a constitutional policy speech in Toronto.

Much of the proposed change is expected to involve a transfer of

History

- 9: Mulroney Sees Big Changes From Canadians' Demands
- 8: prime minister brian mulroney of
- 7: canadian senate
- 6: canadian prime minister brian mulroney
- 5: Canadian Prime Minister Brian Mulroney
- 4: International: Canada's Leader: Mulroney Sees Big Changes From Canadians' Demands
- 3: french-speaking quebec
- 2: quebec special status ottawa ce

Cumulative rank

Average rank

Topic

20 matching articles found

Link Quality

low Max # docs 30

Computer Library materials copyright owned by Ziff Communications Company. Wall Street Journal/Material (c) 1990, 1991, 1988 Dow Jones Inc.

Figure 1. Screen shot of VOIR interface.

have different needs at different points in time. This observation should be reflected in the nature and in the quantity of links that are available at each document. The more information is available about a user's topic of interest, the more appropriate the links can be.

VOIR maintains a query context that represents a user's current information-seeking intent. This context consists of a number of keywords either provided by the user or inferred by the system. A user specifies an initial search topic either by typing in a set of keywords, or by drag-selecting a passage of text already displayed on the screen. These queries reset the search context representation.

When the articles retrieved by a particular query are displayed, some of the query keywords are displayed as anchors (i.e., blue, underlined) in the text. Terms are selected based on the degree to which a term discriminates among documents in a collection. Content-bearing terms are thus likely candidates for linking related articles. In addition, proper names and names of places, companies, products, etc., can also serve to connect related articles; they too become anchors when present in the query text. (See [9] for a more detailed discussion of the anchor selection algorithms.) Throughout the rest of the paper, anchors created based on the statistical properties of query terms will be called *dynamic anchors*. Similarly, the links that are activated by the selection of these anchors (where activation implies the construction and execution of a query) will be called *dynamic links*, or *query-mediated links*.

Context-specific links

When the user selects one of the dynamically-inserted anchors, the system uses that selection to expand the previous query.³ Content-bearing terms in the sentence containing the selected anchor are combined with the terms from the three previous context-specific links to produce a weighted query. The selected anchor is assigned the highest weight; the other anchors in the sentence the next highest; then the rest of the content-bearing words in the sentence, and finally the terms from prior queries. These terms are included to give the query stability (to preserve the search context) when the selected anchor occurs in a sentence with few highly discriminating terms.⁴ The documents retrieved by this link-triggered query are then displayed on the screen, replacing the source documents. The newly-shown articles are again marked up to show anchors derived from the last query.

A balance must be maintained between query stability and query flexibility: if the terms of all prior queries were included into each successive query, the set of retrieved articles would very quickly become fixed. To give the user implicit control over the context of the search, and to allow

natural evolution of search interests, only the three prior queries contribute terms to the context of the next query.

The user may reset this context at any time by selecting a new passage or by typing in a new set of keywords. The system responds to this selection by identifying a new set of links appropriate to the new topic.

Context-independent links

Context-independent links display collections of articles related to the current article independently of the prior context. Whereas context-specific links use the passage containing the selected anchor to expand the query, context-independent links ignore the current context and use the leading paragraphs of the selected article as a query. Relevance feedback is used to retrieve additional articles similar to the selected one. The articles retrieved in this manner are cached; subsequent context-independent link selections for the same article cause the cached articles to be displayed. Although computed dynamically, these links represent a sort of a static graph of the hyperbase. Rather than relying on a narrow context to determine document relevance, these links tend to co-relate documents based on their entire contents. These links may be used to follow predictable paths through the document database.

Visualization

Often the same node is retrieved several times during the course of a browsing session. Although it is possible to suppress the redisplay of previously-viewed nodes, this strategy may prevent the user from observing interrelationships among nodes. On the other hand, users may become frustrated when they keep "rediscovering" the same node over and over. The solution adopted in VOIR is to provide a graphical summary of the retrieval history of each article. This summary appears as a histogram above the article text (Figure 2). Each bar in the histogram represents the relative importance of that article with respect to some query or link selection. The more relevant the article, the higher the bar. The bars are accumulated from left to right, corresponding to the temporal sequence of queries. The colors of the bars represent the three types of links: context-setting queries (drag-selections and typed keywords) are coded red; context-specific links are blue; and context-independent links are green. The colors correspond roughly to the colors used to represent the links: reddish selection background for passage selections, blue imbedded links, and green context-independent button labels. Thus the display of Figure 2 reveals that this article has been retrieved six times out of nine queries (with three gaps), and that it was more relevant to the last query — the context-independent one — than to the first few. This information may be used to determine whether an article has been seen before, and to which other articles it is related.

³ The query that retrieved the article containing the selected link.

⁴ Sentences containing many pronouns, for example.



Figure 2. Retrieval history visualization.

The histogram may be used as an interactive navigation aid. Clicking on any histogram bar will cause the corresponding link or query to be selected in the history list. Double-clicking on a bar will cause the system to restore the corresponding query. Thus dragging the cursor over the histogram with the left mouse button pressed will cause the history list to scroll through the corresponding items.

A global overview of the browsing session is also provided: a scatterplot is maintained, with each point representing a retrieved node. For each node, the cumulative score (the sum of the heights of its bars) is plotted on the vertical axis; the average of heights (when the article was retrieved) is plotted on the horizontal axis. The two axes are not linearly related because the average score (the horizontal axis) is computed over the number of times a node is retrieved, not the number of queries issued by the user. Thus nodes are not penalized for being retrieved infrequently.

This arrangement produces some clusters in the display: the top-right corner contains the landmark nodes (frequently retrieved, highly relevant), and the bottom-right corner contains the highly-specialized nodes (seldom retrieved, highly relevant when retrieved). Figure 3 shows an example of this visualization.




Figure 3. Global overview of the browsing session, highlighting the location of one article.

The global visualization display also supports interaction. When a user positions the cursor over an article's histogram, the point that corresponds to the selected article is highlighted in red in the scatterplot (small box around one of the points in the top-right corner of Figure 3). When the user drag-selects a region in the scatterplot, all articles corresponding to points included in the selected area are displayed as if in response to a link selection.

Ideally, the global overview display should contain semantic clustering (e.g., [4], [18]). That is, points representing documents should be positioned in the space to indicate content similarity of the respective documents. The closer two points are, the more similar the corresponding documents. This type of display would allow users to detect groups (“neighborhoods”) of documents related to a particular subject. Retrieving one

document in such a cluster would suggest to the user which other documents may be of interest.

The decision of which articles to display is not unambiguous from a user interface perspective, however. Although it is possible to pre-compute a clustering of all documents in the database, such a clustering will not reflect the current search context, and thus it is likely that most of the information in the display will not be useful for any particular search. On the other hand, it is possible (and computationally feasible) to re-cluster each result set independently. The problem is that this approach will produce highly unstable displays: the same document will appear in a different place for each query, decreasing the likelihood that the user will gain a better understanding of the organization and scope of the database. These considerations suggest that a hybrid algorithm that combines the stability of a fixed clustering with the context of the current browsing session (perhaps in a fisheye view) should be more appropriate.

EXPERIMENT

An experiment was conducted to test the relative efficiency of the link algorithms and interfaces. A simplified, instrumented version of VOIR (without context-independent links and without the global overview) was used to browse a portion of the TIPSTER collection [11] of about 74,500 Wall Street Journal articles. The experiment was designed to assess subjects' performance when using natural-language queries and query-mediated hypertext links. Three interface conditions were used in the experiment. The *query* condition allowed subjects to select passages in the text and to type keywords; no links were available, although matching query terms were highlighted in yellow in the retrieved articles. The *naïve link* condition provided dynamic links that subjects could use in addition to queries, but subjects were only told that links would show related articles; they were not told how that would be accomplished. Finally, subjects in the *informed link* condition were told how context-specific links worked. The difference between the query and the link conditions existed only in the interface. The links available to subjects were created automatically by the software, and were mediated by queries. The two link conditions were introduced to assess how intuitive users found the dynamic links. Strong differences between the naïve and the informed link conditions would indicate that some training would be required of potential users.

Subjects

Twenty four subjects participated in the experiment. Subjects were either enrolled in a graduate program or had recently completed one. Subjects were familiar with the use of computers for information retrieval purposes; most had at least some familiarity with hypertext interfaces. Subjects were paid \$20 upon completion of the experimental session.

Methodology

The experiment was a between-subjects repeated measures design. Each subject was exposed to one of three interface conditions (query, naïve link, or informed link). All subjects performed searches on the same six topics, presented in random order. The subject's task was to use queries or links to identify as many articles as possible; relevance criteria were given for each topic. Subjects were asked to read the topic description before starting the search session. Each search session lasted 15 minutes, and was terminated automatically by the program. This duration was chosen as a compromise between experimental convenience (short is good) and ecological validity, and to allow comparison with other studies previously conducted in our lab. The experimental software was instrumented to record user-initiated events and system responses. This log data was then used to analyze subjects' performance and behavior.

Topics were selected to represent a range of task difficulty. The number of relevant articles for each topic ranged from 31 to 97 out of 74,500. Relevant articles — determined by expert raters — for each topic were provided by the National Institute of Standards and Technology (NIST) as part of the TIPSTER corpus. Using these relevant articles, recall and precision measures were recorded for each subject in each topic, and within each topic by the type of query used (typed, selected or link). Recall and precision scores were calculated based on three sets of articles: the set of all articles retrieved by the query, the set of articles viewed by each subject, and the set of articles judged relevant by the subject (Figure 4). The time to retrieve the first relevant article was also recorded.




Figure 4. Relationship between sets of articles used to compute recall and precision measures.

Results

Analysis of variance was performed using topic and interface as the independent variables. The dependent measures were recall and precision based on retrieved, viewed and selected articles. Recall and precision scores were calculated by pooling the results of all interactions (links and queries) performed for a single topic. These measures are designed to compensate for backtracking and for repeated queries. There were no significant differences in performance due to the interface; as expected, significant differences between topics were found: of the six measures, the *least* significant effect was $F[5,143]=3.84$, $p<0.05$. There was no significant difference ($F<1$) between interface conditions with respect

to the time to retrieve the first relevant article. Statistically-reliable differences in effectiveness among the different query types (passage selection, typed query and dynamic link) were found. When subjects used dynamic links (regardless of interface condition), they achieved higher recall and precision scores compared with using typed queries or passage selections. See [9] for a thorough discussion of these results.

No learning effect was found by comparing the performance on the first topic with that of the other five, or by comparing the first three to the last three topics. A learning effect may have been masked by the large between-topics variance.

Cursory examination of subject behavior suggested that subjects differed in the strategies they used to form queries. Subjects were categorized into two clusters based on the total number of queries and the numbers of queries of each type (link, passage and typed) they issued per search topic. T-tests showed that subjects who made more queries had higher recall (see Table 1) without loss of precision. Subjects in one cluster tended to spend a lot of time reading articles to determine their relevance to the search topic, while other subjects tended to skim over the retrieved documents, making many quick relevance judgments. Thus the two groups were characterized as “readers” and as “skimmers,” respectively. On average, skimmers had almost twice the judged recall score (0.18 vs. 0.11) without a statistically-significant loss in precision.

Table 1. Recall measures for query strategies (Readers vs. Skimmers).

Measure	df ⁵	t	p <
Retrieved recall	142	-4.20	0.001
Viewed recall	142	-3.34	0.001
Judged recall	74.2	-3.85	0.002

Subjects' browsing strategies (as distinct from query strategy, above) were analyzed further based on their use of different navigation mechanisms available to them. First, the behaviors of the sixteen subjects from the link interface conditions were analyzed; subsequently all 24 subjects' data was used.

Browsing strategies were characterized by patterns of query use by type: link (L), passage (P), and typed (T). A run of degree $n-1$ was defined as n consecutive instances of the same type. Transitions between types were also defined. For example, the sequence PPTTTLP has a P run of length one, a T run of length two, an L run of length zero, and three transitions. Run and transition scores were computed for each subject by normalizing the raw scores to add to 1.0 for each subject-topic pair.

⁵ The fractional degrees of freedom reported on t-tests have been calculated using Satterthwaite's approximation to compensate for unequal variance. See the SAS/STAT User's Guide, Version 6, Chap. 42, p. 1637.

The 16 subjects participating in the link conditions were clustered based on the values of the four variables (three runs and transition) using Ward's minimum variance method. Three clusters were identified: subjects in the first cluster ("Linkers," n=5) preferred to use links; subjects in the second cluster ("Typers," n=3) typed many of their queries; and subjects in the third cluster ("Shifters," n=8) switched between different query types frequently. Table 2 shows the mean scores for the four measures for each cluster. Boldface numbers show the dominant dimension for each cluster.

Table 2. Means for linking strategy clusters.

Measure	Linkers (n=5)	Typers (n=3)	Shifters (n=8)
Query	0.083	0.642	0.190
Link	0.584	0.033	0.065
Passage	0.055	0.026	0.180
Transition	0.278	0.299	0.565

Analysis of variance was performed using cluster categorization as a pseudo-independent variable. Significant differences in performance were found between Typers and the other subjects. Typers tended to have higher viewed recall ($F[2,88]=7.25$, $p<0.005$) and lower viewed precision ($F[2,88]=3.83$, $p<0.05$) than Shifters or Linkers. No significant differences in performance were found between Shifters and Linkers.

A second cluster analysis was performed on all 24 subjects. The distinction between *production* (typed queries) and *selection* (link and passage selections) was used in conjunction with the number of transitions to characterize subjects' performance. Cluster analysis revealed four clusters, as shown in Table 3.

Table 3. Means for strategy clusters.

Type	Selectors (n=3)	Producers (n=6)	Shifters (n=6)	Balancers (n=9)
Production	0.024	0.612	0.168	0.240
Selection	0.810	0.060	0.200	0.428
Transition	0.166	0.329	0.633	0.332

Analysis of variance was performed using cluster categorization as a pseudo-independent variable. Four clusters — "Selectors" (n=3), "Producers" (n=6), "Shifters" (n=6), "Balancers" (n=9) — were identified. Selectors obtained higher retrieved recall scores than others ($F[3,135]=3.58$, $p<0.025$), and higher viewed precision than producers ($F[3,135]=2.74$, $p<0.05$). Producers exhibited higher viewed recall than shifters and Balancers ($F[3,135]=4.02$, $p<0.01$).

Discussion

The results of this experiment suggest that query-mediated links may serve as an effective alternative to natural-language queries, particularly in situations in which it is

awkward to construct many queries. Clear differences in query strategy were found between subjects. "Readers" concentrated on a few articles and made relatively few relevance judgments, while "Skimmers" made many queries and many relevance judgments. The skimming strategy appears to have paid off: subjects who ran more queries obtained higher recall scores without sacrificing precision.

Analysis of browsing (rather than query) strategies that considered subjects' choice of navigation method and sequences of use revealed several clear patterns. Some of these patterns were associated with better performance than others, but some tradeoffs were also observed. Overall, it appears that subjects tend to select different combinations of navigational commands, but these choices result in trading off the different performance measures. Thus these browsing strategy choices may be attributed to intrinsic factors — preference, prior experience, perhaps personality traits — rather than to effects induced by the interface.

Although it was possible to design the experiment to force subjects to use links and to prevent them from reading articles exhaustively, the more naturalistic approach adopted here revealed distinct differences in the way subjects approached a given task. Some subjects reported that they liked the interface because it provided multiple ways of expressing information-seeking intent. This result suggests that designers of information retrieval systems should consider incorporating multiple browsing interface mechanisms to support users with different retrieval strategies.

It is desirable to compare the performance of query-mediated links to static links; such comparisons, however, may be impractical to achieve given the difficulty of statically linking hundreds of megabytes of text.

APPLICATIONS

The VOIR prototype has been tested extensively with portions of the TIPSTER collection of newspaper articles. It has also been used with the HCI Bibliography abstracts converted from Refer format to SGML. The HCI Bibliography is a relatively small collection, containing only about 12MB of text. This size results in extremely short response times, although some documents are too small for effective indexing. Nonetheless, the interface seemed quite effective when used by the author to conduct literature searches.

An attempt to index a collection of UNIX man pages was not successful. It was too difficult to establish a meaningful stop-word list, and it appears that the approach advocated here works well with texts containing large numbers of unique words. Documents in the man collection, on the other hand, tended to have formulaic wording with few discriminating words. Most of the content was borne by syntactic symbols whose definitions were local rather than global. Thus, although it was possible to search the

database for some concepts, iterative link selection was not effective.

We are currently developing query-mediated linking algorithms for browsing the Web. The approach involves inserting additional links into existing Web pages based on topic specifications provided by users; indexing may be done locally to the Web server, or a remote index server may be used. This approach holds promise for simplifying Web-based searching interfaces, and may prove to be a viable complement to static links.

CONCLUSIONS

This paper has described techniques for merging a traditional query-based information retrieval engine with a hypertext interface. Hypertext and information retrieval techniques have been integrated — rather than segregated — by the interface. An experiment was conducted to test the relative effectiveness of query-mediated links versus explicit queries. Experimental results indicate that query-mediated links can be as effective as user-specified queries. Furthermore, there is evidence that individual differences in subjects' preferences can affect search effectiveness. This result suggests that query-mediated links can augment or replace explicit queries in cases where query interfaces may be awkward to use. Empirical evidence regarding browsing strategy appears to answer the question posed in the title. There is a difference between queries and links: the difference is in the choices people make when using the interface.

ACKNOWLEDGMENTS

The author wishes to thank Mark Chignell, David Modjeska, Lynn Wilcox and the anonymous reviewers for their helpful comments. This research was funded by a grant from the Information Technology Research Centre of Excellence of Ontario (ITRC).

REFERENCES

1. Belkin, N.J. and Croft, W.B. Retrieval techniques. Chapter 4 In *ARIST*, M.E. Williams, Ed. Elsevier, 1987, 109–145.
2. Belkin, N.J., Marchetti, P.G. and Cool., C. BRAQUE: Design of an Interface to Support User Interaction in Information Retrieval. *Journal of Information Processing and Management* (1993)
3. Borgman, C.L. Why are online catalogs hard to use? Lessons learned from information retrieval studies. *JASIS*, 37, 1986, 387–400.
4. Chalmers, M. and Chitson, P. Bead: Explorations in Information Visualization. In *Proceedings of SIGIR '92*, (Copenhagen, Denmark, 1992), ACM Press, 330–337.
5. Chignell, M.H., Nordhausen, B. Valdez, F. and Waterworth, J.A. The HEFTI Model of Text to Hypertext Conversion. *Hypermedia*, 3, 3 (1991)
6. Conklin, J. Hypertext: An Introduction and Survey. *Computer Magazine*, 220, 9, (Sept. 1987), 17–41.
7. Egan, D.E., Remde, J.R., Gomez, J.M., Landauer, T.K., Eberhardt, J. and Lochbaum, C.C. Formative Design-Evaluation of SuperBook. *ACM TOIS*, 7 (1), (1989), 30–57.
8. Furuta, R., Plaisant, C., and Shneiderman, B. Automatically transforming regularly structured linear documents into hypertext. *Electronic Publishing*, 2, 4, (1989), 211–229.
9. Golovchinsky, G. What the Query Told The Link: The Integration of Hypertext and Information Retrieval. In *Proceedings of Hypertext '97* (Southampton, UK, April 1997)
10. Golovchinsky, G. and Chignell, M.H. The Newspaper as an Information Exploration Metaphor, *Journal of Information Processing and Management*, in press.
11. Harman, D.K. Overview of the Third Text REtrieval Conference (TREC-3). Harman, D.K. (Ed.) NIST Pub. 500-225, (1995), 1–19.
12. Hertzum, M. and Frøkjær, E. Browsing and Querying in Online Documentation. *ACM TOCHI*, 3, 2, (1996), 136 - 161.
13. Kamba, T., Bharat, K. and Albers, M.C. The Krakatoa Chronicle — An Interactive, Personalized Newspaper on the Web. Georgia Tech Technical Report GIT-GVU-95-25, (1995)
14. Nielsen, J. *Hypertext and Hypermedia*. Academic Press, (1990)
15. Stotts, P.D. and Furuta, R. Dynamic Adaptation of Hypertext Structure. In *Proceedings of Hypertext '91* (San Antonio, TX, 1991), ACM Press, 219–231
16. Turtle, H. Natural Language vs. Boolean query evaluation: A comparison of retrieval performance. In *Proceedings of SIGIR '94* (Dublin, Ireland, 1994), Springer Verlag, 212–220.
17. Turtle, H. and Croft, W.B. Inference Networks for Document Retrieval. In *Proceedings of SIGIR '90*, (1990), ACM Press, 1–24.
18. Wise, J.A., Thomas, J.J., Pennock, K., Lantrip, D., Pottier, M., Schur, A., and Crow, V. Visualizing the Non-Visual: Spatial Analysis and Interaction with Information from Text Documents. In *Proceedings of Information Visualization '95* (Atlanta, GA, 1995), IEEE Press, 51–58.