

Orienteering* in an Information Landscape: How Information Seekers Get From Here to There

Vicki L. O'Day, Robin Jeffries
Software and Technology Laboratory
HPL-92-127 (R.1)
September, 1992

information search
and retrieval (search
process), information
use, intermediaries,
collaborative work

We studied the uses of information search results by regular clients of professional intermediaries. The clients in our study engaged in three different types of searches: 1. monitoring a well-known topic or set of variables over time, 2. following an information-gathering plan suggested by a typical approach to the task at hand, and 3. exploring a topic in an undirected fashion. In most cases, a single search evolved into a series of interconnected searches, usually beginning with a high-level overview. We identified a set of common triggers and stop conditions for further search steps. We also observed a set of common operations that clients used to analyze search results. In some settings, the number of search iterations was reduced by restructuring the work done by intermediaries. We discuss the implications of the interconnected search pattern, triggers and stop conditions, common analysis techniques, and intermediary roles for the design of information access systems.

*Orienteering is a sport in which individuals or teams try to complete a course leading from a common starting point through fixed intermediate points to a common destination. Orienteering events are typically held in semi-wilderness areas or large parks. Participants navigate with a map and compass, using landmarks in the environment to get their bearings and plot a route to the next point in the course. Beginners may be satisfied just to find their way to the end, but getting there faster is considered better.

1 Introduction

An important way of looking at the effectiveness of information retrieval is to see how well the resulting material and the manner of retrieval and delivery match the user's problem situation. As research directions have shifted to include the study of users of information systems, rather than concentrating solely on system characteristics [4, 7], most efforts have centered on understanding how users can express their information needs effectively to the system (e.g., [2, 3, 5]). In contrast, our work focuses on information delivery, rather than search performance – how the results of an information search are digested and used to solve the problem that prompted the search. To understand information delivery requirements, we conducted a study of how library clients deal with the information they get back from human intermediaries.

Our hypothesis, based on an examination of search examples in several domains, was that a single, static view of search results, however appropriately designed, would often not be enough to satisfy real user needs; instead, users would use the results as a basis for exploration, perhaps examining them from different perspectives. We learned from our study that people did do exploration, often conducting a series of interconnected but diverse searches on a single, problem-based theme, rather than one extended search session per task. Like practitioners of the sport of orienteering, our searchers used data from their present situation to determine where to go next. We identified a set of common triggers and stop conditions for further search steps within an extended search. We also learned that people used a set of common analysis techniques on their data, which often involved constructing new artifacts.

The motivation for this study was in our work on information retrieval mediators [9]. We are currently building a toolkit called DIME that will allow people to create, share, and use mediators [8]. A retrieval mediator is a small, special-purpose packaging of the retrieval expertise necessary to handle a prototypical information question in a particular domain. It includes a search strategy, iterative steps if necessary to complete the search task, post-processing of retrieved data, and default choices for information presentation. A user-level information-seeking task would typically involve several mediators. Private mediators will be used to provide a level of individual customization and extensibility in searching. Mediators will be identifiable entities in the system; high-level, task-oriented descriptions of mediators will be searchable by mediator-builders or end users.

In this paper, we describe the results of our study and the implications for the design of the mediator toolkit and other information access systems.

2 Methodology

We conducted an ethnographic study of fifteen clients of professional intermediaries. About half of our informants worked for the same large computer manufacturer and half worked for a variety of other companies. All of our informants regularly conducted mediated searches on financial and business-related topics. We chose this focus in the hope that information use patterns might be clearer if there were some commonality in the language and domain expertise of the informants, but the informants were by no means a homogeneous group.

A variety of professions was represented, including financial analyst, venture capitalist, product marketing engineer, demographer, management consultant, logistics specialist, research assistant, statistician, merger and acquisition specialist, college finance officer, and

sales/planning staff. The job levels of the informants ranged from entry-level to senior management. Their work settings were diverse, including engineering, education, finance, and consulting organizations. Some informants worked for companies with fewer than twenty-five employees; some worked for companies with tens of thousands of employees. Some were sophisticated users of technology, while others only used a small handful of tools such as spreadsheet programs. The backgrounds of the informants were in business, technical, and social science disciplines.

Semi-structured interviews were conducted in the offices of the participants. We audiotaped all of the interviews for which we were given permission, resulting in over 150 pages of transcripts. The length of interviews ranged from forty-five minutes to two hours; the average interview lasted about an hour.

The interview questions were designed to expose the work contexts and specific tasks that triggered search requests and to explore example search requests in detail. Informants were asked how they expressed requests to the intermediary, which formats were used for delivery of results, and what they did to understand and use the information when it arrived. When possible, the researcher and the informant examined together past search results that the informant had saved in a printed form, along with any resulting summaries, graphs, tables, and newsletters that had been produced using these search results.

It must be noted that the data gathered are based on the informants' memories of their searches, augmented by artifacts they had saved, and thus reflect whatever aspects of the searches the informants found memorable. No observations were made of actual interactions with intermediaries or "live" analyses of search results. Our informants were generous in sharing real examples of their past searches, but their current analyses would have contained material too sensitive to share.

In almost all of the work settings, the intermediaries were professional librarians. A few of the intermediaries were trained searchers without library science degrees. Throughout the rest of the paper, we will usually refer to the intermediaries as librarians.

We did two levels of analysis of the interview transcripts. At the first level, we looked for themes or patterns that appeared consistently throughout our informants' search experiences. Two patterns did become apparent. One of these, a pattern of interconnected searches, will be described in this paper. The other pattern, that virtually all of the library clients in our study turned out to be intermediaries themselves of one sort or another, will not be discussed here due to space limitations.

In the second level of analysis, we extracted from the transcripts all of the search scenarios described by the participants in the study. These search scenarios were characterized by three basic "search modes", which we will describe below. To explore the theme of interconnected searches mentioned above, we paid particular attention during our analysis to the transition steps between individual searches in a progressive search. We extracted from the transcripts the triggers and stop conditions that influenced people's decisions of whether to do further search steps. In some settings, we noted that these triggers and stop conditions were modified by an unusual work breakdown between librarian and client. Since we are interested in how people use information returned by a search, we also extracted from the transcripts the set of operations that were used by our informants to analyze their search results.

During this second phase of the analysis, two coders analyzed one-third of the transcripts and reconciled their differences in each of the three coding areas: the search scenarios found

and assignment of each to one of the three search modes, triggers and stop conditions found, and analysis operations found. The differences between coders were few; the reconciliation mainly consisted of adding missed examples to the lists of triggers and analysis operations. The remaining two-thirds of the transcripts were analyzed by one coder.

In the following section, we will describe the overall pattern of interconnected searches and three different search modes, the connections between steps in a progressive search, the impact of changing the role of the librarian, and frequently-used analysis operations.

3 Findings

3.1 Search Modes and Interconnected Searches

Our informants described 66 different search scenarios. Some of these scenarios were clearly one-time, ad hoc searches, while others represented repeated searches done either at regular intervals or whenever the informant worked on a particular task. Thus a single scenario often represented many executions of a search. In our coding of search scenarios, we distinguished between scenarios that were described very specifically, with reference to particular search parameters, and scenarios that were described generically by our informants.

These search scenarios fell into three basic “search modes”: *monitoring*, *planned*, and *exploratory*. We give examples of the three modes below. A small number of scenarios (six) fell into two categories rather than one; we have folded these scenarios into the frequency counts given below by assigning one-half point to each of the two categories involved. Five of the scenarios could not be categorized due to insufficient data.

The three search modes were:

1. *Monitoring* a well-known topic or set of variables over time. (4 generic, 10.5 specific)
Katherine¹, a financial analyst, calls the librarian every quarter to request a search on four competitor companies. She asks for their public earnings announcements, which she uses to track revenue and order growth, and she also asks for any references to the companies in the business literature so she can see if there is “anything unusual about this company.” She creates a quarterly report with tables comparing her company’s performance to the competitors, with explanatory notes at the bottom.
2. *Following a plan* for information-gathering suggested by a typical approach to the task at hand. (9 generic, 12.5 specific)
Larry, a process benchmarking specialist, helps groups inside his company improve their business processes in areas such as order fulfillment. When handling a particular assignment, he follows a typical plan in his library searches. First, he looks for the companies that do an outstanding job with the particular process he’s investigating. To do this, he has the librarian search for any references to that process, giving her a list of synonyms that might be used to describe the process. Once he has identified the companies with the best processes, he follows up by having the librarian retrieve information about the financial performance of each of the companies, since if a company

¹ We have anonymized examples when necessary, and we have not used the informants’ real names in this paper.

isn't doing well there may be something wrong with the process. When he's ready to visit one of the top companies, he has the librarian run a search to find any recent news on the company, so he will feel completely up-to-date before the visit.

3. *Exploring* a topic in an undirected fashion. (8 generic, 17 specific)

Bob, a management consultant, received a request from a client to give advice on certain retail distribution issues. Before getting started with the client's problem, Bob decided to look into the "basic current facts on the industry." He felt that he needed to know "whether what we are talking about is smaller than a bread box, bigger than a house - just size it." He requested general information on the industry from his librarian.

There were about twice as many specific as generic search scenarios, with a stronger bias towards specific scenarios in monitoring searches, where the library clients performed the same search on a weekly, monthly, quarterly, or annual basis and were usually able to describe very precisely what they did. The generic search scenarios may represent more actual searches than specific (non-monitoring) scenarios and may therefore be more fruitful for understanding the search process, but it is helpful to have the data anchored by a plentiful set of specific examples. There were roughly the same number of planned and exploratory searches overall, with fewer monitoring searches. However, each monitoring search scenario, whether generic or specific, represents many actual searches.

We also looked at the distribution of the use of different search modes across individual library clients, and we found that every client did more than one type of search. All but two clients did exploratory searches; planned searches and monitoring searches were done by about two-thirds of the clients.

A pattern of behavior that almost all the library clients followed was to conduct over time a series of interconnected but diverse searches on a single, problem-based theme, rather than one extended search session per task. This orienteering-style approach occurred in each of the three search modes. Some, but not all, monitoring searches led to further investigations, primarily when something new or unexpected turned up regarding the companies or other variables that were being monitored. Virtually all planned and exploration searches were done in several steps. Each stage in the search was typically followed by careful reading, assimilation, and analysis of all of the resulting material, which then triggered new directions to explore. The winding paths taken in these searches could sometimes be partially predicted by the task that motivated the search, but even in apparently routine searches there were often unexpected turns in the journey.

Here is a high-level description of a progressive search given by John, a merger and acquisition expert:

John: What you want is a thorough and efficient way that will cover first of all the leading sources and then second of all more localized sources ... So there is a kind of a general quick and dirty wanting to find out what's out there and then once you do that then you want to go in more specifically and that is where you get into more detailed searches. There is like a 30,000-foot view and then you go into specific areas.

The pattern of looking for an overview first was common to both planned and exploration searches, followed by detailed searches suggested either by the plan or by results of the initial search. This high-level view was an orienting view, often serving as an index into points of

interest. Monitoring searches tended to focus on details from the beginning, since the context of the competitor companies or other monitored variables was already well-understood.

We would like to emphasize that this kind of progressive search is different from the refinement of a query within a search session, which is intended to gradually bring the result set closer to a “perfect” match to the user’s information need. Systems such as those described by Williams [10] and Fischer and Nieper-Lemke [6] allow the user to provide feedback on the relevance of terms appearing in the result set, after which the query is redefined to focus on items with these terms. Such reformulation strategies assume that the user’s information need is fixed throughout the search session, and a dialogue of sorts is taking place between the system and the user to help the system match the user’s fixed need.

In the interconnected searches conducted by our informants, the information need shifted between search requests. Within a session, the information need was of necessity fixed, since librarians conducted the actual searches and there was rarely any interaction between a librarian and a client during a search. But between sessions, analysis of results took place and triggers for further investigation were found.

Though different information needs were reflected in different search steps, each successive search carried with it the context of the problem that prompted the searches and of the previous searches done to attack the problem. For example, there was a continued focus on the content areas that bounded the problem, such as an industry, market segment, or geographic region. Also, when a particular result suggested a follow-up search, the new search was often framed to include the context of the earlier result. For example, a search for information about the past mergers involving a particular company might trigger a search for information about the partners in those mergers. But the search on each of the partners is not comprehensive; instead, it focuses on data that relate to the partner’s deal-making capacity and history. Thus there was a thread of continuity through interconnected searches, even when they wandered among different topics.

It was the accumulation of search results, not the final search result set, that had value for most of our library clients. When people finished searching, they often created summaries of the material they had found, including both overviews and detailed views and analyses. In many cases, they also presented results to the interested parties in their work settings.

This pattern of interconnected searches has also been observed by Bates [1], who described the process as “berrypicking,” analogous to finding berries scattered in bunches, rather than all in one spot. She too points out that information needs change throughout a series of searches and that searchers need an accumulation of results rather than a single target result. But Bates’ emphasis is on how these observations affect search techniques, assuming no intermediaries. For example, people may shift information sources as they move from one set of results to another, or they may browse in a source to find a promising new area to explore. Our library clients used librarians throughout, and our current interest is in information usage patterns and interfaces for effective information delivery rather than search techniques. To continue the berry analogy, our library clients seemed to be making a berry dessert (sometimes with surprise ingredients), rather than simply picking and eating berries. How did they decide which ingredient to hunt for next and how did they know when they were finished?

3.2 Triggers

To learn more about the pattern of interconnected searches, we examined the different types of transitions between one step and the next. The data do not allow us to provide accurate frequency counts for these triggers and stop conditions, since people did not consistently reveal the triggers that occurred in their search scenarios. However, we did collect all the triggers they mentioned, and the following four categories characterize almost all of these triggers:

1. The next step fit with the overall *plan*.
2. Something *interesting* arose and prompted exploration.
3. There was some *change* to be explained.
4. There was something *missing* from the data.

3.2.1 Pursuing a plan

Planned search steps do not need much elaboration. The example of the benchmarking specialist given above can serve as a prototype. Planned searches were predictable in their overall structure from the beginning, and the type of results desired was also anticipated (e.g., a list of top companies). In spite of their predictability, planned searches were still carried out in a step-wise fashion, since analysis of one step was needed to produce parameters for the next step. Though the outline of a planned search was clear from the beginning, sometimes one or more of the other triggers occurred and expanded its scope.

3.2.2 Encountering something interesting

People were alert to any search results they considered interesting or unusual, and they typically wanted to follow up on these results. The library clients usually could not define precisely what “interesting” meant to them, but they felt confident that they would know it when they saw it. John, the merger expert, described his exploratory searching this way:

John: Well, 90% of what you get isn’t useful. I mean, you are crawling through the sand to find the nuggets ... you have to read through the one hundred and fifty citations to find the five articles that you are really interested in. This is especially true when you don’t know exactly what you are interested in ...

Sometimes the search results gave the client a new perspective on material that was already known. For example, the college finance officer examined a report of various financial data on colleges and found a number representing a college’s total allocation for student living. This was a way of packaging data that he had not thought of before, and it caused him to generate the same number for his own and other institutions.

The search results often provided a new angle or twist to the overall investigation, which was considered desirable by the library client. For example, the research assistant was given an assignment to investigate and write a short piece on environmental regulation. She described her search this way:

Linda: Trying to build a story, you have to have the foundation and then the body and try to find an angle that makes sense and is interesting. As you go along you see where pieces are missing and fill those in until you feel that you are close enough ...I knew those fundamental questions and could find those answers easy enough. Interesting angles ... There were some unexpected twists of the regulation ... it had something to do with logging and Germany and sort of international trade friction based on these things.

3.2.3 Explaining change

Learning about a change in the environment can, of course, be seen as a version of encountering something interesting. But the situations in which the library clients did follow-up searches based on a perception of change had a different feel than the above examples of following an interesting new direction. The motivation for the follow-up searches was to solve a problem, rather than to track down an interesting new idea. Also, the data involved in change situations were much more likely to be quantitative, rather than qualitative, perhaps because changes in quantitative data are easier to detect. This does not imply, of course, that the needed explanations for the changes were quantitative in nature; the contrary was usually true.

The library clients who were in marketing or who consulted to marketing provided many of the examples in this category. The marketing engineer regularly followed revenue increases and decreases and searched for information that would help explain them. The statistician monitored certain variables related to product sales and performance, and had to find an explanation when she found that three variables she tracked had begun to correlate, instead of the two that had correlated in the past. In an example involving qualitative information, the planning staff member who did weekly searches on competitors found a reference to a new investor in the industry segment he followed, so he did a search to find out everything he could about this investor and the possible implications of his interest.

3.2.4 Finding missing pieces

Some of the triggers for further search and analysis were due to missing information. For example, a client was trying to find trends in the data returned by a search and discovered that certain numeric data were reported inconsistently over time. Usually there had been changes in the way the numbers were packaged. This explanation was typically found by a phone call to the information provider, rather than through an online search.

Search results sometimes stimulated new connections for the library client when they were combined with other information that was already known, leading to a need for more information to support deeper analysis. The management consultant encountered this when he related search results to his own client-provided information and his general knowledge of the industry involved: "When I fit this with what I already know, it begs a question."

Missing information also triggered further analysis or search when the library clients had clients of their own to whom search results or summaries and analyses were delivered. Sometimes these second-level clients asked for new perspectives on the information that had been delivered, perhaps a regrouping or a different slice through the data. Depending on the structure and granularity of the original information gathered by our library client, such a request might trigger new searches. Sometimes second-level clients asked for an analogous result on a related topic, such as another competitor company. This always required follow-up searches.

3.3 Stop conditions

The stop conditions for interconnected searches did not fall into clear, distinct categories in the same way that triggers did. Instead, we gathered some general impressions of the circumstances under which searches were wound down. People usually stopped searching when there were no more compelling triggers for further search or when they felt that they had done an appropriate amount of searching for the task. In a few cases, there was a specific inhibiting factor to further exploration (e.g., learning that a particular market was too small to be worth examining further).

Though our pool of informants had many different professions, they shared some common values and practices through their training and experience in business affairs. Their searches were usually prompted by decisions that had to be made by themselves or others. There was no expectation that a complete collection of relevant information to support a particular decision could be defined, much less retrieved. Alan, a planning staff member, related a lesson from business school:

Alan: And you know the rule of thumb that you can get 80% of the information fairly quickly and the last 10% will kill you. So the point that they were trying to tell was, is it really necessary to get that last 10% of information? ... And the overwhelming answer is, "no way".

Our library clients were very comfortable with making inferences; missing information is a fact of life in a world with closely-held business secrets. These clients were very satisfied with the work done by their librarians; they had faith that if not much was returned by a search, it was because not much was out there. Timeliness and brevity were strongly-held values, referred to again and again by our library clients, and these attributes were considered more important than trying to cover every angle. While it was important to achieve closure in the cumulative search results, this closure had a very pragmatic interpretation. For example, some searches stopped when the client was finally able to make sense out of numbers of interest. Others ended when a particularly evocative visualization emerged from an analysis of the latest data.

This characterization of a successful search interaction is strikingly different from the precision and recall measures that have often been used to measure a system's success in performing an information search. We found no such exactness in determining whether a match between an information need and a set of results had been achieved. Nor do task performance measures seem to apply easily in this domain; the problems are loosely-defined and often non-routine, and the outcomes have many contributing factors other than information search performance. Thus the clearest measure of system success seems to be client satisfaction with both the search process and the information delivery process.

3.4 Librarians with domain expertise

In two of the work settings of our informants, a restructuring of work had occurred that changed the nature of the interconnected search patterns experienced by clients in these settings. Librarians acquired domain expertise, either through formal training (at the management consulting firm) or experience (at the demographer's independent research organization). The librarians were then expected to do some or all of the first-level analysis of the search results and to produce an overview or other high-level output describing the search results.

The motivation for this approach was clearly given at the management consultant firm. The consultants had recognized the pattern of their interconnected searches, and they felt pressured by the time involved in communicating back and forth with the librarian to bring about the next step. Here is how the management consultant described the change in work responsibilities:

Bob: What we found happen is when they went to go through it they would go, "But there are two missing pieces here. Go back and dig deeper to find those." ...What we found is when we put the information specialist in the position of having to do that first level synthesis, that loop gets closed quicker. They also generate less raw material in the search, because they become more efficient searchers because they understand at least the first level intermediate product.

Professional searchers at this firm are given a short course in basic finance, and they are also expected to learn about particular industries and types of business analysis from the searches they do. When they get a search request, the basic goal of the search is explained. The librarians then do the search and analyze their results. They produce a package with an overview page, a set of slides and graphs with bulleted main points and views of quantitative data, and an annotated bibliography of the source material. As background for meeting preparation or an initial view of a topic, this may be enough. If a more complex synthesis is needed, a member of the management consulting staff does it based on these first-level results.

While this is considered a successful work methodology at the consulting firm, further improvements are being implemented on an area by area basis. When an information need comes up, an associate of the firm and an information specialist are teamed together to solve the need. This has begun to produce good results in the areas where it has been tried.

Bob: Because we don't think we can get the full set of business analytical skills and the full set of information search skills in the same person. We just don't think that that's realistic. But the thought is that if you create a team that has both sets of skills and the request goes to the team, not the individual, so there is not a sense of a hand-off from the information researcher to the analyst but they are in this together ...the sense of ownership is probably a critical thing that changes. Plus our sense is that there are two other things that will happen ... they will progressively disclose work in progress which will be shaped by feedback from the analyst who says, "This is helpful, this wasn't. More of this, less of that" in a much shorter interval. And the second thing is our guess that there will be a learning cycle that will take place and that the analyst will learn more about the mechanics of the search and be able to direct the search more precisely and the information specialist will learn more about the mechanics of analysis and begin to develop a better sense of what is useful and what's not ... We have learned some lessons about training, supervision, structure, etc. So, we are going to implement those and try it in [another work area].

Even when work was not restructured as dramatically as it was in the management consultant firm, some librarians and clients felt that their shared understanding improved the effectiveness of their interactions. Clients tended to latch on to a favorite librarian and return to her, even when any member of the library staff was capable of conducting a particular search. Trust was built; one client said that when he first started doing searches he sat at the librarian's elbow to watch her work and check her results, but later he just left a phone message outlining what he needed. Clients felt that their regular librarians understood their technical terms better than the librarians they used less frequently. In some cases, the librarian had developed a good understanding of the kind of work the client did (such as in the case of the process benchmarking specialist), and both the client and librarian felt that

this helped to target the searches.

When librarians are involved in analysis, a higher-level information product is delivered to the client. This can affect the course of the progressive search path. In the case of planned searches, at least some steps of the plan will have already been carried out. There should be no obviously missing pieces from the result set; if there are, there is a quality problem with the information, according to the management consultant. Some unexpected angles will have been followed up, but others will not have been recognized, since these are often based on a thorough understanding of the problem or domain. The stop conditions will not change in their basic nature, but they will often be encountered sooner. The overall effect is to shorten the path to the goal.

3.5 Analysis techniques

We have described the different types of searches done by our library clients, the basic pattern of interconnected searches that most of them followed, and the triggers and stop conditions that led from one search step to the next. Now we will describe how clients went about assimilating and manipulating the data from a single result set in order to make progress in their tasks. As with triggers, we extracted the analysis examples from the transcripts and categorized them, but there was some ambiguity in attaching analyses to search scenarios and in determining which analyses were generic and which were specific. We did find a larger set of analysis examples and a larger number of analysis categories than we had found for triggers.

Initially, library clients read through the pile of material returned from the search. In most cases, they read paper copies, even when the material had been delivered electronically. A few clients scanned the information, but most said they read it cover to cover so as not to miss a valuable nugget. Most clients annotated the results while they read, to indicate what was interesting or questionable in the result set and to identify likely articles to be retrieved in full or topics to be targeted in follow-up searches.

In addition to reading and annotating, our library clients primarily did six types of analysis:

1. Looking for *trends* or *correlations*.
2. Making *comparisons* of different pieces of the data set.
3. Experimenting with different *aggregates* and/or *scaling*.
4. *Identifying a critical subset* of relevant or unique items.
5. *Making assessments*.
6. *Interpreting* data to find meaning in terms of domain or problem concepts.

Though what exactly counts as an analysis example is open to interpretation, roughly 80% of the approximately 80 analysis examples we found fell into one of these six categories. The remaining analysis techniques included cross-referencing, summarizing (a constructive technique related to several of those listed above), finding evocative visualizations, and a half-dozen others which were only encountered once each.

Trends are a very basic way of thinking about the business environment, and they provide fundamental support for decision-making. They offer a means of both tracking and predicting changes in the environment, a way of viewing the landscape of both qualitative and quantitative data. The need to find trends influenced both the data that were sought (numbers, time series) and the manipulations that resulted. Clients frequently put numeric data into spreadsheets, computed interesting numbers, and then generated a variety of graphs, tables, and other visualizations. As they looked through the visualizations, they tried to find patterns. Clients who did monitoring-type searches usually mapped trends over the long time periods of their searches. Forecasting was a common activity similar to finding trends.

Comparisons were also pervasive. A competitive environment leads to comparisons among companies, deals, processes, and more. The clients' descriptions of their high-level tasks often included comparisons. The data to be compared were often qualitative in nature, so creating the right table or graph was not enough. The data had to be summarized and salient points highlighted. The comparisons usually appeared in new artifacts created by the client, rather than in marked-up copies of the result set.

Aggregation and *scaling* were often done in conjunction with comparisons or trends. For example, a trend might only become apparent when certain industry segments were lumped together or scaled to a particular level of detail. Different buying behaviors might make sense only when customers were clustered by certain characteristics. Aggregation is a means of realizing abstractions, and we believe that along with comparison, this analysis technique is often used beyond the business domain. Again, the library clients experimented with aggregates in the artifacts they created using their search results, such as spreadsheets, graphs, and structured textual summaries.

Identifying a critical subset involved selecting from the result set those elements that met some domain criteria. Sometimes the subset simply represented a collection of interesting data points, as in the example of the sales/planning staff member choosing which paragraphs from his search results should go into the weekly newsletter he produced. Sometimes the subset became the starting point for a new search; for example, the merger and acquisition specialist chose from a set of material on potentially comparable mergers those mergers that were the *most* comparable, and then did further search and analysis on this smaller set. Sometimes a library client formed different subsets in an exploratory manner, to see which provided the most meaningful basis for a comparison operation.

Assessing was a common operation that required clients to dive into the data and use their domain knowledge to come up with meaningful conclusions. Some example questions that the library clients tried to answer as they analyzed their search results were: Is this market declining? Will this company be around in a few years? Is this product a match to this customer? What is contradictory in this information? People sometimes used domain-based analysis techniques such as financial ratio analysis to answer questions like these, but mostly they seemed to use their general experience to come up with reasonable judgements.

Interpreting is similar to making assessments, but it seems to be a more subjective type of analysis. As library clients interpreted their data, they assigned new meaning to raw information. For example, the merger and acquisition specialist looked at search results on previous deals to determine the "psychology" of the company in question: What were their problems and how did they usually go about solving them? While assessing and interpreting, library clients needed to take both broad and detailed looks at their search results.

4 Discussion

We have described how each of the different types of searches conducted by our library clients tended to take place in separate search steps over a period of days, weeks, or months. What are the system design implications of this pattern?

The DIME retrieval mediator toolkit and runtime environment that we are building must take into account the context provided by interconnected searches. There are clear interconnections between successive mediator invocations, ad hoc queries, and requests to intermediaries. Mediators might be particularly useful for representing monitoring and planned searches, since these are repeated over time. However, even monitoring searches can lead to exploratory forays, and planned searches often require analysis operations before the next step in the plan can be undertaken. Thus a mediator must allow different kinds of intervention by the user rather than simply run from beginning to end without interruption.

DIME should provide assistance to users in finding and invoking related mediators when the relationship can be anticipated, as it can in the planned search cases. In a series of mediator invocations or ad hoc queries, we would expect to see some common information sources and common terminology used to express the boundaries of the problem area, so carrying these parameters over from one search to the next should be made easy for users. If query reformulation techniques ([6, 10]) could be extended beyond individual items and terms to include some of the surrounding context of the entire previous search, this would provide significant support to interconnected searching. This might be as simple as applying a filter with context-specific terms (e.g., industry segments or geographical areas) to a similarity-based retrieval set.

A record of an entire interconnected search thread, comprised of both requests and results, should be saved by the system in such a way that it can be deactivated (stored persistently) and activated as the search dies down and then picks up again. When a search thread is reactivated, the user should be able to browse and build on the requests, results, and planned steps of the search thus far. Many library clients had more than one search thread in progress at once, so each thread should have an independent existence. Such a record could also be used as a template to create a new mediator or set of mediators, if the search pattern should turn out to be one that recurs over time.

There are also clear interconnections within the picture that grows from successive search results. The search context and the direction of the search are closely tied to the particular triggers a user finds for further analysis and search. One way of linking search results together would be to allow users to annotate a result set with notes indicating triggers, patterns, and conclusions that are suggested by the data. This would offer the user a kind of interactive search-planning tool, which might evolve into an overview description of the search results. Similarly, as new descriptive artifacts are created using the search results as raw material, coarse-grained links back to the raw material could provide either a bibliography or a form of rationale. For example, an outlining tool with the ability to link back to segments of the search results would be useful.

The analysis techniques used by the library clients could not be automated; their use was based on the clients' understanding of the domain and the applicability of the information to the problem at hand. In some cases, analyses depended on application tools such as spreadsheets and graphics packages. However, we believe that certain generic operations to manipulate information would be useful to customers of intermediaries and mediators. The

ability to perform these generic operations on application data such as spreadsheets would be valuable, but the operations should at least work on the raw search results.

To support aggregation and comparisons, users should be able to extract information chunks, label them, move them around, and create arbitrary groupings out of them. Careful juxtaposition of information is often needed to allow side-by-side examination, expose relationships, and prepare for deeper looks at the data. As with the annotation capability suggested above, these editing operations might also help the user create descriptive artifacts along the way. The processes of identifying a critical subset, making an assessment, and making an evaluation are not well understood by ourselves or others, and more research is needed into the processes themselves before we can predict the kind of interfaces that would best support them. However, the ability to manipulate chunks of information, as suggested above, would probably make the extraction of a subset easier. The ability to create different overlays of thematically-linked user annotations would support all three of these processes.

Though there are many search situations in which automated support for querying is practical and appropriate, we do not mean to suggest that expert librarians or other intermediaries can be replaced. Retrieval mediators can be seen as an extreme example of automated support; the user does not express a query at all, but instead invokes a mediator with suitable parameters. We believe that this is a promising direction in information access research, and in this study we saw examples of searches for which mediators could be created. But mediators are only suitable for repeated queries with a predictable pattern and manageable scope. They would not make sense for unusual queries or browsing situations, for instance.

For the situations in which librarians are the best solution, our findings suggest an alternative working style for clients and librarians that involves a closer partnership than the traditional service-oriented model. With shared goals and frequent communication, partners in some settings can accomplish search tasks faster and more effectively than when each has a completely separate role. The success of this approach depends on the librarian acquiring some domain expertise, which is easier in domains such as business than in highly technical areas. But even in technical domains, librarians have shown an aptitude for picking up terminology, research themes, names of relevant companies and individuals, and the like. We suggest that developing partnerships that use the strengths of both the user (deep domain expertise) and the librarian (deep search expertise) makes sense as a complementary approach to the current trend of empowering end users to conduct their own information searches. These partnerships would benefit from certain collaboration technologies, such as sharable document annotations and remote communication support.

In summary, we have seen that information searches have structure and continuity, even when they are very exploratory in nature. As designers of search systems, we must consider how a user's interaction with the system can be guided by this structure, without losing the flexibility to perform expert analyses of search results and to follow interesting new search directions.

5 Acknowledgements

We are grateful to the library clients and librarians who spent time with us describing their jobs and their library searches. We also thank Lucy Berlin, Bonnie Nardi, and Andreas Paepcke for useful comments on earlier drafts of this paper.

6 References

- [1] M. Bates. The design of browsing and berrypicking techniques for the online search interface. *Online Review*, 13(5), 1989.
- [2] N. Belkin. Ask for information retrieval: Part 1. background and theory. *Journal of Documentation*, 38(2), 1982.
- [3] N. Belkin. Ask for information retrieval: Part 2. results of a design study. *Journal of Documentation*, 38(3), 1982.
- [4] B. Dervin and M. Nilan. Information needs and uses. In M. Williams, editor, *Annual Review of Information Science and Technology: Volume 21*. Knowledge Industry Publications, Inc., 1986.
- [5] D. Ellis. A behavioural approach to information retrieval system design. *Journal of Documentation*, 45(3), 1989.
- [6] G. Fischer and H. Nieper-Lemke. Helgon: Extending the retrieval by reformulation paradigm. In *Proceedings CHI '89*, 30 April-4 May 1989. Austin.
- [7] E. Hewins. Information need and use studies. In M. Williams, editor, *Annual Review of Information Science and Technology: Volume 25*. Elsevier Science Publishers, 1990.
- [8] A. Paepcke. An object-oriented view onto public, heterogeneous text databases. Technical Report HPL-92-84, Hewlett-Packard Laboratories, 1992.
- [9] G. Wiederhold. Mediators in the architecture of future information systems. *IEEE Computer*, March 1992.
- [10] M. Williams. What makes rabbit run? *International Journal of Man-Machine Studies*, 21, 1984.