

## **Information Artisans: Patterns of Result Sharing by Information Searchers**

Vicki L. O'Day, Robin Jeffries  
Systems Technology Laboratory  
HPL-93-19  
February, 1993

information search,  
information use,  
information  
artifacts,  
intermediaries,  
collaborative work

We studied the uses of information search results by regular clients of professional intermediaries. We found that all of the participants in our study acted as intermediaries themselves, sharing information they had received from library searches with others in their work settings. There were four basic models of sharing: updating team members, consulting, broadcasting, and putting information into a shared archive. In many sharing scenarios, the library clients acted as information artisans, creating new artifacts by transforming and enhancing their search results before passing them on. When possible, the library clients delivered their new information artifacts in collaborative settings, to ensure that recipients understood and could apply the results and to allow opportunities for follow-up search requests. These observations suggest that new functionality is needed for information search systems, in order to support the analysis, manipulation, and packaging of search results and collaborative information delivery with intertwined communication and information components.



# 1 Introduction

Many designers of information retrieval systems have focused on improving interaction with search systems by making it easier for users to express and refine their search requests, so the search results will come as close as possible to matching the real information need. In our work, we take a complementary approach. We are particularly interested in the effective delivery of search results, to enable users to analyze and manipulate search results as needed by their particular problem situations. To learn more about information delivery requirements, we conducted an ethnographic study of how library clients deal with the information they get back from human intermediaries.

We found that *all* of the informants in our study turned out to be intermediaries themselves, though they were selected for the study simply because they were regular library clients. These clients initiated a variety of searches, usually to inform decisions. Even when they had the authority to make decisions on their own, the work practices of their settings led them to share their search results with others, either by request or on their own initiative.

Acting as intermediaries required these clients to move beyond the problems of dealing with information to focus on communication issues. The need to communicate what they had learned required dedicated attention to the composition and packaging of the information and the setting in which it was delivered. Sometimes search results were passed on unchanged, but often they were refined or transformed in some way by the library clients. These transformations were creative efforts, distinct from the work of analyzing the search results for the problem at hand. New information artifacts were synthesized, based on the raw material obtained from the library.

The practice of communicating search results, often by creating new information artifacts, has implications for the design of search systems and, more generally, for the design of many tools through which information is manipulated. Effective information delivery mechanisms must not only allow the searcher to easily read, analyze, and preserve search results and conduct follow-up searches, but must also provide flexible ways of sharing both raw search results and new artifacts derived from them. Our purpose in this paper is to outline the set of common information-sharing practices of our informants and to illustrate the need for more effective support for information delivery. We will describe the models of sharing that the library clients used, the issues involved in creating new artifacts out of search results, and the practice of delivering the new information in collaborative settings.

## 2 Related Work

Most work in information retrieval is concerned with improving the efficiency of the searcher's session with the information system, rather than with the course the information takes once it is in the hands of the searcher. An exception to this approach is provided by Taylor [12], who uses a "value-added" model to analyze how and under what conditions information is transferred and used. In his model, all information processing systems, technological and human, add value to the information they deliver. The basic question is how these systems can be designed and used so the particular values that are added match the needs of the decision-makers who consume the end product. Examples of values Taylor proposes include *reducing noise*, *mediation* between difficult-to-use information tools and the consumer, and *adapting* information to bring it closer to the problem. Our study verifies Taylor's model in that the library clients added exactly these kinds of values when they passed search results

on to others, especially when they created new artifacts such as summaries.

We have focused our attention on the creation and use of information artifacts. In general, by examining artifacts and their uses, we can begin to understand system design requirements [1, 6, 8]. In the language of activity theory, Bannon and Bødker [1] note that artifacts *mediate* the activities that people engage in. (See [6] for an overview of the principles of activity theory.) That is, artifacts govern to some extent the paths that people may take as they engage in activities to accomplish their goals. Through their intrinsic properties and historical patterns of use, artifacts provide ways of sharing and dividing work. Bannon and Bødker stress that it is important to take the *use process*, not the artifact itself, as the object of study, since an artifact outside of its context of use is not meaningful. Our study is an attempt to do this for a particular set of artifacts, library search results and the documents that are derived from them.

Norman describes a *cognitive artifact* as “an artificial device designed to maintain, display, or operate upon information in order to serve a representational function” ([8], page 17), and points out that such artifacts change the basic nature of the tasks for which they are used. In addition to changing the task for the artifact user, a design and construction task is added for the creator of the new artifact. Norman suggests as a design principle for cognitive artifacts that the “surface representation” of a new artifact should allow the person using it to work with exactly the information needed for the task, no more and no less. The library clients used this principle as they created streamlined, problem-based representations of their search results to use in decision-making processes.

### 3 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 the rest 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; most interviews lasted about an hour.

The interview questions were designed to expose the work contexts and specific tasks that

triggered search scenarios and to explore example search scenarios in detail, from the search request to the analysis and use of search results. 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.

In analyzing the interview transcripts, we saw two themes or patterns that appeared consistently throughout our informants' search experiences. One pattern was that of search clients as intermediaries, which we will describe in this paper. The other pattern was that of interconnected searches, described in [9]. We will briefly summarize the second pattern here to provide context for the discussion of intermediary activities.

People usually conducted a series of interconnected but diverse searches on a single, problem-based theme, rather than just one search session per task. In a given search scenario, a library client operated in one of three different *search modes*: monitoring, planned, or exploratory. In monitoring searches, the library client chose a set of variables (such as competitor companies or specific technologies of interest) and at regular intervals requested a search on these variables. The frequency of monitoring searches varied from once a week to once a year. In planned searches, the client executed a multi-step information search plan suggested by a typical approach to the task at hand (such as researching a company that might be a candidate for a merger). Typically, the client knew what *kind* of information he or she wanted back from each search step (e.g., basic financials on the merger candidate or a list of comparable merger transactions). In exploratory searches, the library client explored a topic in an undirected fashion, using the interesting aspects of each set of search results to suggest a new direction for the next step.

Depending on the purpose and mode of a search, the kind of results obtained, and the working relationships of the library client, there were a variety of ways in which search results were shared. We will describe the most typical approaches to sharing results in the next section.

## 4 Models of Sharing Search Results

When we say that all of the library clients we interviewed were intermediaries themselves, we are using the word intermediary to include different kinds of information-sharing activities. We do not only think of an intermediary as someone who receives explicit information requests, initiates online searches, and passes back the results that were supplied by the online system, though some intermediaries did follow this pattern. We saw four basic models of sharing of search results among the library clients, which are described below. Since most library clients (all but two) did more than one kind of sharing, the numbers given for how many people fit each model add up to more than the total number of informants in the study.

1. Sharing results with other members of a *team*. [9 people]

Library clients who shared with their team members seemed to define their teams as groups of people who were working towards a common goal and who regularly communicated about their individual and group progress. A person might belong to several ad hoc teams, each created to solve a particular problem. The work of a team was divided among the members, but the team met to make plans and draw conclusions together.

In a team, people were usually considered local experts on the topics they had searched. They passed on to other team members detailed selections from the search results or high-level summaries, and answered follow-up questions when necessary. The merger and acquisition specialist described his role as an information broker:

**John:** ...to a certain extent what you are doing, you are getting the rest of the team up to speed on what things are going on. Early in the project, people are looking for superficial information. How big is [a candidate company for a merger]? How big is their business? Where is it located? What are they doing? What are their products? So it's really sharing the information with other people.

**Interviewer:** And then later, the more nitty-gritty, as-needed stuff?

**John:** That's right. They look to us as the repository of the information, you know. I mean, they are not retaining it. It's, "That's nice, what's your phone number?"

2. Self-initiated *broadcast* of interesting information encountered in search results. [8 people]

When library clients encountered unexpected but interesting tidbits in their search results, they circulated these tidbits to people across the organization who might be expected to find them useful. This kind of routing relied on the client's understanding of other people's current responsibilities and interests. Ad hoc broadcasts usually involved excerpting the item of interest and passing it on verbatim.

Some people had formalized the information provider role and offered regular (e.g., weekly or quarterly) news or analysis updates to others. This occurred particularly with the results of monitoring searches. The search results were put into a standard format chosen for that particular search scenario, such as a spreadsheet for quarterly financial results or a newsletter for the latest news on emerging technologies and competitors' activities. People who acted as regular sources of broadcast information added their own annotations and structure to the material they forwarded.

3. Acting as a *consultant*, handling repeated or ad hoc information search requests made by others. [11 people]

Many clients regularly played the role of consultant, conducting information searches in response to a request from someone else and then passing the results back to the requester. Sometimes the consulting relationship came about because of the library client's position within the organization, as in the case of a member of an umbrella marketing group which was a resource for various product development projects. Sometimes consulting requests came about because of the client's special area of expertise, such as statistical analysis or finance. Several library clients had a small number of distinguished clients of their own from whom they frequently received information requests. In each case, these distinguished clients were high-level managers who relied on others to conduct searches and analyze data on their behalf.

In two cases, people were consulted because they became known as people who were particularly effective at finding information. Raymond was a financial analyst who fit into this category:

**Raymond:** What I have been doing is working with [the librarian] and trying to sort of be the contact through to [the library], where rather than have a bunch of divisions call [the library], they could maybe call me and work with one person. ...since a lot of these messages have gone out with my name on it then I have been getting a lot of calls from divisions.

Some consultants passed results on unchanged; others did extensive analysis and manipulation of the information before passing it on. The people who were more likely to transform search results before communicating them were those who were more senior (typically more experienced) and those for whom consulting was a formal part of the job, as opposed to a sideline or volunteer activity.

In most cases, library clients passed on search results to others within their own organizations. But four clients also passed search results to outside interested parties. For example, the management consulting firm passed selected search results on to its clients, and the venture capital firm used search results to prepare material on markets and individual start-up companies for the investors from whom the firm attracted capital.

Though these sharing activities fit into the consulting category, we distinguish them because a higher standard of information refinement and packaging was demanded when the information went to outsiders. This does not mean that all raw search results were transformed before being included in an external report or presentation, just that extra care was taken to ensure that the right material was included and the right level of analyses were done.

4. *Archiving* potentially useful information into a group repository. [3 people]

Several library clients who were involved in competitive analysis belonged to groups that maintained an archive of current information on competitors. Search results in this area were routinely shared through the archive. In one case, the archive consisted of an on-line “slide show” that walked the user through the latest information on key competitors; in the other cases, the archive was a paper-based file of printed-out search results, clippings, and other material.

There were a few situations in which search results were not shared in any way, even in an abstracted form. This occurred primarily when the searcher was the only decision-maker who would be using the information. In this case, the results were retained as rationale for a decision, but not explicitly communicated to others. This situation was unusual among the library clients we talked to, when the search results were successful. (Sometimes searches led to dead ends or the library client didn’t have time to do needed follow-ups; in these cases, the material stayed with the client.) Only three people gave examples of successfully-completed search scenarios whose results were not passed on somehow. Even people in high-level positions belonged to management teams or had outside clients with whom they discussed decisions and therefore shared results.

These different kinds of information sharing required different amounts of contextual information to be associated with the search results to make the results useful to the recipients. Ad hoc broadcasts of search result excerpts required very little context; these items were self-contained, and their applicability was expected to be obvious to the recipient. Data acquired in monitoring searches were placed into a structure that was similar from one update to the next, such as a spreadsheet model that was simply updated by the addition of a new column when quarterly financial results were obtained.

The sharing of search results that took place in consulting and team scenarios was driven by the problem that had motivated the search itself. Consulting brought the context of a question already asked; team sharing brought the context of a group problem on the table. In these cases, more complex analyses took place and a higher-level information product was

created by the library client. Raw results were embedded in slides, tables, and the like. For example, one of the financial analysts used financial data, excerpts from annual reports, and quotations from press releases to create table and text-based summaries of the positions of competitor companies.

In some sharing scenarios, the library clients simply passed on material exactly as it had come to them from their librarians, with little or no additional packaging. But in most cases, the clients became “information artisans” who crafted and delivered new information products. These products were derived from the results of their own searches and were developed with attention to the needs and interests of the recipients. In the next section, we will describe some of the methods and issues involved in creating new artifacts out of raw search results.

## 5 Information Artisans

The raw material on which new artifacts were based was similar from one search example to another. It was either purely textual (e.g., articles or press releases) or a combination of numeric and textual (e.g., financial statements). The online sources searched by the librarians did not include graphical images, for instance, though the online sources sometimes pointed to print versions which contained graphics.

From this material, people created a surprising variety of new representations. These included annotated excerpts, summaries, reports, spreadsheets, charts, graphs, slides, visualizations of three-dimensional data, special-purpose computations or indices, and combinations of all of these. Data from other sources were sometimes merged in with search results. Some people used sophisticated tools to construct these artifacts (such as when the statistician experimented with different 3D visualizations using her graphics package), but others relied on manual techniques such as photocopying, retyping, and editing, as search results moved back and forth between paper and electronic copies.

Above all, these new representations offered problem-based structuring and abbreviation of the original search results. The merger and acquisition expert described his presentations of search results this way:

**John:** Usually you are abstracting it, you are not passing it on verbatim to other people so it’s not sort of cut and paste kind of stuff. You’re abstracting it... People want to know what you learned, not what you know.

On the other hand, several library clients pointed out that it was important to make the original source material available. Carol, a sales/planning staff member, described some of her intermediary activities:

**Carol:** I’m probably the main gatherer, and then after that I will act as sort of like summarizer and write something up giving a more detailed description about what I found and then reference people to the information. If it is important, then I will just file the whole thing, the pack, all the information I have, along with this letter and put it in our files there for future reference. If it is not as important as I thought I will reference them to, “Well, I got them from this source, this source, and this source,” and those sources are not going to go away.

The same search result material can be structured so it is useful for different purposes. In the management consulting firm, the librarians did the first level of analysis and produced a set of slides and graphs with high-level points. (See [9] for a description of this unusual



work arrangement.) In this setting, the search results passed through the hands of a chain of information artisans, each of whom made a contribution to the information's usability. Here is a description by the management consultant:

**Bob:** Generally [the package] has a one or two-page summary on the top that says, here is the main points. Then a set of slides/graphs, if you will, of some of the basic quantitative data, and then the last thing attached to it is sort of an annotated bibliography that does, okay, if you want to go beyond the main points in the top-level data of the various articles and internal [company] publications we found, here is the most relevant. And here is kind of the major message in each one of those and then you can pick and choose... For a more complex synthesis of it, which happens sometimes in preparation for a meeting, we would get one of our associates, typically someone more senior, who would put that together with more of a narrative flow, more of a story with some conclusions around it. That's almost always the case with something that we are going to share with the client.

A flexible, integrated set of tools is needed to support the variety of formats and methods people use to construct new artifacts out of material they have received. It should be possible to organize, annotate, extract, and merge search results. We have begun to think about information constructs we call "information collages," which are custom arrangements of selected information chunks. Complex collages might include multiple views of the same basic information (such as an overview and a detailed view) and scripts to guide readers through the elements of the collage. Prototype or template collages could be created for often-reused information structures, such as the tables created for the results of certain monitoring searches or the management consultant organization's standard presentation format described above.

Information delivered from the library should be accessible from within the tools people typically use for analysis. In most of the search scenarios we saw, the library clients had to retype search results when they wanted to include them in spreadsheets, statistical packages, graphics packages, and even electronic mail tools. Part of the problem was that the tools lacked an import capability. Another part of the problem was that people often chose to have library search results delivered on paper (e.g., for faster delivery, since fax is faster than electronic mail in some settings) or they moved the results to paper right away for easier reading and annotating of the material. People also need the ability to export the artifacts they have created using a variety of tools into the various distribution channels they use: electronic mail, small-group discussions, live presentations, and of course, paper-based mail.

As information is passed from the library to the library client to downstream clients, it should carry with it pointers to sources and pointers to raw data that may have been filtered out along the way, since later artisans may want to refer back or include parts of the original material. People do sometimes need to follow up as they do deeper analyses or related searches; providing more complete information would make this easier.

## 5.1 Sharing Analysis Results

In [9], we described the six most frequent types of analysis the library clients did on their search results. These six analysis types account for roughly 80% of the approximately 80 analysis examples we found. The first three analyses involved direct and detailed manipulations of the raw search results: (1) looking for *trends* or *correlations*, (2) making *comparisons* of different pieces of the data set, and (3) experimenting with different *aggregates* and/or *scaling*. The last three analyses were perhaps more demanding cognitive activities, and required library clients to examine their data from both broad and detailed perspectives: (4) iden-

tifying a *critical subset* of relevant or unique items, (5) making *assessments* or judgements, and (6) *interpreting* data to find meaning in terms of domain or problem concepts.

Each type of analysis suggests different presentation styles for communicating the results to others. Trends and correlations were usually shown with graphs. Comparisons were suggested by juxtaposition of tables, graphs, or textual descriptions. For example, one of the sales/planning staff members produced a weekly newsletter composed of press releases, analyst reports, and other pieces of text. He arranged the text items in clusters, using juxtaposition and annotation to highlight the similarities and differences among different reports on the same topic.

Library clients experimented with different aggregates and scaling factors (for example, clustering customer data or industry segments by certain characteristics) to see where the best correlations could be found. In the artifacts they created, some clients chose to expose the different tacks they had taken during analysis, rather than only the one that had produced the “best” results. For example, a financial analyst created a document out of a series of graphs, each of which aggregated the raw data in a different way. Many of the graphs showed no discernible correlations or trends, but those that did stood out clearly from the rest.

The interpretations and conclusions produced by the last three types of analysis were not as closely related to particular types of presentations as the results of the first three types of analyses. In some domains, special-purpose information structures were created to show conclusions. For example, a sales/planning staff member used a four-quadrant “SWOT” chart (strengths, weaknesses, opportunities, and threats) for organizing competitive analysis data gathered in his searches. The venture capitalist created a “burn-o-meter” graph from the raw data on his start-up companies, to show when each company would need a new influx of capital. In other domains, however, the insights that resulted from complex analyses were represented using more general techniques: written narrative, supporting details, and interactive discussion.

## 5.2 Communication Challenges

New artifacts were designed with their consumers’ knowledge and purposes in mind. The library clients knew the positions, interests, current tasks, and organizational environments of the people with whom they shared search results. They used this knowledge in selecting what to include and what to emphasize in the new artifacts. However, people still found it difficult to convey the essence of what they had learned from their searches. They themselves had read the whole set of search results with all of its accumulated context, and they typically tried to produce something that was much shorter and was also closely tied to the original problem or question. It was hard to do this and still retain the full flavor of the original material. As the planning staff member who produced the weekly newsletter said, “...but you could tell if you had read those articles and the analysts talking about it. It is easier for me to feel it than it is for me to describe it to you”.

Several library clients distinguished between the work of understanding the material on their own and the work they did to create effective artifacts for others. The statistician said about graphs, “Not always do I look for the graph first; it may be the final step for me. It is the means of communication, not the means of analysis”. She described some of the thinking she went through in deciding how to convey information:

**Adrienne:** We have things like all the forecasts, everybody loves forecasts, models, and so you

want representations for models where I would be typically expressing mathematics... But conveying those models to other people, few people like series of equations... if you are talking about all those engineers being mathematical, but nowadays they are more graphic...

In the business domain, timeliness and brevity are strong values, so the library clients tried to condense the information they had gathered into the shortest possible message. The choice of visual representations for data was sometimes related to the need to be succinct, as the research assistant pointed out:

**Linda:** We like to use graphs on reports. It conveys information, it catches people's attention, and it breaks up the text. Layout is very important to us in our reports because we are dealing with people in the corporate world, the strategic planners and other executives who are very busy and we try to present information... the most effective, quickly-readable way possible.

A financial analyst was careful to include in his reports only the kind of information that he was best qualified to obtain and analyze:

**Raymond:** Well I try not to, I guess my purpose is to try to add some information that divisions don't get elsewhere. I mean I am not an engineer and I always get nervous quoting somebody from trade journals... I feel that this is unique information and it goes to higher levels in the organization, so I basically try to make it brief and not repeat a lot of things that a lot of other people are doing.

Communicating information search results to others was a separate effort from collecting and analyzing the results on one's own. It often meant experimenting with different communication mechanisms to find the one that worked best. In many cases, written artifacts were not enough, though when recipients were not colocated the written versions had to suffice. Even when library clients had carefully crafted new representations for their search results, they augmented their new documents with meetings or interactive presentations to make sure that the information was received and understood. In the next section, we will describe some of the variations we saw on this practice.

## 6 Collaborative Delivery of New Artifacts

The delivery of new information artifacts was sometimes done through inter-office or electronic mail, especially when the recipients were scattered in remote locations. But when the recipients were local, the delivery of the information was usually in person and collaborative. These sessions were used to walk through printed copies of the new material that had been created by the library client, check the library client's understanding of the problem, check the recipient's understanding of the results, jointly interpret the results, and provide opportunities for feedback and requests for further searches.

Refinement of search problems took place in consulting scenarios in which the library client who provided the consulting did some analysis before passing the results on. Checking periodically with the downstream client ensured that the search path followed the right direction, and it also allowed the downstream client to understand the results as they grew, rather than only seeing a final version. The statistician described how she interacted with one of her clients in a particular search scenario:

**Adrienne:** Well, first while we were doing some of the refining, I would start and I would take them a chart and say, "Look, this is this, this is where we are, and this is when it ends." And they

go, “Okay, I can handle that.” And when there was something that I knew would be of particular need or interest I handed that to them and [said], “Look what’s happening,” and we would talk about it for a little while. But that wasn’t the whole story. When I finally found out what was happening here, I went and told them because I knew they were concerned and I knew that was relevant but I was also the whole time collecting this whole story which yes, you start, “Let me tell you a story that your data told me...” Some of the people want written documents and I do written documents... and oral back and forth delivery because there is a lot of, “Do you understand it or not understand it.”

In this scenario, the library client used the face-to-face sessions partly to guide the data collection and analysis and partly to make sure that the recipient of the results understood them and could apply them in the original problem domain.

The collaborative sessions described by the library clients always included an information artifact that was used as a reference or a focus of attention. When the meetings involved only two people, as in the case of the statistician above, the two participants looked at a single copy of the document together. In small-group meetings, each participant might have his or her copy of the document, but it was examined in a coordinated way by the whole group. For example, one of the financial analysts was a member of a cross-functional team which met monthly to discuss progress on an ongoing project. The analyst conducted library searches between meetings, analyzed the results and put selected results into a spreadsheet model, then guided the group through the model during the meetings. The model had different elements of interest to different participants; some of the information was targeted at the management representative, some at manufacturing, and some at accounting. But all of the participants were exposed to all of the material, to enable them to collaborate on decisions that would affect all parties. In large groups, shared information was often cast as a slide presentation, with copies of the slides for participants to take away.

A similar pattern of collaborative information use is described by Luff, Heath, and Greatbatch[7], based on their observations of paper- and screen-based information uses in three settings: a medical clinic, an architectural design firm, and the control room of the London Underground. The researchers found that in each setting, people commonly used documents not only for the asynchronous exchange of information, but also to support synchronous, real-time collaboration. Even documents intended for individual use ended up being used by others in the work settings. The researchers suggest that computer-based information-sharing mechanisms might be designed to support colocated as well as distributed work. We believe that such support would also be useful to the information artisans we interviewed.

Our library clients were experts on the data they had collected, since they had read and digested it carefully, but their downstream clients (team members, requesters of consulting services, or recipients of broadcast information) were often more expert in at least some aspects of the problems or decisions to which the information was being applied. Thus the interactive sessions offered opportunities for joint interpretation of the results. Each person had a piece of the puzzle, and discussion and experimentation were sometimes needed to work through the implications of the search results. For example, the statistician brought a collection of eight to ten visualizations of her search result data to a meeting. Together, she and her client examined these visualizations and talked about what they revealed in the problem domain (which in this case was market trends for a particular industry). These sessions offered a learning opportunity for both parties.

The delivery of search results sometimes triggered requests for follow-up searches. These

might be different "slices" on the same material, requiring search and analysis similar to what had gone before. That is, the downstream client would ask the original library client to "do the search again", but to substitute a different set of parameters, such as another company or industry, as the focus of the search. Some follow-up requests were meant to probe more deeply in the same general topic area, since library clients tended to acquire a reputation as experts in areas they had investigated.

Over half of the library clients emphasized the importance of delivering search results collaboratively. Those who delivered information mainly through paper or electronic mail usually did so because the recipients were geographically scattered. Although the interview data are not completely clear on this point, it seemed that two library clients delivered information non-interactively to recipients who were colocated because of the organizational distance between the sender and receiver. A junior person in the organization may be less likely to sit down and hold a discussion with a senior manager than he or she is to send some requested information through electronic mail.

Collaborative delivery was in general preferred by the library clients in our study, but there is clearly a tradeoff between the time and energy consumed by this practice and the ready accessibility of lower-bandwidth information delivery mechanisms such as electronic mail. If electronic delivery were more flexible and expressive, it might be the case that people would actually prefer it to collaborative sessions. More work is needed both in inventing new ways of sharing and exploring information and in evaluating the differences in applicability of asynchronous and synchronous delivery.

It is clearly the case that the technical requirements of the current style of interaction between library clients and the recipients of their new artifacts were high and would be difficult to duplicate for non-colocated groups. The presence of the artifacts as props is critical to in-person, collaborative discussions, and so is the ability of the presenter to get clear feedback about the recipients' level of understanding and acceptance of the information. Both colocated and non-colocated groups need support for examining a new information artifact together, while retaining the ability of individuals to adjust their views of the information, take notes, and retain private copies for their own records. Current technologies such as shared screens would help provide views of the artifact to remote participants, but only if the users could point at the material, annotate it, move about easily within it, and so on. In addition, the presenter must have sufficient communication bandwidth to be able to detect and interpret confusion on the part of the recipients. That is, the communication channels should support both shared views of the information and shared views of the people involved. Asynchronous communication mechanisms would need to provide rapid turnaround (which is especially important for joint interpretation of a document's implications) and expressive responses.

Kraut, Egido, and Galegher examined "high-quality real-time interactions" in their study of scientific collaborations [5] and identified similar requirements: support for the exchange of information created in very different formats (e.g., handwritten notes, text files, graphical images), joint manipulation of data, and full backchannel and feedback mechanisms.

A number of researchers have created new communication technologies for distributed workers that also include ways of sharing information [2, 4, 3, 10, 11, 13]. One of these projects, Cruiser [3, 10], was designed to provide people with opportunities to maintain awareness of others or engage in informal communication; information-sharing features were added to Cruiser after user feedback indicated that they would be useful. ClearBoard [4], Liveboard

[2], VideoWhiteboard [11], and We-Met [13] were meant from the beginning to support people engaged in *task-focused work* (we have borrowed this term from [3]).

Each of these systems primarily stresses the dynamic creation of new artifacts by two or more people, rather than the import and examination of pre-existing information artifacts. The latter capability is crucial to the interactions of the library clients in this study. Liveboard is the most flexible in offering a variety of ways to bring in existing information; users can run applications on Liveboard, do a slide show, and use a text editor, in addition to creating new material using the whiteboard feature. ClearBoard is unique in allowing a pair of collaborators to maintain eye contact as they draw, which is certainly critical to high-quality real-time interactions. ClearBoard also stresses seamless integration of the information space with the communication space, which we feel is not necessary to support our library clients; it should not be awkward to move back and forth between communication activities and information-manipulation activities, but we believe that the two visual fields (for people and information) need not overlap as completely as they do in ClearBoard..

For the purposes of our library clients, we emphasize that neither information nor interpersonal communication should be the main focus of the technology; they must share the spotlight. The prototypes described above are interesting and promising, but none of them brings together all of the features that would be needed for the information-sharing sessions between the library clients and their customers. These people need multi-user (not just pairwise) meeting spaces into which the participants can bring different kinds of artifacts to pore over and sometimes change together, and which allow them to read and respond quickly to one another's verbal and non-verbal signals of engagement and learning. It may be that some combination of the technologies involved in these prototypes offers the best solution for supporting high-quality real-time interaction.

## 7 Summary

We conducted a study of regular library clients in a variety of work settings and found that each one of them shared their search results with others. The library clients shared with their teams, consulting clients, scattered people who might be expected to find their results interesting, and archives. They often crafted new artifacts out of the search results in order to abbreviate and customize the raw data and add to its value for themselves and others. Many library clients customarily delivered information to others in collaborative settings, to ensure that the recipients would be able to understand and apply the information.

When we began the study, we were looking for system requirements related to information delivery and usage, but we found that communication needs must also be taken into account. The information artisans in our study described sharing scenarios that included analysis, synthesis, and communication components; this is where information and communication needs meet. In turn, information and communication technologies must come together to enhance the information-based interactions of these and other information artisans.

## 8 Acknowledgements

We thank the librarians and library clients who spent time sharing their search scenarios and tracing the paths that their search results took.

## 9 References

- [1] L. Bannon and S. Bødker. Beyond the interface: Encountering artifacts in use. In J. Carroll, editor, *Designing Interaction*. Cambridge University Press, 1991.
- [2] S. Elrod, R. Bruce, D. Goldberg, F. Halasz, W. Janssen, D. Lee, K. McCall, E. Pedersen, K. Pier, J. Tang, and B. Welch. Liveboard: a large interactive display supporting group meetings, presentations, and remote collaboration. In *Proceedings of CHI '92*, 3 - 7 May 1992. Austin.
- [3] R. Fish, R. Kraut, R. Root, and R. Rice. Video as a technology for informal communication. *Communications of the ACM*, 36(1), 1993.
- [4] H. Ishii and M. Kobayashi. Integration of inter-personal space and shared workspace: ClearBoard design and experiments. In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW)*, 31 October - 4 November 1992. Toronto.
- [5] R. Kraut, C. Egidio, , and J. Galegher. Patterns of contact and collaboration in scientific research collaborations. In J. Galegher, R. Kraut, and C. Egidio, editors, *Intellectual Teamwork*. Lawrence Erlbaum, 1990.
- [6] K. Kuutti. Activity theory and its applications to information systems research and development. In H.-E. Nissen, H. Klein, and R. Hirschheim, editors, *Information Systems Research: Contemporary Approaches and Emergent Traditions*. Elsevier Science Publishers, 1991.
- [7] P. Luff, C. Heath, and D. Greatbatch. Tasks-in-interaction: Paper and screen based documentation in collaborative activity. In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW)*, 31 October - 4 November 1992. Toronto.
- [8] D. Norman. Cognitive artifacts. In J. Carroll, editor, *Designing Interaction*. Cambridge University Press, 1991.
- [9] V. O'Day and R. Jeffries. Orienteering in an information landscape: How information seekers get from here to there. In *Proceedings of INTERCHI '93*, 24-29 April 1993. Amsterdam.
- [10] J. Rosenberg, R. Kraut, L. Gomez, and C. Buzzard. Multimedia communications for users. *IEEE Communications Magazine*, May 1992.
- [11] J. Tang and S. Minneman. VideoWhiteboard: Video shadows to support remote collaboration. In *Proceedings of CHI '91*, 3 - 7 May 1992. Monterey.
- [12] R. Taylor. Information values in decision contexts. *Information Management Review*, 1(1), 1985.
- [13] C. Wolf and J. Rhyne. Communication and information retrieval with a pen-based meeting support tool. In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW)*, 31 October - 4 November 1992. Toronto.