Fourth Quarter, 1996

Technology   SB Strategy & Business ...Booz, Allen & Hamilton

Finding the Knowledge Needle in the Data Haystack
(Page 1 of 9)▶

By David Berreby

When the coaches of the New York Knicks gathered last season to analyze a game against the Charlotte Hornets, one of their assistants made an interesting point: Charlotte's Glenn Rice, normally a forward, had played as a shooting guard and missed only one of six jump shots. That was much better than the Hornets' average success rate of 50 percent. The Knicks might want to think about why Mr. Rice was able to do so well against them in the shooting-guard position and come up with a strategy for dealing with it before they faced Charlotte again.

The Knicks coaches take such observations seriously. Last season, this same assistant pointed out that forward Charles Smith, who the Knicks had playing offense, actually did better against the Houston Rockets when he played defense. And that games in which forward John Starks made a lot of attempts were games the Knicks mostly lost. All of which were useful and shrewd perceptions. Not bad, considering that the analyst has been working for the Knicks for only a couple of years. And that, in fact, the analyst isn't even human.

The observations were made by an experimental program called Advanced Scout, developed for the National Basketball Association's coaches by scientists at I.B.M.'s Thomas J. Watson Research Center in Yorktown Heights, N.Y.

Using statistics kept by the coaches on play-by-play sheets, the Scout program looks for variations from the usual patterns of play and points them out in plain English sentences. With a CD-ROM containing digitized video of the game it studies, the program can even bring up the exact moments it wants to highlight and replay them on the laptop computers that the coaches bring with them on the road.

The Scout program is only one of a host of experimental projects and preliminary commercial ventures aimed at answering an increasingly important question for any large enterprise: As operations generate more and more data, what patterns, trends and relationships are hiding in the stream?

It is a particularly crucial question to companies and their managers, who "know" vast amounts about their businesses, in the form of data and documents to which they have access. Now researchers are focusing on computer products to help them really know what they know.

With the addition of Internet access -- and the creation of massive "intranets," closed to the outside world but available to all within a company -- managers must cope not only with data they themselves generate but also with millions of bits of information now available to them on a network.

Whether or not the Microsoft Corporation succeeds in fully melding the Internet into its Windows operating system, as it has recently announced it would, the trend is clear: Toward a seamless web of information, accessed on a computer, spanning yesterday's memos to yourself, comments from the Singapore office on your memo and sales figures for every product like yours for the last decade, not to mention the complete works of Shakespeare and the latest stock market moves. What was once a trickle of information coming into the office is now, in the words of Michael B. Spring, an associate professor of information science at the University of Pittsburgh, a firehose.

That can lead to a paradox: extra details can obscure patterns and make it harder to get to the useful facts.

Being able to call up any split seconds of play in a whole game off a CD-ROM, for example, could just lead to overload and confusion. So coaches -- and their managerial equivalents in corporations -- will need programs to track down the trends and correlations that count.

Eventually, when computing power gets good enough, such programs are likely to become intellectual prostheses, helping managers -- and their employees -- grasp and manipulate information, and patterns of information, in the flood of data.

❯

Fourth Quarter, 1996

Technology　　　　　　　Strategy & Business
Booz, Allen & Hamilton

Finding the Knowledge Needle in the Data Haystack
◀(Page 2 of 9)▶

"The first decade of computer development work was dedicated to information processing," said Aron Dutta, a principal at Booz-Allen & Hamilton in New York. "The second was dedicated to storage. The third was dedicated to bigger and faster processing, but now, in the fourth decade, the emphasis is on representation."

Mr. Dutta is the architect of Booz-Allen's "knowledge-management system," a computer network that links some 7,000 consultants working in 30 countries. The goal of the system is to let every consultant tap the experience and the insights of any other -- typing plain English questions into the computer and getting relevant documents, names of experts and, eventually, even a computer-generated simulation of how a particularly respected expert would look at the problem. And all of this, Mr. Dutta added, is tailored to a particular format and even industry jargon that the questioner prefers.

"The tool should help the users perceive something, see its relevance, define a context for it and take action," Mr. Dutta said. "If I query this stuff, why should I care where the information lives, or how it's represented?"

When such systems are widely in place, executives will be able to see hundreds of people working separately to create an international engineering standard or a complicated contract -- thanks to programs that keep track of what each worker has added, what each worker knows about the contributions of others and what relationships are forming among different documents as they are written and rewritten.

The executives will then have the ability on Tuesday morning to know everything interesting about the first 15,000 people who responded to the free offer in the ad campaign launched on Monday afternoon. They will be able to look at data about their companies' performance and watch the material glow, spin and even sing in ways that allow them to grasp important information as quickly as they can perceive a pixel turn from blue to red.

The new technology should be especially good at relating attributes -- noticing that product complaints this week are coming from central Michigan, or that people with contact lenses are buying the latest Three Tenors album. Finding such relationships fast is going to be an important aspect of marketing to the worldwide middle class of North America, Europe and Asia.

"Volkswagen gives away a mountain bike with its Jetta model because it discovered that buyers interested in Jettas tended to be sporty, outdoor types," said Robert Look, a marketing representative for I.B.M.'s Visualization Data Explorer, an already released software product. "Those are the kind of correlations that visualizing information helps you to find."

For instance, he said, the I.B.M. program turned up a distinct overlap between people interested in buying Subaru Outbacks and people who said they were in technical fields. "So Subaru knew that it was worth sending reps to the American Meteorological Society convention," Mr. Look noted.

The day is coming when marketing feedback will be instantaneous and ubiquitous, added Mr. Dutta of Booz-Allen. "It's now possible to get 1-800-FLOWERS on a Web page and send flowers to your mother. In a couple of years, you'll get a message as soon as you log on, saying, 'It's your mother's birthday tomorrow. Do you want to send flowers?' "

❮　❯

Fourth Quarter, 1996

Technology   Strategy & Business
...Booz, Allen & Hamilton

Finding the Knowledge Needle in the Data Haystack
◀(Page 3 of 9)▶

So the Knicks coaches are on a path that many other kinds of managers are sure to follow.

"Pretty soon, almost everybody will have information about their jobs, and they will need to mine it," said Inderpal Bhandari, who conceived of Advanced Scout and leads the group that works on the project. "We've got data collection automated. The salesman has instant sales figures; the security guard in the building has all the building logs. Now we are approaching the day when the salesperson in the field, for example, calls up sales information and looks for patterns that might help form a strategy for next week's calls."

As operations generate more and more data, what patterns, trends and relationships are hiding in the stream?

In an article much quoted by researchers in information management, Vanevar Bush, who had been the Federal Government's chief of science research during World War II, proposed 51 years ago that people would one day be able to use an "enlarged intimate supplement" to memory, a machine in which a person would store all his books, records and correspondence for easy, instant access. The technology for a real version of this device, which Mr. Bush called a "memex," is almost in place.

Managing and analyzing information, says Brewster Kahle, head of the Internet Archive in San Francisco, is going to be the hot central skill of the knowledge-based corporate future.

"In corporate culture, the finance guy used to be the guy with glasses counting the beans, but then in the 80's you started seeing C.F.O.'s running companies," Mr. Kahle said. "Maybe in the next decade, the librarians will start running companies."

Mr. Kahle, who was one of the founders of the Thinking Machines Corporation and later developed the Wide Area Information Server (the system that makes searching fast and easy on the World Wide Web), founded the archive last March with the goal of saving every bit and byte on the Internet for future historians. Since all the Web pages, newsgroups and other features of the Internet add up to about 1 to 10 terabytes of data (a terabyte is a million megabytes), Mr. Kahle expects that the project will lead to better ways of managing information on the gargantuan terabyte scale.

Among those ways, he suggested, will be devices that combine the role of library and researcher -- replacing the passive archive that just sits there with a thinking memory bank that knows what's in all the records.

"There's going to be a new science of handling huge amounts of information," Mr. Kahle said. "No one could ever read every book in the library, but computers can."

❮ ❯

Fourth Quarter, 1996

Technology   Strategy & Business
    Booz, Allen & Hamilton

Finding the Knowledge Needle in the Data Haystack
◀(Page 4 of 9)▶

Research has burgeoned in the last several years and the field is, as usual, changing quickly. "Things are happening really fast, now that everybody has discovered the Internet," said Cathy Marshall, a researcher at the Xerox Corporation's Palo Alto Research Center in California. "It's almost scary to work on development in this area."

Yet some themes are emerging.

One is that the torrent of information will require that the same volume of computer screen give more information. This is often described as a quest for new metaphors -- ways of representing information more densely, getting more bang from each pixel. (See accompanying article.)

"This whole issue of designing the appropriate metaphor is really the art to the science," said I.B.M.'s Mr. Look. Computers don't now provide the many different kinds of information that our senses get in the real world -- "like being able to tell from the color of paper if something's old," Ms. Marshall said -- but they are soon going to start taking advantage of those extra channels of human perception.

Mr. Spring, of the University of Pittsburgh, went back to research on perception done by psychology labs, making an inventory of human perceptions that might be put to use. The key, he said, is to find attributes that are "pre-attentive": traits people will take in without having to focus.

"Ask people to open a newspaper and find the word 'cat' on a page and it's hard to do," he explained. "But ask them to find the one letter in red ink and they find it instantly. In order to find a particular word, you have to process characters. But color you get right away. In the same way, certain kinds of sounds are pre-attentive. The most pre-attentive voice range is that of a young girl. That's why on fighter planes the missile-attack warning is a young female voice. It cuts through distraction but it's non-threatening."

According to Mr. Spring's research, people should be able to discriminate 156 different hues of a single color; approximately 60 levels of brightness; more than 15,000 distinct positions within a normal line of sight; some 70 shapes; 100 different relative sizes; and around 20 levels of loudness. That's a lot of empty channels that a computer display could fill with information. And that's not even counting relative opacity and texture, where Mr. Spring couldn't find any statistical research.

Already, he said, "I've got students doing the stock market, with positions represented as little spheres that chime or glow when you should turn your attention to them because something's happening." And with more computing power, he wants to develop a prototype system for representing constantly changing information using a virtual reality world. This is one form of dealing with information by representing abstract data as objects in space.

"Say a corporation has 30 offices," Mr. Spring said. "We collect information -- on gross sales, staff complaints, profitability, size of R.&D. budget, to name a few examples -- every day. We can assign attributes to those variables. Say the size of a sphere is gross revenues, its color is its profitability, the sound it makes is staff complaints, the texture might be rougher the more managers it had working there. Distances in space might represent the distance from the home office."

❮ ❯

Fourth Quarter, 1996

Technology   **SB** Strategy & Business
...Booz, Allen & Hamilton

Finding the Knowledge Needle in the Data Haystack
◄(Page 5 of 9)►

With goggles and earphones, Mr. Spring said, an executive could go "flying" among these spheres. "The program represents daily reports over 6 months, so as I run it, the spheres are changing. I might set it to run 6 months of data in 9 minutes. That means I can walk into my office, put on these goggles and get a sense of 6 months of performance in minutes. If I see a sphere go from blue to red, or start making more noise, then I want to be there and take a closer look.

"Of course," he went on, "no one with any brains is going to make a decision based on that, but guiding decisions is not the goal. The goal is to focus my attention where it needs to be."

What Mr. Spring's "virtual library" of spheres shares with already launched research projects, and with commercial applications, is an unprecedented flexibility in the presentation of the information.

"With our product, people tailor their view of the information to themselves," I.B.M.'s Mr. Look said. "A metaphor might be perfect for one person and horrible for the next."

With borders between desktop machine and network, and between corporate information and the Internet, all blurring, there is no need for everyone to see the same information the same way. Indeed, standardized ways to present information might even be considered a drawback.

Records on maintenance on an aircraft, for example, mean different things to engineers worried about metal fatigue and accountants worried about overtime. When libraries were physical spaces, the standardized form for information -- the library card catalogue, the insurance form with 14 boxes filled out by each client -- was the only hope of keeping track of everything. But in computers, information is kept as zeros and ones of code, and is then represented by a visual metaphor (some researchers prefer the term "user illusion"). That means the standardized presentation of information can give way to a make-your-own-sundae approach.

"Companies are seeing more and more of that on their intranet -- not needing to mandate that everything look one way or be done one way," said Ms. Marshall of Xerox PARC.

In fact, said Mr. Dutta of Booz-Allen, cultures vary so much -- not only from country to country, but from industry to industry and even from subdiscipline to subdiscipline -- that it is essential for information to change clothes when it crosses from one person's office to another. Japanese computer displays feature much more red and blue than do those in American offices, for example.

❮ ❯

Fourth Quarter, 1996

Technology   Strategy & Business
        Booz, Allen & Hamilton

Finding the Knowledge Needle in the Data Haystack
◀(Page 6 of 9)▶

"What we call 'business process re-engineering' is what they call a 'growth engine' in Japan," Mr. Dutta said. "My colleague in Japan may want to know, 'What are the best practices for describing a growth engine?' Someone in the United States might ask, 'What is the best practice for re-engineering for profitability?' They should each get the same result, displayed in the format and colors each is most comfortable with."

Ms. Marshall developed VIKI -- Visual and Kinesthetic Information -- a computer system designed to "spatialize information," which allows for the combination of lumping and arranging, pattern finding and randomness that researchers say workers will need to manage the information stream. For example, Ms. Marshall explained, if VIKI were applied to organizing Web sites, it would allow you to group and arrange sites in any shape you wished.

"Suppose you're doing a business analysis and you and your co-workers are gathering up Web references to other products," she said. "With VIKI, you can group and arrange that list of Web sites. You can decide to represent everything about video products as round, for instance, and all frequently checked ones as red."

Of course, the most efficient metaphor for information canonly present the user with the things the designers thought of. If they didn't think to make the spheres grumble when there's labor unrest, the user won't see it.

And so an equally important and lively line of research is aimed at helping people find the surprises in information -- the patterns and shifts they didn't expect or even conceive. The goal here is to make sure the user can get in under the metaphor and find what he wants or can't afford to miss. Projects in this field go by a variety of names, including data mining and information farming, reflecting different philosophies about the work involved. (Just as the metaphors that appear on screen are important, so are the ones that guide the people who write software.)

The tricky aspect of the work is that a program must strike a fine balance between just presenting a jumble of facts and presenting only what's already known -- the categories with which the inquiry started.

"When you mine data to discover knowledge, by definition you don't know what you're looking for," said Mr. Bhandari, the father of Advanced Scout. "So you can't have a system that commits you too early." Ms. Marshall put it this way: "It's important to be able to play with information, to be able to look at it in lots of different ways."

Data mining, Mr. Bhandari added, "is a technique that allows you to automatically identify and extract hidden patterns in mounds and mounds of data." The technique had been applied manually for decades in fields where it was worth the human hours to pore over information -- for example, in the oil industry's use of geological records.

With goggles and earphones, Mr. Spring said, an executive could go "flying" among these spheres.

"I realized four to five years ago that desktops were becoming more powerful and the Internet was spreading, and so I realized that the average person -- Joe Sixpack -- would be in a position to get a lot of data and mine it," Mr. Bhandari recalled. He and his team set out to create an analytical tool that could take any sort of input and find exceptional patterns in it. He wanted the device to take questions and give results in plain English.

❮ ❯

Technology   Strategy & Business
...Booz, Allen & Hamilton

Finding the Knowledge Needle in the Data Haystack
◀(Page 7 of 9)▶

In 1994, he got in touch with the Knicks because he had read that coaches were beginning to store their game records in data bases, and he figured coaches "fit the mold of the average user."

The experiment began shortly after the Knicks lost a game to Houston. The program swallowed coded versions of the play-by-play sheets that the coaches kept on each game, plus "the standard statistics -- the stuff you see in the paper," Mr. Bhandari said. "We set it up and we let the Knicks' guy play with it."

The program uses what Mr. Bhandari calls "attribute focusing" -- sifting through patterns, looking for anomalies, then sifting through those, looking for exceptions to patterns it finds in the first set of exceptions.

"It builds one level like that upon another until it gets to the most interesting level," Mr. Bhandari said. "This circumstance led to a different-from-normal shooting percentage. Then this subset of circumstances is even more different.

"That year," he continued, "they had a forward named Charles Smith. The Knicks had him cast in an offensive role. But the program pointed out that in games where he blocked a lot of shots, Houston usually lost. They were impressed. "The program, Mr. Bhandari said, is designed to relate any set of variables to any other. With certain kinds of statistics, like the basketball play sheets or surveys of customer satisfaction, there is no need to modify the program. Where the data are not already tidily arranged, as in, say, weather forecasting, "we would have to spend a few hours with you asking you what kind of questions you want to answer," he added. "All you need is to be able to table data. Then the program can relate the rows to the columns."

In other words, he said, "without much adjustment it can look for patterns in surveys of customer satisfaction. You get the results of focus groups on, say, the reliability of the product." Once you note a trend in responses, he said, the program should be capable of taking you from the abstraction to the raw information that underlies it -- for example, the part of the video in which a focus group makes the comments that are relevant to the trend. The program, he said, could also track "the behavior of buyers in response to different marketing approaches. Or medical insurance. You want to know if people are getting the right treatment. So you can mine medical information for the effects of different treatments on subsequent health and costs."

❮ ❯

Technology   Strategy & Business
...Booz, Allen & Hamilton

Finding the Knowledge Needle in the Data Haystack
◀(Page 8 of 9)▶

A different approach is what Ms. Marshall calls information farming. If data mining is burrowing into a mountain of facts for hidden nuggets, then information farming is staying in a fixed space, letting things change within it and watching and harvesting the information that is useful.

What holds constant is not the particular goal you apply to specific piles of information ("What's in here that can suggest a good marketing strategy for spring?") but rather the general characteristics of what you are interested in ("Keep me abreast of new developments in pharmaceuticals and check on how the Knicks are doing").

In information farming, you "use the computer as a personal information space, where you cultivate your understanding," said Mark Bernstein, a senior scientist at Eastgate Systems Inc. in Watertown, Mass. Eastgate is a leader in research on hypertext, the kind of linking via highlighted words that is the organizing principle of the World Wide Web.

"Fact-checking is information mining," Mr. Bernstein continued. "Journalism is information farming. If a wonderful thing you find isn't what you expected, that's fine.

"Everyone does information farming, if they're thinking, and everyone does information mining, if they have any analytical bent," he said, though he admitted the two schools tend to keep to themselves. "All these terms carry political connotations," he added. "Information farming is Jeffersonian. Information mining is Hamiltonian." A concrete example of an information farming task, Mr. Bernstein said, is now familiar to most people with access to a computer and a modem: keeping track of World Wide Web pages. Think of your presence on the Web as a kind of space, into which Web pages come. Maintaining that space so that you can harvest the information you need is information farming.

"Even a casual user will quickly reach the point of wanting to keep track of a couple of hundred Web sites," Mr. Bernstein said, but browsers like Mosaic and Netscape offer only "bookmarks" -- you can record a Web site's address onto a list of sites you know you will want to revisit.

"That's fine for seven to nine titles -- seven to nine being the number of different items that fit comfortably into a person's short-term memory," Mr. Bernstein said. A long list, arranged in the order of your visits, is like a big pile of laundry in the corner of the room. Without any kind of order to sort out the "maybe interesting" from the "I need to check in here every couple of days," the only thing a user can do is the computer equivalent of rummaging for clean socks.

Eastgate Systems' entry, Web Squirrel, tracks what you like to use. "Related Web sites can be grouped into neighborhoods," Mr. Bernstein said, "and different neighborhoods can be connected." The program uses Ms. Marshall's VIKI to allow people to set up their space as they please, he said. "How you define a neighborhood is up to you. The operative metaphor is spatial -- you create a geography that's specific to your mind, then project the geography of your mind into the computer space."

❮ ❯

Fourth Quarter, 1996

Technology   **Strategy & Business** ...Booz, Allen & Hamilton

Finding the Knowledge Needle in the Data Haystack
◀(Page 9 of 9)

Of course, much of the data people move around in the virtual world is not statistics, but words. Here the future of information management looks likely to belong to hypertext.

At the University of Pittsburgh, Mr. Spring and his colleagues and students are working on an experimental system called Cascade. Here, as elsewhere in information research, one goal is to get the computer display to convey more than it does now. Where a Web browser might offer words up in two colors -- one for a link that has been clicked, one for a link that has not -- Cascade has 20 colors, allowing users to indicate the type of document a link connects to, the source of comments attached to that document, the age of the document, the frequency with which it is accessed and so on.

"If you have a thousand people commenting on proposals for a new international standard, that can be pretty useful," Mr. Spring said.

Indeed, color coding is a clear example of a metaphor that makes information-getting easier, he said. In a Cascade document, new comments attached to a paragraph are blue; the ones a user has already seen are yellow. "That lets you find new comments immediately, without having to read anything," he explained. "Just find the blue bands in a scroll bar that's mostly yellow and you're instantly at what's new." (For the color blind, equivalents of such coding can be worked out using other attributes, like size.)

Cascade is also searchable "semantically," Mr. Spring said. "For example, if you search 'dog,' you'll get puppy, collie, pit bull and Lassie as well."

It has been said that the impact of new technology is overestimated in the short run and underestimated in the long run. Computers have yet to eliminate travel or paper or inefficiency, but their day is still young.

In the next few years, Mr. Spring said, "expect bigger screen territory on the desktop, lots of processing power and a ubiquitous high-speed network, available any place where you now find a bar napkin to write on or a pad by the phone."

Computer processing power, he noted, has been doubling every year and a half, a trend which, if it continues, in just a few decades promises information of such completeness and computer helpers of such intelligence that the concept of thinking could take on connotations we can't conceive today.

Yet the fact remains that computers won't be able to replace the human agent. Computers will be able to help people perceive and think as never before, but they won't be able to take over the job.

"The questions can't be automated away," Mr. Bernstein said. "After all, when I sit down in front of the machine to use it, the goal is not educating the computer, it's educating me.">

Illustrations by Ward Schumaker

Making a Better Metaphor

Everything, from Shakespeare to symphonies to shoe sales in September, is made of the same stuff when it is stored in a computer a series of on-and-off instructions, in the form of zeros and ones.

What the person in front of the machine sees, then, is a representation of that code, known as a metaphor or sometimes called a user illusion. Now that there is so much more to see, a heated search is on for better metaphors. And that has sent the experts back to the fundamentals of how perception and thinking work.

"Metaphors guide your thinking," said Cathy Marshall of the Xerox Corporation's Palo Alto Research Center in California. But, she added, "it's not clear what they facilitate and what they prevent. That's a subject of research."

Some of the results of the recent research boom have been surprising. For example, "people think of screen icons as nice things that take us away from typing and code, back to the days of arts and crafts," said Mark Bernstein of Eastgate Systems. "But the icon is very complicated and we've got a long way to go before we've mastered it."

Until the last decade, Mr. Bernstein said, "people thought icons were efficient because people understood them better than text. But that's *not* what icons are doing. A lot of icons are unintelligible to people from other cultures. A frame house with a roof and windows for 'Home' doesn't mean anything in parts of Africa and Asia, where houses don't look like that. Or take the symbol for stop -- hand out, palm forward, fingers extended up. This is not a polite gesture in some parts of the world. In fact, the rule in making international icons is, no hands."

"What icons turn out to be good for is just being small," Mr. Bernstein said. "You can get a lot onto a screen. That's their real advantage."

What researchers have come to realize, he said, is that a computer icon, a metaphor for information in the machine, arrives as a complex cultural object. It's a metaphor already. "The nervous guy with sweat flying off his head, that's a trope," he noted. "That's not real."

Some Rules for Creating Metaphors?

Don't neglect the wisdom of the ancients. Some scholars of hypertext have taken a look at the representations of spatial relationships in Egyptian hieroglyphics, whose creators had thousands of years to work out conventions for using those

relationships to organize information.

Many software people also adhere to traditions of the mighty and still-thriving Cult of Disney. "Animators figured out long ago that people respond warmly to features like a human infant's: big head, big eyes, small mouth," Mr. Bernstein said. (In fact, the Harvard evolutionist Stephen Jay Gould established years ago that Mickey Mouse had indeed evolved to be cuter and cuter as years passed, presumably in response to selection pressure from the marketplace.)

"So when you go to design the representative of an intelligent agent, you use that knowledge," Mr. Bernstein said.

Researchers trying to fill the screen with the most comprehensible material have found themselves going back to psychology research to discover just how the human brain processes its perceptions. Mr. Bernstein expects that vast arrays of information, like Brewster Kahle's Internet Archive or depictions of the Web, will probably use "a strategy of zooming and panning over a two-dimensional landscape, rather than letting people fly through a three-dimensional world. Because it turns out most people aren't well oriented when flying. They have trouble taking in the consequences of rotation."

As you might expect, metaphor researchers look hard at other representations of space -- at art, maps and architectural drawings, for example.

"We're bringing in the familiar to do that," Ms. Marshall said. "Bringing in perceptions from the physical world."

Further Reading

Surfers of the Internet can check out work mentioned in this article on the following Web pages:

- Michael B. Spring's Cascade
- Eastgate Systems' Web Squirrel

Mr. Spring's discussion of his research and the field in general can be found elsewhere on the Web, at two sites:

- The Virtual Library, Explorations in Information Space (PostScript document)
- The Document Processing Revolution (PostScript document)

For the linear-minded, some books worth consulting are:

- "Cyberspace: First Steps," edited by Michael Benedikt (MIT Press, 436 pages, $15.95, 1991).
- "Life on the Screen: Identity in the Age of the Internet," by Sherry Turkle (Simon & Schuster, 347 pages, $25, 1995).
- "Hyper/text/theory/" edited by George P. Landow (Johns Hopkins University Press, 377 pages, $16.95, 1994).

◀