<< Return to article

Search Beyond Google

Google reigns supreme as the search engine of choice—but for how long? A pack of startups—and Microsoft—are developing technologies to find what you want, faster.

Illustration by Tim Bower

By Wade Roush
March 2004

If employees at Google are anxious about the future, you wouldn't know it from a visit to the company's headquarters. Since last fall, when talk of an initial public offering got investors salivating, the organization has been under unusual scrutiny: some observers have called it "the hottest company on the planet," while others claim it's a business in leaderless disarray, with competitors crowding in and major customers on the verge of defection. But the Google complex in Mountain View, CA, is as outwardly carefree as any college campus. The main lobby is a study in shagadelic kitsch, with a baby grand piano, a spinning party light, and a row of neon-bright lava lamps arranged in the same blue-red-yellow-blue-green-red sequence as the company's familiar logo. The cafeteria pulses with rock music, shouted conversation, and the sounds of geeks slurping free gourmet food. Upstairs, in the cubicle farms, programmers chitchat across walkways littered with toys, Segway transporters, and the occasional canine.

It's only when I sit down in a quiet conference room with Google director of technology Craig Silverstein that the giddy dot-com mood turns more serious. Now that companies like Google and Internet ad agency Overture have demonstrated that displaying subject-specific paid ads alongside the results on a search page is a real moneymaker—contributing to an estimated $2 billion in industrywide revenues in 2003—a pack of wannabes are investing in search software they say will give users more pertinent results than Google's, faster. I ask Silverstein whether Google's famous focus on better technology will keep it ahead of all that competition. His answer

is circumspect.

"It's very easy to move from one search engine to a better one," he says. Google pays hundreds of researchers and software developers, including more than 60 PhDs, to man the front lines in this technology war, explains Silverstein, who is himself on extended leave from his doctoral studies in computer science at Stanford University. But he acknowledges that's no guarantee of victory. "We hope the next breakthrough comes from Google—but who knows?"

Who knows, indeed? According to Reston, VA-based research firm comScore, Google has a large lead over its rivals in U.S. audience share, accounting for 77 percent of all searches in August 2003 (including searches conducted at AOL and Yahoo!, which used the Google search engine). But in the search industry, innovation is a wild card. "In 1999, you could have said that AltaVista had pretty much finished off the search market," notes Whit Andrews, a research director at technology advisory firm Gartner. "In 1997, it was Inktomi. In 1995, it was Yahoo!. You never know in the search business when there's somebody down the street who is going to make you look like yesterday's news."

Google is vulnerable partly because it has few of the infrastructural advantages, like AT&T's once exclusive ownership of most of the telephone network or Microsoft's control of PC operating systems, that typically help to perpetuate dominance. (Indeed, press reports in January indicated that Yahoo! might soon drop its relationship with Google and turn to its own search technology.) And the company's claim to fame—the ability of its search algorithms to find the most relevant results, based on their popularity—may be growing stale. "When Google first launched, they had some new tricks that nobody else had thought about before," says Doug Cutting, an independent software consultant who wrote some of the core technology behind search engine Excite and has designed search tools for Apple Macintosh computers. But plenty of other search engines now offer intriguing alternatives to Google's techniques, Cutting believes.

For example, there's Teoma, which ranks results according to their standing among recognized authorities on a topic, and Australian startup Mooter, which studies the behavior of users to better intuit exactly what they're looking for. And then there's the gorilla from Redmond: Microsoft is turning to search as one of its next big business opportunities. Its researchers are devising a new operating system that melds Google-like search functions into all Windows programs, as well as software that scours the Web for definitive answers to questions you phrase in everyday English. Meanwhile, Yahoo! launched its own research laboratory in January, and Cutting himself is building an open-source alternative to Google (*see "Keeping an Eye on Google"*). "Nowadays," he says, "I'm not convinced [Google is] markedly better."

Whichever technology hooks tomorrow's Web surfers, its builder will earn enormous influence—and handsome profits. Some 550 million search requests are entered every day worldwide (245 million of them in the United States). By 2007, the paid-placement advertising revenue generated by all these searches will reach about $7 billion, says Piper Jaffray analyst Safa Rashtchy. Yet surveys indicate that almost a quarter of users don't find what they're looking for in the first set of links returned by a search engine. That's partly because the precious needles of information we seek are buried under a haystack that grows by some 60 terabytes

every day. And it's why brutal competition in the search industry is certain to continue, especially as search companies usher in a host of advanced technologies, such as natural-language processing and machine learning. "Over the next five to ten years," says Rashtchy, "we could see massive improvements that provide orders-of-magnitude increases in relevancy and usage." And it's the competition to deliver those improvements—much more than the success or failure of Google's rumored IPO, expected by many to happen this spring—that is likely to determine how we will be navigating the Web a few years from now.

Pulling Rank

By nature chaotic and decentralized, the Web screams out for tools to help people hunt down documents no matter where they reside. Say you want information on treatments for scurvy in the 18th century: without a search engine, you have no way of knowing that the data you need is stored only in places like a cryptically named file (www.jameslindlibrary.org/trial_records/17th_18th_ Century/lind/lind_kp.html) on a server at the library of the Royal College of Physicians in Edinburgh, Scotland.

"We hope the next breakthrough comes from Google—but who knows?" - Google director of technology Craig Silverstein (Photo by Debra McClinton)

When you type "scurvy" into a search box at Google or MSN or Ask Jeeves, however, you're still not touching the actual file at the Royal College. You're merely rifling through the search company's index of the Web—a huge list assembled by software "spiders" that crawl through thousands of pages every second, copying keywords, phrases, titles and subtitles, links, and other descriptive information. Once a fragment of information lands in the index, it's usually compressed, assigned a "weight" or importance, and stored in a database for quick retrieval. The search terms you enter are compared against this index, and links to pages that contain one or more of your terms are displayed in order of relevance.

How a search engine determines that relevance is the secret sauce. Google rocketed to prominence in 1999 largely due to PageRank, an algorithm invented by founders Larry Page and Sergey Brin that was the first to capitalize on the massive interlinking of Web pages. Each link is, in effect, a vote made by the author of one page for the contents of another. Page and Brin realized that if their index were big enough, they'd be able to assess a page's importance by counting the number of other pages that linked to it. They took other factors into account as well, such as the pertinence of the text surrounding the links and the linking pages' own popularity. But their groundbreaking insight was that the Web is a giant popularity contest—and that the most-cited pages will probably be the most useful. The technique worked fiendishly well, and Web users voted with their clicks. Between June 2000 and January 2004, former top dog AltaVista, which ranked results largely according to the number of times a page mentioned the user's search keywords, dropped from eighth place in overall Web traffic rankings to 61st, while Google climbed from near-invisibility to fourth place, according to data from research firms Media Metrix and Alexa. Google has so pervaded the Web that its very name was selected by the American Dialect Society as the most useful new word of 2002.

Despite its advantages, PageRank has a few flaws. Just as earlier search engines could be fooled by pages peppered with thousands of keywords in "invisible" white-on-white type, an unscrupulous site owner who wants his Web address to appear higher in Google's search results can easily publish thousands or even millions of junk pages that contain links to his site, artificially raising its rank.

(Google says it has ways of counteracting such attacks, but won't discuss them.) The same loophole in PageRank allows "Google bombing"—a recent phenomenon in which bloggers make a humorous or political point by creating so many links to a given site that it comes up first when users type a specific term into the Google search box. Google bombers protesting the war in Iraq, for instance, managed to make George W. Bush's White House biography the first-ranked result under "miserable failure."

More bothersome to some critics, however, is PageRank's obsession with fame. A legitimate page that matches a Google user's search terms perfectly may get buried in search results simply because there aren't enough other pages pointing to it, notes Daniel Brandt, a Web developer who runs a critical site called Google Watch. A page's relevance to an individual user, Brandt and other critics argue, may depend on more than its popularity. "Just because the rest of the planet thinks that this is the number one travel site doesn't mean it is the number one travel site for you," says Liesl Capper, founder and CEO of Sydney-based upstart Mooter, who believes she just might have a better way.

Putting a Stamp on Search Results

Enter the same search term into ten different search engines, and you're likely to get ten conflicting sets of results. That's partly because the search companies' spiders crawl different subsets of the Web; but more importantly, it's a reflection of the unique principles at work in each company's ranking algorithms. Here's how three search engines handle the term "**stamp collecting**."

	Ask Jeeves/Teoma	Google	Mooter
1.	**American Philatelic Society**	**Coin and Stamp Collecting** (About.com)	**Stamp Link**—Philately, Stamp Collecting's Best Site in Its Category
2.	**Joseph Luft's Philatelic Resources on the Web**	**Joseph Luft's Philatelic Resources on the Web**	**Warragul Philatelic Society (Stamp Collecting)**
3.	**Linns.com**: The website of the world's largest weekly stamp newspaper—Linn's Stamp News	**American Philatelic Society**	**Postal history**, philatelic covers and stamps for sale
4.	**Stamp Link**—Philately, Stamp Collecting's Best Site in Its Category	**BNAPS Stamp Collecting for Kids**	**Great Britain Philatelic Society**
5.	**Philatelic.Com**	**Linns.com**: The website of the world's largest weekly stamp newspaper—Linn's Stamp News	**Stamp Collecting, Philately, Stamp Auction**
	The top-ranked page has the highest "authority" — essentially, the most links—among communities of Web sites about stamp collecting. It's validated by references from resource pages (link collections from experts and enthusiasts—in this case, stamp collectors) and link	Google officials won't discuss how the Google engine arrives at rankings for specific sites. Patent documents and published papers, however, show that Google ranks pages according to how often other pages link to them. Google also takes into account such factors as the	Mooter first groups results into clusters or themes. The pages shown above appear in the "philatelic" cluster, ranked according to how frequently the search keywords and the cluster name appear

| popularity measurements similar to Google's. The runner-up, Joseph Luft's Philatelic Resources, has fewer and less qualified referrals from experts on the subject. | relevance of the referring pages and the text surrounding the links. Presumably, collectstamps. about.com is the most cited page on this subject in Google's index. | on each page, among other factors. Mooter "learns" the user's intent by noting which clusters and pages are clicked on, and reranks the results to reflect the apparent pattern of interest. |

Sources: Ask Jeeves, Google, Mooter. Results current as of January 5, 2004 (Ask Jeeves, Google), January 13, 2004 (Mooter).

A Starburst of Ideas

I'm lunching with Capper on a brilliant early-winter day in San Francisco. She's in town to call on potential investors and customers. "People who control the flow of information have a subtle but pervasive power," she tells me earnestly. "Someone has to hold that power, and it is important that the people who do are those who consciously try to have a positive impact, and who give power back to the individual." Mooter aims to do that by making Web searches easier and more personal. Capper grew up in Zambia, studied psychology in South Africa, and founded a chain of early-childhood-development centers before emigrating to Australia in 1997 and choosing search technology as the place to make her next impact. She set up shop in downtown Sydney and hired Jondarr Gibb, an experienced software architect, and John Zakos, a graduate student writing his Griffith University doctoral thesis on the applications of neural-network theory to Internet searches.

The personal touch: Mooter's Liesl Capper says search engines should "give power back to the individual." (Photo by Icon inc., Melbourne)

The three have mixed their ideas on psychology, software, and neural networks to create a ranking algorithm that learns from the user as a search progresses. Before dumping a long list of links on a user, Mooter analyzes the potential meanings and permutations of the starting keywords and, behind the scenes, ranks the relevance of the resulting Web pages within broad categories called clusters. The user first sees an on-screen "starburst" of cluster names. A search on the name Paul Cézanne, for example, yields clusters such as *art, artists, Cézanne, France, galleries*, and *famous paintings*. That's the psychology part. "When you do a traditional search, you get your millions of results, and your mind does a conceptual grouping," says Capper. "But our minds are hard-wired to process only three to five kinds of information at once. We decided not to override that but to work with it."

Then comes the learning part. To develop a more precise understanding of what the user is probably looking for, the Mooter engine notes which clusters and links get clicked and uses that information to improve future responses. Suppose a user enters the term "dog," clicks on a cluster called "breeds," and then spends a lot of time looking at sites about Schnoodles (a popular Schnauzer-Poodle mix). When the user clicks on a new search result, Mooter will personalize the ranking to reflect this apparent pattern of interest, which might, for example, lead to sites about "dogs" plus "breeds" plus "Schnoodles" appearing higher. A refined set of results appears on every page; the engine continues to adjust the rankings based on the user's behavior.

The whole idea is to give people the results they want in as few clicks as possible. "Two clicks and we already have a very good idea of where you're heading," says

Capper. When Mooter's beta site debuted last October, Capper didn't expect it to be noticed outside Australia. But traffic from around the world has been so heavy, she says, that the company has had to install more Web servers to keep the service running.

Spend much time talking to search-industry insiders and you'll realize that there are almost as many ways to rank search results as there are pages on the Web. Google's supposed overreliance on popularity was one of the inspirations behind Teoma (pronounced tay-o-ma), founded in 2000 by computer scientist Apostolos Gerasoulis and colleagues at Rutgers University in New Jersey. Teoma's search software now powers Ask Jeeves, the number six search site. Google "looks at the structure of the Web, but that method doesn't go down to the next level," says Paul Gardi, Teoma's senior vice president for search. "When you get down to the local level, you will find that links cluster around certain subjects or themes, very much like communities." For instance, pages on "home improvement" don't simply link upward to more popular pages; they also tend to link to each other, forming circles around prominent sites like Hometime.com, Homeideas.com, and BobVila.com.

The Rutgers scientists designed Teoma (Gaelic for "expert") to find those subject-specific communities and exploit their wisdom. Before the Teoma engine presents the results for a given set of keywords, Gardi explains, it identifies the associated communities and looks for the "authorities" within them—that is, the pages that community members' Web sites point to most often. Teoma tries to verify the credibility of these authority pages by checking whether they're listed on resource pages created by subject experts or enthusiasts, which tend to link to the best pages within the community. It then ranks search results according to how often each page is cited by authority pages.

IBM and other organizations experimented with similar authority-based ranking systems in the late 1990s, but Gerasoulis says their approaches could take hours to slog through all the pages out there. Gerasoulis's proprietary technique does the same thing in about a fifth of a second. Ask Jeeves dumped its previous search provider and switched to Teoma's technology in 2001, and its query volumes jumped 30 percent per year in 2002 and 2003.

Hard as it may be to believe when you're looking at a dozen pages of search results, today's search engines ignore most of what is out there on the Internet. Software spiders have difficulty indexing content that is protected behind sign-up forms or stored in databases such as product catalogues or legal and medical archives and only assembled into Web pages at the moment users request it. This so-called Deep Web may amount to as much as 92 petabytes (92 million gigabytes) worldwide, or nearly 500 times the volume of the surface Web, according to the School of Information Management and Systems at the University of California, Berkeley.

Mining the Deep Web is the mission of another fresh face in the search business— Chicago-based Dipsie. "Google and Teoma only index about 1 percent of the documents out there," says Jason Wiener, Dipsie's founder and chief technology officer. Wiener, a self-taught programmer who ran a San Francisco Web development company until the dot-com crash, has spent the last two years building a more nimble crawler, one that can get past forms and database interfaces. Say you're wondering about the standard equipment on a Mercedes 55SL convertible. At Cars.com, drilling down to the page with detailed product information will take

about six steps. Dipsie, however, will have indexed the entire Cars.com database in advance, so it can send you to the same page with a single click. "We don't handle anything that requires authentication with a username and password, but we do almost everything else," Wiener says. He claims that by the time Dipsie's search site becomes publicly available this summer, its index will include 10 billion documents—triple the current size of Google's index.

So while Google is still king of the hill, the hill itself is now crawling with competitors with their own bright ideas. "Google knows this," says Gartner analyst Andrews. "They were born at Stanford, and they know there are students in Stanford's classes who are saying, 'Hey, I've got an idea—what if we take this algorithm and stitch it together with that algorithm?' They've got to either hire the young turks or defeat them."

Microsearch

But if there is one software company that knows how to hire young turks and turn their ideas into market-dominating products, it's Microsoft. Name any hot corner of computer science, and the company Bill built is likely to employ at least one or two of the field's leading investigators: after all, the five Microsoft Research labs around the world employ more than 600 researchers. And when Microsoft smells a big market, it usually moves with full force to stake its claim.

There's nothing blue-sky about Microsoft's forays into information retrieval, the discipline from which the search engine sprang. The company has already won a 97 percent market share in PC operating systems and a 90 percent share in office software; search is one of the last big pieces of the computing landscape that Microsoft does not dominate. And a survey of R&D projects at the company confirms that it sees enhanced forms of search as key to its business growth. As the release of the next version of Windows, code-named Longhorn, grows nearer—a test version will be ready later this year—researchers and product developers are accelerating efforts to make Web searching an integral part of it.

One of the flashiest pieces of software in the works promises to allow you to enter your questions in simple English and get a direct answer back. The company believes search users shouldn't have to worry about selecting the right keywords, linking them together with the right Boolean operators (*and, or, not*, etc.), and scrolling through page after page of search results. Instead, says Microsoft researcher Eric Brill, search engines should understand and answer questions in natural language.

Take Microsoft Research's AskMSR program, which Brill and his colleagues have been testing on Microsoft's internal network for more than a year. At its core is a simple search box where users can enter questions such as "Who killed Abraham Lincoln?" and, instead of getting back a list of sites that may have the information they seek, receive a plain answer: "John Wilkes Booth." The software relies not on any advanced artificial-intelligence algorithm but rather on two surprisingly simple tricks. First, it uses language rules learned from a large database of sample sentences to rewrite the search phrase so that it resembles possible answers: for example, "___ killed Abraham Lincoln" or "Abraham Lincoln was killed by ___." Those text strings are then used as the queries in a sequence of standard keyword-based Web searches. If the searches produce an exact match, the program is done, and it presents that answer to the user.

In many cases, though, the program won't find an exact match, but only oblique

variations on the text strings, such as "John Wilkes Booth's violent deed at the Ford Theater ended Lincoln's second term before it had started." That's okay, too. As its second trick, AskMSR reasons that if "Booth" frequently appears in the same sentence as "Lincoln," there must be an important relationship between them— which allows it to posit an answer, even if it's not 100 percent confident (*see "Q: How Does Question Answering Work?" below*). "We are tapping into the redundancy of the Web," explains Brill. "If you have a lot of places where you are somewhat certain that you have found the answer, the redundancy makes it more certain." As the Web grows, so will its redundancy, making AskMSR ever more powerful, Brill reasons. While plans for AskMSR aren't definite, Brill believes the code will see the light of day, perhaps as part of a future Microsoft search engine.

Another Microsoft Research effort is less concerned with how search engines work than with how and when users need information. "Right now, when you want to search for information, you basically stop everything you're doing, pull up a separate application, run the search, then try to integrate the search result into whatever you were doing before," says Microsoft information retrieval expert Susan Dumais. "We are trying to think about how search can be much more a part of the ongoing computing experience."

Toward that end, Dumais is developing a program called Stuff I've Seen that's designed to give computer users quick, easy access to everything they have viewed on their computers. The interface to the experimental program, which will influence the search capabilities in Longhorn, is an always available search box inside the Windows taskbar. Enter a query into the box, and Stuff I've Seen will display an organized list of links to related e-mail messages, calendar appointments, address book contacts, office documents, or Web pages in a single, unified window. One emerging feature of Stuff I've Seen, called Implicit Query, would work in the background to retrieve information related to whatever the user is working on. If you're reading an e-mail message, for example, Implicit Query might display a box with links to the titles and e-mail addresses of all the people whom the message mentions, and to all of your previous e-mail from the sender. To make the software even more useful, Dumais is working on adding an item to the two-button mouse's standard Windows right-click menu that would be labeled "Find me stuff like this" and would search both personal and Web data for information related to a highlighted name or phrase.

AskMSR, Stuff I've Seen, and related projects are all part of a larger shift in technology strategy at Microsoft, one that could position the company to convert hundreds of millions of Windows users around the world to its own search technology, much as it wrested the Web browser market from Netscape back in the 1990s. The crux of this transformation is the new Windows File System, or WinFS—the very heart of Longhorn. Under the current Windows file system, each software application partitions its allotted storage space into its own peculiar hierarchy of folders. This makes it nearly impossible, for example, to link a chunk of information such as the name of the author of a Word document with the same person's address or phone number in Outlook. WinFS, by contrast, has at its core a relational database: an orderly set of tables stored on your hard drive where all the data on your computer can be searched and modified by all applications using a standard set of commands.

If Longhorn includes tools based on Stuff I've Seen and allows them to

communicate directly with a Web search engine, it could create the "single search box" dreamed of by software makers—the gateway to all the information you need, whether inside your PC or out on the network. Gartner's Whit Andrews, for one, is looking forward to Microsoft's new software. "Bring it on!" he says. "I am sitting here looking at my e-mail. If I want to look you up, I've got to remember to go Google you. But what I really want is to find out if I have talked with you in the past. So I want to right-click and search globally, search my e-mail and contact folders, search U.S. Search.com [which sells access to information stored in public records]. Who has that advantage? Microsoft is there, and for the low-price stuff that consumers aren't going to throw a whole lot of money at, they are in a terrific position."

Q: How Does Question Answering Work?
A: Like This

Microsoft's **AskMSR** software turns plain-English questions into formal search queries and polls the Web
for consensus answers.

1. Question	How many eggs are in a baker's dozen?
2. Rewrite query	"There are" + "eggs in a" + "baker's dozen" "A baker's dozen has" + "eggs" "baker's" + "dozen" + "eggs"
3. Collect search results and filter (for example, ignore results that do not resemble an answer to a "how many" question)	"A dozen usually has 12 eggs, so how many eggs does a baker's dozen have?" "The Baker's Dozen Cookbook" "Why are 13 eggs called a baker's dozen?" "13 eggs make a baker's dozen."
4. Extract answers from text and present most likely answers	13 eggs (81 percent likely) 12 eggs (7 percent likely)

Source: Microsoft

Meanwhile, Back at the Googleplex

I ask Google technology director Craig Silverstein whether Microsoft's search buildup keeps him up at night. He acknowledges that Microsoft and Google are exploring some of the same technical territory but contends that because Google is so much smaller than Microsoft (1,000 employees versus 55,000), it can act more nimbly on its ideas. And despite its smaller size overall, Google has more researchers devoted primarily to search than Microsoft. Silverstein also points out that each of Google's several hundred software developers is required—as part of the job—to spend 10 percent of his or her time on far-out personal projects, which provide a continuous flow of creative ideas.

Some of those projects surface at Google Labs (labs.google.com), a section of the Google site where the public can try out—and comment on—search-related software that's still in development. Google Viewer, for instance, animates results so that they scroll up the screen like movie credits. Voice Search lets you enter a search by telephone if you happen to be away from your desk, then retrieve the results online later. The Google Deskbar installs a permanent Google search box in the Windows taskbar; results appear in a small, temporary window, so users don't have to launch their Web browsers every time they want to look something up.

But none of the Google Labs prototypes represents an innovation of the magnitude

of Page and Brin's original PageRank algorithm. Nor are they in the same league as Microsoft's effort to reinvent Windows and integrate the applications that run on it. While Silverstein and his colleagues will talk about the efficiency of Google's more than 10,000 Web servers and the passion and drive of Google's programmers, they won't say how the company hopes to improve PageRank, or what new technologies might counter threats such as Teoma and AskMSR. So in the end, there's little outward proof that Google has the new ideas it will need to retain its market share. Says open-source programmer Doug Cutting, "Google has a whole lot of people trying to come up with monumental advances, but we haven't seen them. I think if they had them, they would show them."

One thing Silverstein does like to talk about is his long-range goal for search technology, which he believes is still in its infancy. "It's clear that the answer [to search] is not a ranked list of Web sites," he says. No one expects to approach a librarian, ask a question about the Panama Canal, and get 50 book titles in response, he argues. Silverstein thinks information retrieval experts should aim high, building software that is every bit as good at pointing users toward the specific resources they need as a well-trained reference librarian. That, of course, will require major advances in fields such as probabilistic machine learning and natural-language processing—and Google continues to hire some of the best new PhDs in those areas, including four recent graduates from the Stanford laboratory of Daphne Koller, a leading machine-learning researcher (*see "10 Emerging Technologies That Will Change Your World,"* TR *February 2004*).

But will all that talent be translated into tools people can use? Google itself appeared seemingly from nowhere, rapidly overshadowing other prominent search engines such as AltaVista. And if there is one message spread by the priests of the dot-com boom that still holds true, it is that people's desire for faster, more efficient ways to do things trumps brand loyalty every time. If rivals like Ask Jeeves and upstarts like Mooter or Dipsie achieve even part of their visions of better ranking algorithms, simpler interfaces, and larger, more comprehensive indexes, they could take a big bite out of Google's business. Microsoft's sweeping overhaul of the Windows environment, meanwhile, promises to change the very concept of search for the vast majority of computer users.

The good news for Internet surfers is that competition will make search utilities an even more helpful part of our daily lives. Without search tools, the Web's riches would be just as inaccessible as the tablets, scrolls, and hand-copied tomes of the pre-Gutenberg age, and as the Web itself grows, so does our need for better ways to penetrate it. But which technologies will provide the access we crave—and who will profit most from them—are questions that not even the best search engines can answer.

Wade Roush is a *TR* senior editor based in San Francisco.