What Do You Need to Know to Use a Search Engine? Why We Still Need to Teach Research Skills

Daniel M. Russell

■ *For the vast majority of queries (for example, navigation, simple fact lookup, and others), search engines do extremely well. Their ability to provide answers to queries quickly is a remarkable testament to the power of many of the fundamental methods of AI. They also highlight many of the issues that are common to sophisticated AI question-answering systems. It has become clear that people think of search programs in ways that are very different from traditional information sources. Rapid and ready-at-hand access, depth of processing, and the way they enable people to offload some ordinary memory tasks suggest that search engines have become more of a cognitive amplifier than a simple repository or front end to the Internet. As with all sophisticated tools, people still need to learn how to use them. Although search engines are superb at finding and presenting information — up to and including extracting complex relations and making simple inferences — knowing how to frame questions and evaluate their results for accuracy and credibility remains an ongoing challenge. Some questions are still deep and complex, and still require knowledge on the part of the search user to work through to a successful answer. And the fact that the underlying information content, user interfaces, and capabilities are all in a continual state of change means that searchers need to continually update their knowledge of what these programs can (and cannot) do.*

What does it mean to know something in the age of the search engine? To put it another way, how does a search engine transform human cognition, knowledge, and learning?

Web search tools are in many ways probably the most sophisticated AI systems that most people use on a daily basis. They're expressly designed for straightforward walk-up use without instruction, and in the vast majority of search tasks they perform admirably. The back-end AI systems parse the query, perform analysis of online content, do machine learning, have multiple kinds of modeling, and in the end, the user is shown an answer or a list of possible places to look on his or her own.

As a consequence of all this AI technology, it is now simple to look up nearly any particular piece of information within seconds. Need to know the population of Japan? The number of elementary schools in the United States? The signers of the Declaration of Independence and where they lived at the time? The distance from Earth to the sun? Given the rise of information in easily accessible online formats, these kinds of queries are fast, accurate, and available on your mobile device 24 hours a day, 7 days a week.

What's more, answers to queries like this can be made up of different media types. The diversity of different media types that can answer questions radically changes what we think of as content. Video makes certain things very different and simpler to learn than with traditional media; finding three-dimensional CAD models or online slide sets extends an ordinary answer to a question. Think of learning how to do origami by reading instructions versus watching an online video. Likewise, physical therapy exercises are much more easily followed on video, and even surgery procedures can be learned (and then performed) solely by studying online videos (Koya et al. 2012; Richard Santucci, personal communication, 2014[1]).

Most importantly, information technology changes the way we think, particularly in scholarly pursuits. This has also happened historically (see, for instance, Blair [2010] and Weinberger [2011]). From the introduction of printing, through cataloging systems, databases, and search engines, as we change the methods of organizing information to decrease the time to access information, we also change the fundamentals of the way we conduct research and think about knowledge more generally (Russell et al. 1993). As research scientists, we now search for literature, code, and data primarily online through search engines.

Looking up pieces of information has never been simpler, thanks to three forces: (1) the growth of content on the world wide web, (2) the increasing competence of search engines to index that content in sophisticated ways, and (3) improvements in the capability to parse queries expressed in question forms. As our society turns increasingly into a mobile, always-connected, always-on culture, the amount of time it takes to access information continues to decrease. The advent of widely available search engines is one of the success stories of software and hardware engineering. AI systems and techniques are so deeply enmeshed in their architecture that AI engineering can claim a large part of the credit for search engine successes.

By the end of 2014, more than 3 billion people had access to the Internet, meaning that they had the power to ask any question at any time and get a multitude of answers within milliseconds.[2]

However, with this ability comes the task of distinguishing between accurate, credible, true information and misinformation or disinformation. That skill, which was once in the hands of socially approved editors, publishers, librarians, professors, and subject-matter specialists, has now passed into the hands of everyone who is searching for the information. It's the searcher's challenge now.

Yet we're not doing an especially good job of teaching these skills to our students. It has become evident that high school students in the United States, when required to perform simple research tasks using multiple web resources, have difficulty selecting search keywords effectively, determining the credibility of a website, and discerning the bias of an Internet article (Hargittai 2002a, 2002b; Badke 2010). What's more, student online research skills seem to vary according to net family income, which is correlated with high use of the Internet at home and school. The skills needed to determine the credibility of available information mean that each of us needs to become an expert on understanding what the information we're finding actually means, how it's created, and where it's coming from (Leu et al. 2014).

At the same time, there has been a huge shift in the way we work and the way we think about things. Programmers seem to spend about as much time searching for code and development support as they do actually writing code (Umarji, Sim, and Lopes 2008). It's well-known that MDs and other knowledge-intensive professionals rely on Internet-scale search to maintain their command of relevant information (Hughes et al. 2009). And we all rely on a quick search to answer the simpler, smaller, less important questions that come up in our lives all the time.

The bigger question is this: In a world where we can do an online search for nearly any topic, what does it mean to be a literate and skilled user of information? What does a knowledge-based research skill mean? Is there still a role for advanced research skills of the kind once traditionally taught in libraries? I suggest that the answer is yes—there is still a need for instruction in this skill set.

Have our professional research skills kept up with the shifts in technology? Do they need to continue to improve as search engines become ever more capable of processing content?

Search Engines are Knowledge Engines

It's important to realize that search engines as we think of them now — Google, Bing, Yahoo, Naver, Cesnam, Yandex, Watson, Siri, Cortana, Google Now, and others — aren't just the text-mashing or link-analysis engines of the early 2000s. They are rapidly evolving into knowledge engines that do richer and deeper analysis, in addition to providing knowledge-based functions that rely on intelligent inference and semantic analysis. As a way to think about this new breed of knowledge tools, let's call these computer knowledge engine programs KEs, so we don't carry along the burden of "knowledge" or the biasing effects of what it means to be intelligent. The KEs we have built to date are impressively sophisticated now, and are continually adding to their knowledge competency. What does this mean for us, their users?

KEs Change the Way We Think About Knowledge

KEs provide many functions to their users, going well beyond the late-1990s model of text web-document indexing and giving access to a universe of images, maps, high-resolution Earth imagery, three-dimensional objects, local movie times, and videos (to name just a few). A KE's role has evolved into one where many different kinds of information are rapidly available with a simple interaction. Once available exclusively on desktop computers, hand-held devices can now reliably provide voice interactions to navigate from place to place, ask information-seeking questions, or to show family photos. With all of these different kinds of information resources, the temptation is to believe that all of human knowledge is available through a KE. The KEs give an implicit sense that the world's information space is "flat," and widely available through search, regardless of where it is, or how that information is organized (Zhang 2008, Rowlands et al. 2008). The most common expression is that ". . . just about everything is available via [a web search] these days…" even though that isn't even close to being true (Holman 2011).

The mental model users have of KEs isn't just that of a tool for searching web pages, but as a way to search for content and ask questions of what's available online—in public online content, as well as personal content. Although the perception is that "everything is available," the reality is that the breadth of content available also sometimes makes it difficult to find exactly what you're seeking, especially if it's in a highly crowded term space.

There's an ongoing public debate about the net influence of KEs on whether they make us collectively smarter (Thompson 2013), less intelligent (Carr 2010, 2012), or if their profound effect on the way we think is due to the quality and depth of information available on the web (Weinberger 2011).

Despite the arguments back and forth, it has become clear that people really do think differently when they realize that information can be quickly searched for, rather than simply remembered (Dror and Harnad 2008). As Sparrow, Liu, and Wegner (2011) show through their studies, simply knowing that information can be reliably found online (say, on a KE) changes the way that we learn information and can later recall it. These studies point out that KEs effectively become reliable partners in a form of external cognition (Hutchins 1995), analogous to the transactive memories that we have long used to remember who among our colleagues is a specialist in a particular domain.

But there's a trade-off. These findings suggest that human memory takes advantage of external memory and cognition systems in well-practiced, automatic ways. We learn what the KE "knows" and when we should attend to the information easily available in our computer-based memories. In essence, as Sparrow, Liu, and Wegner (2011) write, "We are becoming symbiotic with our computer tools, growing into interconnected systems that remember less by knowing information, than by knowing where the information can be found." That is, we become dependent upon online tools to the same degree we are dependent on the knowledge we gain from our friends and coworkers, and suffer recall deficits if those friends, coworkers, or KEs are not available.

KEs Do More Than Text Retrieval

In addition to using KEs as a memory amplifier, KE users are rapidly shifting away from thinking about queries as simple textual matches with synonym expansion to content on the web. This is a trend we see with increasing amounts of semantic and structured-data markup—a query can often pull an answer out of the context of the original data setting. For example, when we search for an error code or symptoms of a product malfunction, KEs can frequently extract the relevant portion of the page and present it as an isolated factoid (Paşca 2007; Gruber 2008).

Information Services

During the past few years, KEs' range of capabilities to answer questions has grown dramatically with improved query-parsing methods, better synonym handling, more robust text parsing, deeper knowledge analysis, and improved machine-learning techniques that map from queries to destinations. These continuing improvements have given rise to a view of KEs as question-answering systems, and not simply as advanced text fragment finders. For example, Google's Knowledge Vault (based on contributions from Freebase augmented with contributions from knowledge extraction methods) now has ~1B facts, each with an estimated probability of being true that is ≥ 0.9 (Dong et al. 2014). Despite this apparently large size, repositories such as Freebase or Knowledge Vault are still far from complete. For example, in Freebase (the largest open-source knowledge base), 71 percent of people in the system have no known place of birth, and 75 percent have no known nationality. Coverage for less common relations or predicates can be even lower (Bollacker et al. 2008).

What's more, KEs are increasingly becoming more than just sites where information/knowledge retrieval takes place, but they're also becoming more proactive (presenting information in anticipation of need, such as a phone number just before the meeting starts) and more task-centered (sending messages, making reservations, playing music on command as well as providing lyrics). With recent releases of proactive (aka predictive) information systems, users can discover information pushes of calendar alerts, weather, flight times, sports scores, transit directions, local restaurants, and others (Bohn 2012).

Computational Services

KEs also let users search for services that do different kinds of knowledge work, in essence becoming a kind of cognitive amplifier. Wolfram Alpha, with its sophisticated mathematics engine, is probably the best known of these specialty services. (On Alpha, for instance, the query [integrate sin x dx from x = 0 to pi] gives the numeric answer of 2, while [integrate sin x dx] gives back the symbolic expression –cos(x)

+ c). There are a great number of other kinds of online tools that can be found on the web through a KE search. This changes the way we think about knowledge: we know such knowledge-based tools exist, and we learn the skill of finding them. There has been a profound shift in the way people think about knowledge because these special computational services exist and are easily found by a KE. With tools such as calculators (mathematics, mortgage, great-circle route, body-mass-index), data "mashups" that allow easy searches over recombinations of multiple data sources (pricetracking, Twitter trends map, historical weather data on maps), reverse dictionaries, part-of-speech and grammatically marked-up current content or archives searches (Fraze.it), searchers now think of a KE as an assemblage of data, text, and information services, all of which can be used easily.

Web Content Constantly Changes

It comes as no surprise to learn that content on the Internet grows, disappears, and changes form rapidly. Just as importantly, not only are millions of new pages created daily, but the kinds of content change as well. New media types and aggregations continue to emerge regularly (such as Pinterest-tagged image collections, or professional question-answering sites like StackOverflow), and often in large quantities with impressive coverage on specific topics.

Content access also changes frequently due to copyright or policy-level changes (such as when access to a body of content is withdrawn, or when access rights to a data set are changed). This change in content stability can lead to "content surprise" when online content suddenly disappears or is replaced by newer content that doesn't preserve the old material. Google constantly adds new books to its corpus as new material is added through relationships with publishers or from scanning operations. Sometimes those contents are removed or have newly reduced access as well due to changes in copyright status.

There is also an increasing trend to create new kinds of information — with large amounts of data comes the opportunity to identify and extract new data held in the old data sets by reanalyzing the content, a process called datafication (Mayer-Schönberger and Cukier 2013). For example, German and Czech scientists reanalyzed large amounts of Google Earth imagery to identify the compass orientation of deer and cattle to discover a surprising north-south orientation bias (Begall et al. 2011), and 18th and 19th century ship logs are being datafied to recover weather data from those years for improved long-term weather modeling (Küttel et al. 2010).

Not only can data be analyzed for new signals that were there all along (but previously unidentified), data can now be reanalyzed to improve the original data itself. In other words, by reanalysis (such as through tweaks in the underlying OCR algorithm), the base data can be continuously improved. Just because the data set has been analyzed once doesn't mean there isn't more to be gained by reanalysis with improved algorithms. For the searcher, this means that questions that were once unanswerable with a given data set might now be useful with the newly reanalyzed data. An issue for the user is knowing which data sets have been upgraded and that reissuing queries would prove fruitful.

There is also an ongoing discovery of data resources that have been present, but not indexed or otherwise unavailable for searching. For instance, many images on the web have associated EXIF data that record time, date, exposure, image-specific unique IDs, camera serial numbers, and location information. Some KEs now index this information and make it searchable, making it possible to discover images taken by a particular camera, at a particular place, at a particular time. Not only does reanalysis give rise to new data, but it also exposes previously invisible data to indexing.

Overall, there has been an important shift in the way people think about information content: not only is it large, rapidly and continuously available, but it grows and changes moment by moment. The old mental models of knowledge as a slowly growing, slowly evolving repository are growing more out of date by the second (Thompson 2013). Not only does the content change, but the range of questions that can be asked does as well.

The Information World Is Not Flat

There is a tendency to believe that all the world's information is available through a simple KE search. While indexing the world's information is a goal, today it is still necessary to understand the structure of information resources. Searching for academic research papers is much more efficient if you use one of the scholarly information collections, rather than just searching on the global, open web. This selection of a resource to search is a kind of search scoping needed to include the appropriate kind of result. The information space isn't smooth, but has distinct structure. The more you know about that structure, the more effective you can be as a searcher. In fact, this repeats the advice given to reference librarians to "understand the range of information resources" available (Bopp and Smith 2011). Now, instead of knowing all of the reference books on the shelf, a skilled searcher will know the range and scope of online reference sites and tools, and understand how to find them.

In summary, a huge problem for KE users is knowing what's possible. This suggests an active, ongoing learning strategy on the part of the searcher. Even taking a class on information search will be a good beginning, but lessons learned there will quickly stale in the face of continuous, ongoing change in the KEs and content space. How can future searchers stay skilled and aware of what's possible?

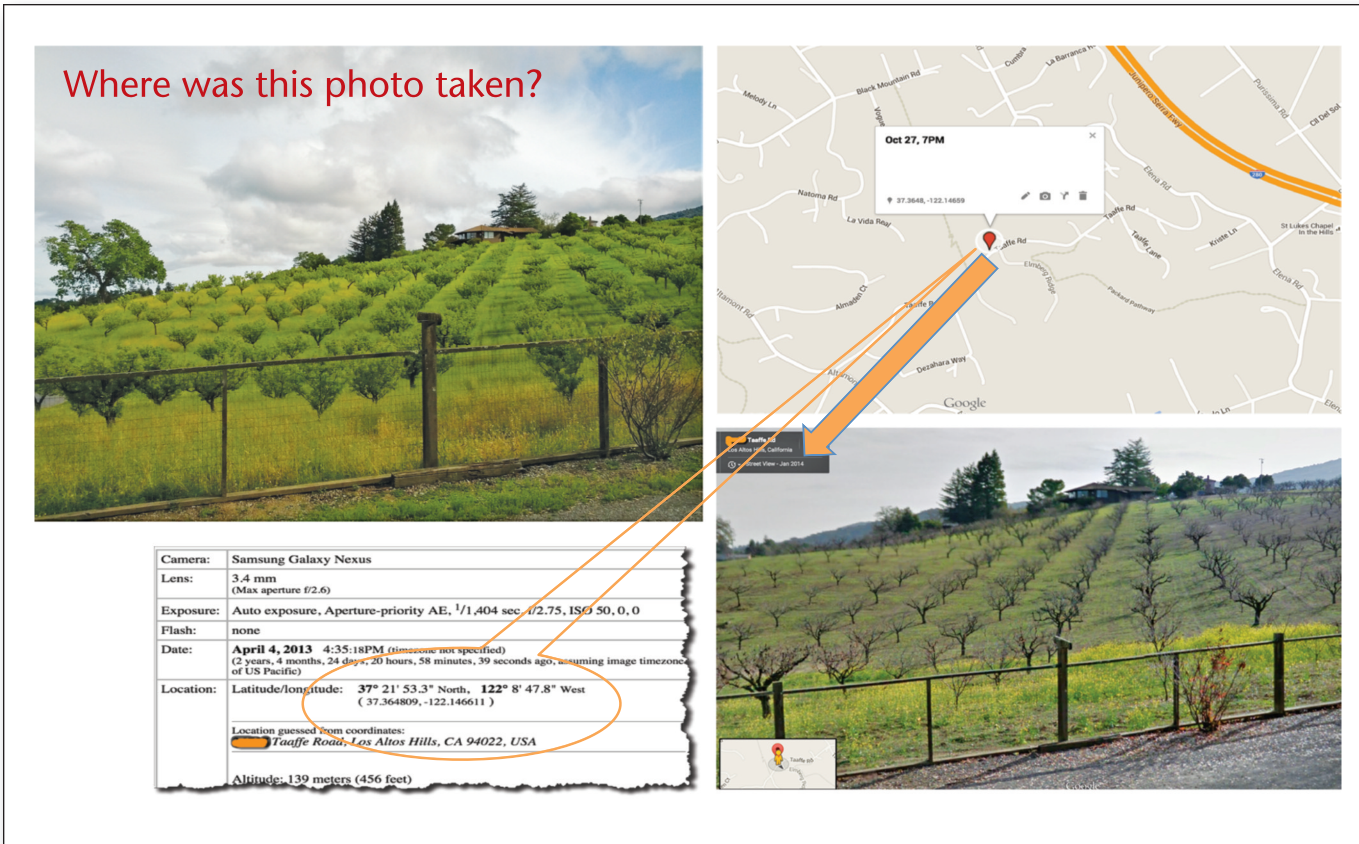

Figure 1. Where Was This Photo Taken?

Metadata information, such as the EXIF data captured on many digital photographs, can be used in unexpected ways to search for information across different kinds of resources. Here, EXIF data (lat, long) is taken from the photo and used to find a location in a map. That location can then be used to switch to a street level view to confirm that the photo was taken at the place claimed.

KEs Are Complex Systems

While KEs work to make the user experience of search simple and transparent, the fact is that a KE is a complex system that can sometimes behave in ways that don't match the mental model of the user. Users need to learn how to work with KEs and understand what they can do. The underlying knowledge base changes rapidly, but so too do the ways in which the KE system itself operates. A KE is sufficiently complex (for both algorithmic and data size reasons) that many of its behaviors are unintuitive.

Automation Surprise

All sufficiently complex systems seem to have inexplicable, magical behavior. KEs often have this property as well. When a complex system behaves in a way that's incongruous with user expectations, a mismatch between user and system models of what's going on arises that is often called automation surprise (Palmer 1995, Rushby 2002).

In KEs, searchers occasionally find themselves seeing results that don't make sense—they're strikingly irrelevant to what they expected to see in response to a query. Usually, this momentary confusion arises from inadvertently switching between content types (such as searching a News corpus while expecting to be seeing Web corpus results). This automation surprise is termed "stuck in a mode" and commonly comes about when the searcher is doing one task, then switches tasks without making the corresponding change in the modality of the search interface (Bredereke and Lankenau 2002).

The World's Knowledge Constantly Changes

In addition to constantly changing web content, we live in a world where knowledge constantly changes, even facts that have been considered verities of long standing and accuracy. While agreed-upon knowledge has always changed with new discoveries, here both the frequency with which new knowledge is being created, and the speed with which that knowledge makes its way into the canonical record have changed (Gleick 2012).

We live in an age of rapid knowledge discovery, and yet the transmission time to get content into schoolbooks is long. Many people now look to online references for information that is much more timely and more reflective of current understanding than traditional printed texts.

Understanding How Extraction and Inference Works in KEs

Going beyond sophisticated text search, modern KEs also actively extract information by processing text sources, looking for named entities, dependences between entities, coreference resolution (within each document), and entity linkage (which maps proper nouns and the coreferences to the corresponding entities already in the knowledge base). Accuracy of the extracted knowledge is then fused with supervised machine-learning methods to improve accuracy (Dong et al. 2014). Searchers need to understand when information found from a search is inferred versus accessed by text lookup.

One of the key components of KEs in the near term future will be their ability to make increasingly accurate inferences. IBM's Watson, for instance, combines multiple sources of knowledge to provide results based on deep reasoning, incorporating taxonomic reasoning when creating answers to queries (Kalyanpur et al. 2012). For the searcher with a complex question, understanding the basis from which the answer to a KE query is derived is essential. Thinking critically about where and how information is derived is an essential part of the research process.

Understanding the Answer Requires Understanding the Question

What is the distance from Earth to the sun? Depends on how and where in the elliptical orbit you measure. When asking that question are you seeking the distance from the center of Earth to the center of the sun? Or the distance of the closest point of the surface of Earth to the nearest point of the sun. Depending on which definition you use, the distance may vary by as much as ~700,000 kilometers (~435,000 miles).

This suggests that KE users need to be able to understand the basis on which KE inferences are made and the results offered up as authoritative. Again, sophisticated users of KEs need to understand that answers are driven by a consensus model and may be based on older information sources.

Reference librarians answer questions of this kind by conducting a reference interview in order to clarify the details of the library patron's question (Bopp and Smith 2011). Usually a librarian can, through knowledge of language and culture, understand that a question about the book *War and Peach* is actually a question about the Tolstoy novel *War and Peace*. Surprisingly, the AI spell correction system in KEs will rewrite "War and Peach" to the correct title. Does it work in all cases? Queries for [the Great Gatsbee] will spell-correct, but [the Great Gadfly] will not (not because the edit distance is too far, but because there's a great deal of content on the web with the string "Great Gadfly," and that overshadows any automatic correction).

Reference librarians also negotiate the boundaries of meaning between cultures. If a patron asks a question such as "Who is the president of Germany?" the librarian has to realize that the term "president" in the United States and "president" in Germany don't quite align. In all probability, the patron is asking about the chancellor of Germany (or less probably, who the Ministerpräsident of a German state is).

A large part of the skill in using a KE today is to take on the role of the reference librarian and recognize when the searcher has hit the edge of what the KE can answer. That is, one of the skills a searcher needs is to know when to extend, refine, and guide the search process. This means recognizing when responses given by the KE aren't lining up with each other, and when multiple resources must be consulted to draw an accurate and plausible picture. For the question [Who is the Ministerpräsident of Germany ?], a skilled searcher must look at the results carefully and quickly learn that this is an ill-formed question. It's a bit like asking [who is the president of Nevada?]. Both questions when posed to KEs will give answers and links to pages with content about Ministerpräsidents and German states (or Nevada and presidents of colleges there). Both times the searcher needs to recognize this lack of a real answer and dig more deeply to understand the question and reframe their search.

We Need to Understand What a KE Can Do

A striking characteristic of KEs is that they evolve rapidly in the range of capabilities they offer. Offering new capabilities is often seen as a competitive advantage. What's more, the range of capabilities will continue to constantly change as new aspects of content become available. In the history of scholarship, this is a new information landscape. Historically, information sources (and their access methods) have changed slowly over time (Blair 2010).

Since mid-2011, several KEs have offered the capability to search their image corpus by using an image as the query. This search capability allows us to search for images by similarity to a given picture. From a user's perspective, it means we can search not only for an entire image, but by understanding a bit about how it works (by computing a signature over the entire image), we realize that we need to search for images that have specific kinds of properties. Thus, for a large, complex image with many parts, the like-

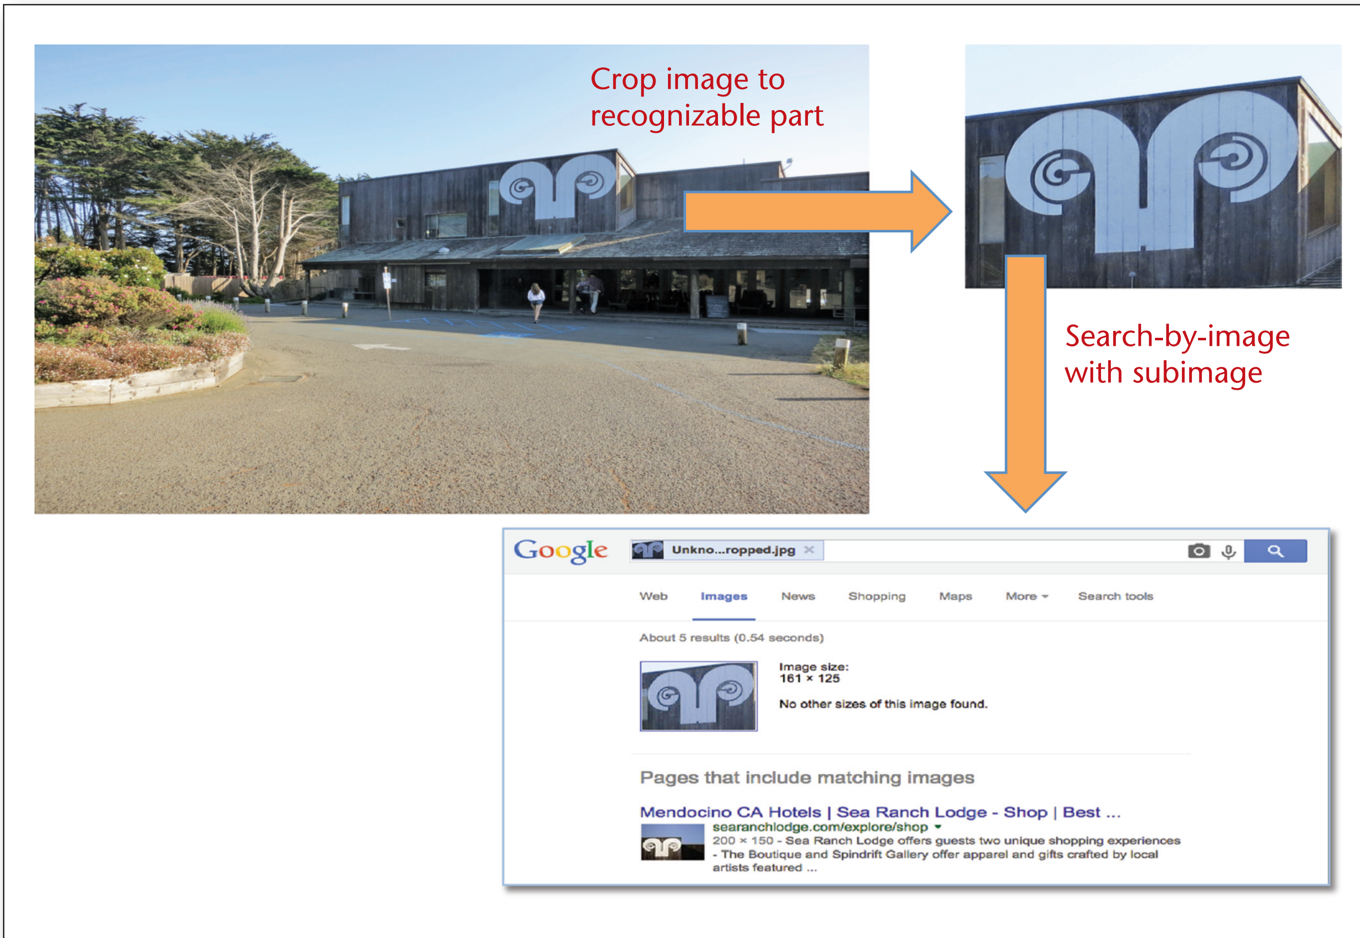

Figure 2. What Is the Logo in the Large Picture?

Knowing a bit about how Google's search-by-image mechanism works suggests that cropping down the large image to just the salient part is more likely to produce a match. The large image has many features, whereas the cropped subimage will more probably match an already existing logo image somewhere in the crawled images.

lihood is that the entire image won't be recognized, but if you crop the image to a part that might well be considered important, then you can find the answer.

Compounding these complications, people even don't understand many basic search capabilities and properties. A repeated finding from studies of KE users it that much of what can be done to use a KE is not widely understood or used (Hargittai 2002a, 2002b). The most dramatic example is that ~90 percent of the U.S. KE-using population does not know that it is possible to search for a string of text on a web page (Ma, Mease, and Russell 2011). Surprisingly, most people search the page visually, scrolling line by line to locate the information they need.

Part of understanding what a KE can do is understanding properties of the underlying information space. For example, the content over which the search is taking place is often specific, keyed to particular terms, and has a global extent. Even though it's not discussed much, key ideas such as relative term frequency and specificity are important pieces of search knowledge. This shows up when users with common names search for their names and are surprised by the lack of information about them. Such users need to understand that John Smith is a common name, as is José Lopez or Arun Gupta.

When systems evolved slowly, it was relatively simple to learn most of the capabilities of a tool. A researcher could easily know all of the functions and capabilities of a traditional research database and the nuances of its interface. But we now see that the underlying content, the capabilities, and the UX change frequently. How users think about and use AI systems may be affected by functional fixedness. They get stuck on a well-known or common use of the KE and don't consider alternative methods for solving a problem using new tools at hand (Duncker and Lees 1945).

In a recent study I conducted (October 2014), just under 30 percent of Amazon Mechanical Turkers calculated a simple mathematical expression incorrectly. Why? This wasn't a test of mathematical knowledge, but a test to see if they knew that KEs have built-in calculators. Those who got it wrong did so because they computed it by hand or by using manual calculators, this despite having just been primed to use a KE for search tasks in the previous question. This suggests that not only do the study participants not know that calculators are built into many KEs, but also that they don't know that it's possible to search for such a tool. At the same time, the KE interface for most KEs has a strong clue built into it: when you enter an arithmetic expression, a calculator opens up on the web page. Despite this obvious affordance, nearly one-third of skilled web KE users don't think to use the tool in this way.

What Does This Mean for Research Professionals?

As professional researchers, we live in a time of impressive change. Not only do we use few of the content search methods of a few years ago (imagine research life without scholarly content indexing!), but as described in this article, the content and methods are changing as well.

We Have to Recognize That Change Is a Constant

The KEs will change their user interfaces, adding and removing capabilities, and the underlying available information available will change. As a consequence, the ways in which we ask questions will change. We have to learn the conventions of the KEs and the landscape of information possibilities.

Need to Understand Coverage and Limitations

Knowledge workers who use KEs every day will need to understand what's in the realm of possibility and understand the assumptions that are built into the search processes. And the KE providers need also to provide simple ways to discover what's possible and understand the extent of possibilities and limitations. For example, there is currently no stemming in current KEs for Turkish, so search over documents in Turkish is very dependent on the forms of the words being used. This will doubtlessly change as text analysis methods improve with time, but it's currently a limitation on how well KEs search in heavily inflected languages. In the future, KEs need to become proactive about the ways they support experienced researchers in their uses.

Similarly, KE coverage is impressively large, but we have to overcome the illusion of omniscience, particularly with students learning to do their research online. The web is not the sum total of human knowledge. While we, as a culture, are putting more and more content online, and more content is "born digital," there is still content that's offline and will be unavailable for the foreseeable future. What's more, copyright and policy issues will keep content tied up in unsearchable ways, while corporate issues will continue to affect the availability of information.

Recognize That There Are Differences Between KEs

The large, general-purpose, broad-range KEs (for example, Google, Bing, and others) provide superb coverage and the ability to provide depth on topics of particular interest. They're extremely good at covering web content, news, text resources, images, videos, maps, and so on. They are less good at providing in-depth search services in specialty topics (for example, mathematics services, domain-specific context indexing, and others). The information landscape is not flat, nor are KEs completely universal in their coverage and competence. Each of the KEs offers a large number of different kinds of knowledge services to users.

Due to local policy or legal restrictions on what kinds of knowledge can be served, KEs will always be slightly different in their behavior from place to place. This isn't just an odd property of implementation, but a deep observation about the nature of social and political factors at work. Just as the content of encyclopedias was never consistent across national boundaries (Aaron Burr, the American traitor; Aaron Burr, the British hero), so too will KEs necessarily serve different versions of knowledge depending on where the query is issued and the knowledge received. Maps are currently different depending on where they're viewed (contended national boundaries always look different from the other side of the border dispute), and this is true for contentious data sets as well (Bowker and Star 2000). illustrate this well in their book about classification systems. Not only do medical categories vary substantially from place to place, but their use and interpretations do as well.

Understand How KE Search Interfaces Work

Currently, KEs are getting better at answering questions. This is clearly the direction of future KE development, but thus far there is no real working discourse model. At the moment when a question can't be answered (for example, [Daniel Russell doctoral advisor]), KEs don't signal that they lack the knowledge to give an answer, but fall back on giving a web search set of results instead. It's an important difference, one that's worth noticing.

A skilled KE user knows that the text abstracts (also known as snippets) for each web result are algorithmically generated without deep semantic processing. Effectively, the snippet composition system selects out fragments of text that score highly with respect to the interpreted form of the query. Those fragments are then concatenated together with ellipses, sometimes leading to an unintended interpretation when read as a summary of the page. If you know this about snippets, the correct reading is clear and straightforward—but this model isn't explained, and isn't widely understood by KE users.

In other words, a KE user has to learn to interpret the subtle signals that are often implicitly expressed in the interface design. Mostly (through iterative design and testing many variations on a theme) the UI designers arrive at a solution that works for most people in most cases. But for critical readings and for complex research tasks, the UI needs to be read with some skill and understanding. This includes attending to changes in the UI, as well as stepping in and questioning search results when an error seems possible. This is just the realm of critical thinking applied to using KEs.

Occasionally the result of deep search processing will actually be incorrect. Until recently, in response to the query [when was the Declaration of Independence signed] some KEs gave an answer of "July 4, 1776." While this is widely believed and commonly represented on many web pages, it's not correct—that's when the Declaration was approved. Signing took place weeks later, with most delegates signing on August 2, 1776.

As has always been true, skilled researchers will second-source their answers, particularly (as in this case) they see evidence of discrepancies in the different web pages that are sources of the information.

Ask Good Questions that Match the Capabilities of the KE

Searchers need to be sophisticated about what they are asking and thus what kind of answer to expect: the world is complicated and not all simple questions have simple answers. Example: When was the *USS Constitution* built? A: The keel was laid November 1, 1794. It was first launched on September 20, 1797 (but it accidentally stopped short of the water). It finally landed in the water and was commissioned on October 21, 1797. Even simple questions can have unexpectedly complex answers.

The increasing sophistication in representing world knowledge online also implies that asking the right questions will become more of a skill. A common error made among beginning searchers is to pose queries that have a built-in bias, a kind of leading question. This is fairly common among K–12 students who don't yet understand the basics of web search; in this case, that results are rank ordered depending on the terms in the query. So a query like [is the average length of an octopus 25 inches?] will give web links that look supportive of the supposition baked into the query (that is, that octopi average 25 inches in length), but only because there are so many positive hits that mention the terms "octopus" and "25 inches" on the same page. The KE doesn't really understand the question, but gives pages that best match the query, with its biases built in.

Read Carefully

Just as with the skill of reading snippets today, reading the answers generated by KEs carefully is an essential skill, particularly learning when new UI idioms come into play. For example, for a simple question like [what are the languages of Eritrea], the answer will be displayed as "Eritrean Official Languages: Tigrigna, English, Arabic" even though 6 other languages (with large, distinct populations) are also spoken there. If you miss the word "Official" in the answer, you''ll expect the answer to match your question and will miss the 1 million Tigre-speaking population of Eritrea. Likewise for the languages of South Africa—for the same query about South Africa, if you overlook the pulldown UI element in the interface, you might be forgiven for thinking that there are only 6 official languages, when in fact there are 11.

This isn't a critique of KE user interface design as much as it is a recognition that designs will continue to evolve to reflect the changes in underlying content and to show the results of new analytics that surface new kinds of information, and that the changing legal and political climate will influence information availability. For the user, this is important knowledge. The searcher needs both to be aware of the continuous evolution and to develop a kind of operational resilience in the face of ongoing changes.

Notes

1. Richard Santucci, M.D., FACS, is a urologic surgeon who teaches reconstructive surgery methods extensively throughout the world. His online videos have been used by surgeons in remote locations to learn procedures that are otherwise impossible to learn.
2. For current data, see InternetLiveStats.com/internet-users.

References

Badke, W. 2010. Lots of Technology but We're Missing the Point: Providing Only the Tools Isn't Enough. *eSchool News* (April 29). (www.eschoolnews.com/2010/29/lots-of-technology-but-were-missing-the-point)

Begall, S.; Burda, H.; Cerveny, J.; Gerter, O.; Neef-Weisse, J.; and Nemec, P. 2011. Further Support for the Alignment of Cattle Along Magnetic Field Lines: Reply to Hert et al. *Journal of Comparative Physiology A, Neuroethology, Sensory, Neural, and Behavioral Physiology* 197(12): 1127–1133.

Blair, A. M. 2010. *Too Much to Know: Managing Scholarly Information Before the Modern Age*. New Haven, CT: Yale University Press.

Bohn, D. 2012. Google Now: Behind the Predictive Future of Search. *The Verge* (October 29). (www.theverge.com/2012/10/29/3569684/google-now-android-4-2-knowledge-graph-neural-networks)

Bollacker, K.; Evans, C.; Paritosh, P.; Sturge, T.; and Taylor, J. 2008. Freebase: A Collaboratively Created Graph Database for Structuring Human Knowledge. In *SIGMOD, Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data,* 1247–1250. New York: Association for Computing Machinery. dx.doi.org/10.1145/1376616.1376746

Bopp, R. E., and Smith, L. C. 2011. *Reference and Information Services: An Introduction*. Santa Barbara, CA: ABC-CLIO.

Bowker, G. C., and Star, S. L. 2000. *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA: The MIT Press.

Bredereke, J., and Lankenau, A. 2002. A Rigorous View of Mode Confusion. *Computer Safety, Reliability, and Security: Proceedings of the 21st Internation Conference.* Lecture Notes in Computer Science Volume 2434, 19–31. Berlin: Springer. dx.doi.org/10.1007/3-540-45732-1_4

Carr, N. 2012. Is Google Making Us Stupid? What the Internet Is Doing to Our Brain. *The Atlantic* (July–August).

Carr, N. 2010. *The Shallows: How the Internet Is Changing the Way We Think, Read, and Remember*. New York: Atlantic Books Ltd.

Dong, X.; Gabrilovich, E.; Heitz, G.; Horn, W.; Lao, N.; Murphy, K.; and Zhang, W. 2014. Knowledge Vault: A Web-Scale Approach to Probabilistic Knowledge Fusion. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 601–610. New York: Association for Computing Machinery. dx.doi.org/10.1145/2623330.2623623

Dror, I. E., and Harnad, S. 2008. Offloading Cognition onto Cognitive Technology. In *Cognition Distributed: How Cognitive Technology Extends Our Minds,* ed. I. Dror and S. Harnad, 1–23. Amsterdam: John Benjamins Publishing. dx.doi.org/10.1075/bct.16.02dro

Duncker, K., and Lees, L. S.. 1945. On Problem-Solving. *Psychological Monographs* 58(5): i. dx.doi.org/10.1037/h0093599

Gleick, J. 2012. *The Information: A History, A Theory, A Flood.* New York: Vintage Books.

Gruber, T. 2008. Collective Knowledge Systems: Where the Social Web Meets the Semantic Web. *Web Semantics: Science, Services, and Agents on the World Wide Web* 6(1): 4–13. dx.doi.org/10.1016/j.websem.2007.11.011

Hargittai, E. 2002a Second-Level Digital Divide: Differences in People's Online Skills. *First Monday* 7(4)(1 April). dx.doi.org/10.1002/asi.10166

Hargittai, E. 2002b. Beyond Logs and Surveys: In-Depth Measures of People's Web Use Skills. *Journal of the American Society for Information Science and Technology* 53(14): 1239–1244.

Holman, L. 2011. Millennial Students' Mental Models of Search: Implications for Academic Librarians and Database Developers. *The Journal of Academic Librarianship* 37(1): 19–27. dx.doi.org/10.1016/j.acalib.2010.10.003

Hughes, B.; Joshi, I.; Lemonde, H.; and Wareham, J. 2009. Junior Physician's Use of Web 2.0 for Information Seeking and Medical Education: A Qualitative Study. *International Journal of Medical Informatics* 78(10): 645–655. dx.doi.org/10.1016/j.ijmedinf.2009.04.008

Hutchins, E. 1995. How a Cockpit Remembers Its Speeds. *Cognitive Science* 19(3): 265–288. dx.doi.org/10.1207/s15516709cog1903_1

Kalyanpur, A.; Boguraev, B;. Patwardhan, S.; Murdock, J. W.; B. Boguraev, Lally, A.; Welty, C.; Prager, J.; Coppola, B.; Fokoue, A. 2012. Structured Data and Inference in DeepQA. *IBM Journal of Research and Development 56*(*3/4*): 10:1 – 10:14.

Koya, K. D.; Bhatia, K. R.; Hsu, J. T.; and Bhatia, A. C. 2012. YouTube and the Expanding Role of Videos in Dermatologic Surgery Education. *Seminars in Cutaneous Medicine and Surgery* 31(3): 163–167. dx.doi. org/10.1016/j.sder.2012.06.006

Küttel, M.; Xoplaki, E.; Gallego, D.; Luterbacher, J.; Garcia-Herrera, R.; Allan, R.; Barriendos, M.; Jones, P. D.; Wheeler, D.; and Wanner, H. 2010. The Importance of Ship Log Data: Reconstructing North Atlantic, European, and Mediterranean Sea Level Pressure Fields Back to 1750. *Climate Dynamics* 34(7–8): 1115–1128. dx.doi.org/10.1007/s00382-009-0577-9

Leu, D. J.; Forzani, E.; Rhoads, C.; Maykel, C.; Kennedy, C.; and Timbrell, N. 2014. The New Literacies of Online Research and Comprehension: Rethinking the Reading Achievement Gap. *Reading Research Quarterly* (14 September). dx.doi.org/10.1002/rrq.85

Ma, L.; Mease, D.; and Russell, D. M. 2011. A Four Group Cross-Over Design for Measuring Irreversible Treatments on Web Search Tasks. In *Proceedings of the 44th Hawaii International Conference on System Sciences* (HICSS), 1–9. Piscataway, NJ: Institute for Electrical and Electronics Engineers.

Mayer-Schönberger, V., and Cukier, K. 2013. *Big Data: A Revolution That Will Transform How We Live, Work, and Think.* New York: Harcourt.

Palmer, E. 1995. Oops, It Didn't Arm. A Case Study of Two Automation Surprises. Paper presented at the 8th International Symposium on Aviation Psychology, Columbus, OH, April 24–27. (www.gbv.de/dms/tib-ub-hannover/227206517.pdf)

Paşca, M. 2007. Lightweight Web-Based Fact Repositories for Textual Question Answering. In *Proceedings of the 16th ACM Conference on Information and Knowledge Management,* 87–96. New York: Association for Computing Machinery. dx.doi.org/10.1145/1321440.1321455

Rowlands, I.; Nicholas, D.; Williams, P.; Huntington, P.; Fieldhouse, M.; Gunter, B.; Withey, R.; Jamali, H. R.; Dobrowolski, T.; and Tenopir, C. 2008. The Google Generation: The Information Behaviour of the Researcher of the Future. *Aslib Proceedings* 60(4): 290–310. dx.doi.org/10.1108/00012530810887953

Rushby, J. 2002. Using Model Checking to Help Discover Mode Confusions and Other Automation Surprises. *Reliability Engineering and System Safety* 75(2): 167–177. dx.doi.org/10.1016/S0951-8320(01)00092-8

Russell, D. M.; Stefik, M. J.; Pirolli, P.; and Card, S. K. 1993. The Cost Structure of Sensemaking. In *Human Factors in Computing Systems: INTERCHI '93 Conference Proceedings: Bridges Between Worlds*. New York: Association for Computing Machinery. dx.doi.org/10.1145/169059.169209

Salomon, G. 1997. *Distributed Cognitions: Psychological and Educational Considerations*. Cambridge, UK: Cambridge University Press.

Sparrow, B.; Liu, J.; and Wegner, D. M. 2011. Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips. *Science* 333(6043): 776–778. dx.doi.org/10.1126/science.1207745

Thompson, C. 2013. *Smarter Than You Think: How Technology Is Changing Our Minds for the Better*. New York: Penguin.

Umarji, M.; Sim, S. E.; and Lopes, C. 2008. Archetypal Internet-Scale Source Code Searching. In *Open Source Development, Communities, and Quality, IFIP 20th World Computer Congress, Working Group 2.3 on Open Source Software,* 57–263. Berlin: Springer. dx.doi.org/10.1007/978-0-387-09684-1_21

Weinberger, D. 2011. *Too Big to Know: Rethinking Knowledge Now That the Facts Aren't the Facts, Experts Are Everywhere, and the Smartest Person in the Room Is the Room*. New York: Basic Books.

Based Question Answering. In *Proceedings of the 23rd International Conference on World Wide Web,* 515–526. New York: Association for Computing Machinery. dx.doi.org/10.1145/2566486.2568032

Zhang, Y. 2008. Undergraduate Students' Mental Models of the Web as an Information Retrieval System. *Journal of the American Society for Information Science and Technology* 59(13): 2087–2098.

Daniel Russell is a research scientist at Google where he works in the area of search quality, with a focus on understanding what makes Google users happy and productive in their use of web search. As an individual contributor, Russell is best known for his studies of cognitive sensemaking behavior of people dealing with understanding large amounts of information. He is also known for work on the large, interactive IBM BlueBoard system for simple shoulder-to-shoulder collaboration, and his contributions to online education in the form of the Massive Open Online Course (MOOC) called PowerSearchingWithGoogle.com, which has taught search skills to more than 600,000 students. Before joining Google, he also held research positions at IBM's Research Almaden Research Center (San Jose, CA), Apple's Advanced Technology Group (ATG), and Xerox PARC. Russell has also been an adjunct lecturer in AI at University of Santa Clara and at Stanford University, and is currently an adjunct faculty member at the University of Maryland, College Park. Russell received his B.S. in information and computer science from the University of California, Irvine, and his M.S. and Ph.D. degrees in computer science from the University of Rochester (1983).