

CHAPTER 1

Introduction

Information seeking is a core life activity. People seek information for a number of reasons, including to find facts (Marchionini and Shneiderman, 1988) and to learn and facilitate effective decision making (Marchionini, 2006). Searchers may also have other motivations, such as pleasure and enjoyment (Wilson and Elweiler, 2010). Historically, people have sought information through interactions with others, either synchronously through in-person or telephone dialog or asynchronously through written communication. Increasingly, information seeking is being conducted via *automated* search systems such as Web search engines. Meeting searcher requirements across the full spectrum of search goals that individuals may have is challenging with such generic search systems. Nonetheless, search engines are offering an increasing range of reactive and proactive search services to meet and anticipate searcher needs. Helping searchers with specific search tasks requires targeted search support, as well as different methods and criteria under which to evaluate the performance of search systems in different circumstances. For example, fact-finding tasks may require only a single resource or direct answer, and strong system performance may be evidenced by low searcher engagement and short task completion times. In contrast, when the goals involve exploration, learning, or enjoyment, richer search and exploration support may be needed (e.g., diverse results or query suggestions, dynamic information visualizations), and increased engagement over a much longer timespan may be a positive indication of system performance. Understanding the nature of the search task(s) being attempted is important in both designing and evaluating the support offered by search systems.

Effective interactions with automated search systems are a critical aspect of successful searching. These interactions can range in complexity from basic text query entry and result selection to rich gestural and spoken dialog interactions. Search tasks can also continue longitudinally and increasingly, these tasks transcend devices, domains, and applications. The range of search interactions, and the support that systems offer for performing them, is expanding given technological advances. Search systems are becoming more intelligent and more aware of their users' interests and intentions, as well as of their surroundings when performing searches. This enables these systems to anticipate people's needs more accurately and to work symbiotically with searchers

to support task completion directly. Developing systems capable of serving as cognitive prosthetics to amplify human intelligence was proposed by visionaries some decades ago (Bush, 1945; Licklider, 1960; Engelbart, 1962). Advances in user modeling, semantic understanding of queries and documents, machine learning, and artificial intelligence mean that this vision is closer than ever to being realized. That said, there is scant literature on the new wave of technology and its relationship with the design of next-generation search systems. I address this shortcoming in this book, where I aim to summarize past and present research and development in the important area of search interaction, but also to spend significant time discussing a future enabled by a number of recent technological advances, as well as changes in the attitudes and expectations of searchers, which is affecting how searcher data is collected and used, and the types of search support offered. I use the phrase “next-generation search systems” since I believe that we are on the cusp of a revolution of how people engage with search technology and how search systems will empower searchers. Given the importance of these recent developments, I adopt a perspective that is more systems oriented than conceptual. I also focus primarily on the Web because of the wealth of recent research in this area and the opportunity for advances at scale to enhance quality of life for individuals worldwide. In principle, much of the work covered herein could be applied to any document collection or search domain, including digital libraries and organizational intranets.

1.1 Historical Approaches

Searches are often motivated by an incompleteness (Ingwersen, 1992; Mackay, 1960; Taylor, 1968) or a “problematic situation” (Belkin, 1982a, 1982b) in the mind of the searcher that develops into a desire for information. The recognition and acceptance of an information problem typically resides at the beginning of the information-seeking process (e.g., Ellis, 1989; Marchionini, 1995; Wilson, 1997). The problem can be internally motivated (e.g., curiosity) or externally motivated (e.g., an assignment). It may be characterized by a gap (Dervin, 1977), a visceral need (Taylor, 1968), an anomaly in a searcher’s knowledge state (Belkin, 1982a, 1982b), as a defect in a mental model, or as an unstable collection of noumena (Marchionini, 1995). Once the problem has been accepted, it must then be understood and defined. To do so, it must be limited and labeled, and a framework for the answer must be constructed. This is defined as the “conscious need” (Taylor, 1968). During this process, attributes of candidate solutions emerge that will ultimately guide search interactions. This process leads to the development of Taylor’s “formalized need” and the possible articulation of an information-seeking task. The searcher defines the problem internally as a task with properties that allow progress to be judged and a particular search strategy to be selected.

Although much of the search interaction that I consider in this book involves richer modeling and support for search interaction behavior, I begin by focusing on the traditional lookup model that has been fundamental in the design of support for searching. This model focuses on query-document matching and is used to represent search activity in a range of applications, including generic Web search. The “10 blue links” method (Broder et al., 2010) is the dominant interaction model currently used in the



Figure 1.1. Lookup-based information retrieval model (adapted from Bates, 1989).

development of database management systems and in major commercial Web search engines offered by search providers such as Google, Yahoo!, and Microsoft. Figure 1.1 depicts this retrieval model. The development of methods and tools to support the retrieval of relevant content per this model drove many of the early advances in Information Retrieval (IR) research (e.g., Van Rijsbergen, 1979; Salton and McGill, 1983), and continues to this day in settings such as the annual Text Retrieval Conference (TREC) (Voorhees and Harman, 2005) and other evaluation fora. The components of the lookup model are the collection being searched (*Documents* in Figure 1.1), a representation of the documents that are stored the collection (*Document Surrogates*, usually as an inverted index for rapid document lookup), the underlying information need of the searcher (*Information Need*), and a query statement (*Query*, traditionally in textual form provided by the searcher at query time). To use search systems employing this model, searchers provide a query statement representing their information needs, and the system returns a ranked list of document surrogates comprising titles (the “blue links” referred to by Broder et al., 2010), result captions, uniform resource locators (URLs), and other relevant information (e.g., Web page size, topical category, most recent page modification timestamp) for examination and selection by the searcher. Beyond ranked lists of search results, other methods such as result clusters and alternative visualizations of results are also available and will be discussed in more detail later in the book (Chapter 6).

The lookup-based model has promoted our understanding of IR in many ways (e.g., as the basis by which systems are evaluated at the Text Retrieval Conference; Harman, 1993; Voorhees and Harman, 2005). Lookup tasks are usually suited to analytical search strategies that begin with carefully specified queries and yield precise results with minimal need for result set examination and item comparison, such as fact-finding or question-answering scenarios (Marchionini, 2006a). A common assumption behind this model is that the information exists in the collection and searchers need only to retrieve it by formulating an appropriate query, with or without assistance from the search system. Under this paradigm, the query is treated as a one-time conception of the searcher’s information need. However, many real searches contain multiple query iterations, post-query browsing, detailed result examination, and ill-defined/exploratory needs, none of which are adequately captured in this model. Hypothesis generation, decision making (an aspect of information use), and the impact of search interaction on searcher knowledge in the short and long term are also important and require support from search systems – yet have only been given limited attention in prior research on the search process, especially in Web search settings where the focus is on outcomes that are easier to measure, e.g., relevance and satisfaction.

As Kuhn (1970) noted, major models that are central to a field eventually begin to show inadequacies as testing leads to improved understanding of the processes being studied. This applies to the lookup-based interaction model as a basis for

information-seeking research. The model has come under increased scrutiny for quite some time (e.g., Bates, 1986a,b; Belkin et al., 1982a,b; Ellis, 1984; Ingwersen, 1992; Kuhlthau, 1993; Marchionini, 1995). The model inadequately represents how humans interact with search systems and the potential dynamism of information needs during a search session, and also ignores important factors such as task context and information use (Bates, 1989; Ingwersen and Järvelin, 2005). Information seeking tasks are also tackled over time across many search episodes (Kotov et al., 2011; Agichtein et al., 2012) using disparate resources, and is often interwoven with other activities, meaning that searchers may be engaged in multiple information-seeking tasks simultaneously (Spink et al., 2006).

Interactive information retrieval (IIR) focuses on how people use search systems to retrieve information from the indexed corpus (i.e., the focus is largely on information *finding*) (Ruthven, 2008). Human-computer information retrieval (HCIR; Marchionini, 2006) has emerged as an important subdiscipline focused on the role of searcher and their context on the search process. Information seeking depends on the cognitive representation of a system's features, which is largely determined by the conceptual model that system designers provide through the search interfaces that they develop. Other determinants of successful searching include searchers' knowledge of the task domain (i.e., their domain or subject-matter expertise), the information-seeking experience, and the physical setting (Marchionini and Shneiderman, 1988). To support more effective search interactions, research in HCIR leverages advancements in search user interface technology and an improved understanding of people's search strategies developed by information scientists and others, via studies of library patrons and reference librarians, but more recently via large-scale surveys, panels, and log analyses.

To help searchers, interactive search systems offer support such as relevance feedback (RF), information visualizations, and query suggestions. RF (cf. Salton and Buckley, 1990) allows searchers to provide implicit or explicit feedback about relevant information and uses these judgments to enhance subsequent searches. Information visualizations (e.g., Card et al., 1999) use graphical techniques to visually represent large-scale collections of nonnumerical information, and help searchers attain new insights in support of decision making or other related complex mental activities. Visualization techniques such as highlighting salient content in documents can also capitalize on human abilities to recognize important cues and content when presented with them, rather than needing to examine information items in detail (Ahn and Brusilovsky, 2013; Lehmann et al., 2010). Query suggestions (also known as "related searches") (e.g., Efthimiadis, 1996; Koenemann and Belkin, 1996; Jones et al., 2006) provide recommendations about query terms to add to an existing query, or specific queries to issue, that assist searchers with the challenging process of query formulation.

Oddy (1977) and Belkin et al. (1982), among others, questioned the requirement for searchers to represent their information needs in a query understandable by the system. Indeed, systems such as *I³R* (Croft and Thompson, 1987), *Bead* (Chalmers and Chitson, 1992), and the *Ostensive Browser* (Campbell and Van Rijsbergen, 1996) offer "query-less" interfaces, where searcher needs are conveyed by means of examples from their browse behavior rather than textual descriptions. Research on implicit feedback (Joachims et al., 2005; Kelly and Belkin, 2004; Kelly and Teevan, 2003; White et al., 2005b) has shown that interaction behavior (mainly document retention activities such

as saving, bookmarking, and printing, as well as SERP click-through) can be used to build enhanced representations of information needs for use in query refinement or future retrieval operations. The notion of implicit feedback has been extended to Web search interaction, where clicks on search results can be used as signals of user interest that have utility as features of machine-learned models for general result ranking (Joachims, 2002; Agichtein et al., 2006) or personalized result ranking (Dou et al., 2007; Teevan et al., 2011b).

1.2 Next-Generation Search Interaction

The lookup model assumes that searchers are engaged in simple tasks with known information objectives. For many tasks, information needs are dynamic and only emerge as a result of learning through reflection regarding the information acquired during the search process (Bates, 1998; Marchionini, 2006). Studies of information behavior also focus primarily on user studies as a way of learning about information behavior. Over the past two decades we have witnessed an emergence of studies of human search behavior at Web scale. Studied search behaviors range from individual query statements (Silverstein et al., 1999; Jansen et al., 2000) to fine-grained interactions such as eye gaze and mouse cursor movements (Huang et al., 2011; Buscher et al., 2012). The focus has shifted to search tasks (Jones and Klinkner, 2008) and to trails that extend beyond just search engine interactions (White and Drucker, 2007). Log analyses of this nature have largely been restricted to researchers in industrial settings given restrictions on the sharing of large quantities of anonymized usage data that is required to make this type of research on naturalistic search behavior feasible. To pursue studies of this nature, academic researchers have needed to be creative in obtaining access to data, either using data collected from university websites (e.g., Joachims [2002] used data from arXiv.org, a service run by Cornell University), creating their own rich logging infrastructure (Lagun et al., 2011; Guo et al., 2013), or even devising games to help collect data from human participants (Ageev et al., 2011; West et al., 2012). In Chapter 12, I will discuss the many issues surrounding access to large-scale behavioral data, and its impact on the scope and the nature of the research that is possible in academia and industry.

New directions in search interaction are being pursued in parallel with the development of specialized search systems. Some interaction methods may be appropriate for particular scenarios more than others. Examples of this include search engine verticals such as image search, in which “infinite scrolling” (Farago et al., 2010) has replaced pagination as the predominant mechanism to explore deep into sets of retrieved items. Activities involving exploratory data analysis may benefit from views of the data such as *TreeMaps* (Shneiderman, 1992) or a closer (dynamic) coupling between queries and retrieved results (Ahlberg and Shneiderman, 1992). Next-generation search systems will recommend appropriate tools for the current search task, as well as allow searchers to control and customize search support to make the human-machine roles during searching more cooperative (Marchionini, 2006b). In offering any new functionality, appropriate consideration must be given to searchers’ desire to focus on task completion rather than learning new systems (Carroll and Rosson, 1987); in other words, additional functionality must be intuitive to use and not detract from the task completion process.

As searcher expectations for the availability of search support evolve, next-generation search systems will need to handle more complex tasks and the sophisticated search scenarios (Wilson et al., 2010). These systems will require capabilities to detect and model search tasks from search activity (Lucchese et al., 2013) and to track these tasks over time (Kotov et al., 2011). The following is an example of a typical search situation that highlights how the capabilities in next-generation search systems will help people learn and take action in the world.

Example search scenario: Elizabeth is undergoing treatment for breast cancer. Evidence of Elizabeth's health concern was apparent in her online behavior months before the first oncologist visit. Queries were seen for [small lump in breast] and other symptoms. Elizabeth's intelligent personal assistant associated her long-term activity (e.g., queries, Web page visits, social media posts, physiological signals from her wearable device), with an above average risk of breast cancer. It did this using a diagnostic model trained from the long-term activities of populations of consenting individuals. Model reliability was verified in stringent testing by medical professionals and extensive trials with breast cancer patients. The model was run against a log comprising Elizabeth's long-term activity and a personalized alert was generated about her risk level and the importance of swift medical attention in her case. Thankfully, she followed the agent's advice. A diagnosis was professionally rendered, Elizabeth updated her secure, cloud-based profile – linked by choice to her electronic health records – and used applications on her mobile device to track and monitor her course of treatment over time, as well as to connect with other breast cancer patients. This additional information is used by the search system to tailor the results that Elizabeth encounters when searching about breast cancer, as well as to find additional information proactively, such as the outcomes of recent clinical trials. In addition, the system helps her understand the nature of complex content, guides her through the stages of treatment, and aid decision making, for example selecting between treatment alternatives based on combining professional medical advice and content mined and aggregated from sources on the Web.

The preceding example depicts a situation in which search systems can play a more central role in rendering assistance proactively as well as reacting to searcher demands. Many of the technologies described in this example are available in some form in present-day computer systems. However, in current search systems, much of the responsibility for the searching, tracking, and sensing is placed on the searcher; I envisage that much of this burden will be shifted to the system over time. Access to this data allows systems to custom tailor the search experience, monitor people's task progress, and identify noteworthy patterns, trends, and anomalies. As is clear from the example, search systems need to be able to support not only information finding but also topic learning (Marchionini, 2006a). To do this effectively, the system also needs to be able to understand a searcher's level of domain knowledge – perhaps through explicit solicitation of this information, or studying their online search behavior

(White et al., 2009). There are many such complex scenarios where search systems could offer quite significant help.

Long-running search tasks such as that highlighted in the preceding example are common (Kotov et al., 2011; Agichtein et al., 2012). Agichtein et al. (2012) demonstrated that over 50% of search tasks observed in Web search engine log data transcend multiple search sessions (although decoupling active tasks from persistent interests remains a challenge). Support can be provided to help people store and rehydrate search tasks across long-running search episodes (Morris and Horvitz, 2007; Donato et al., 2010; Morris et al., 2011), as well as providing direct support for task resumption by considering “standing queries” that represent persistent interests. Standing queries are useful for information filtering (Allan, 1996) or cases where people continue the same tasks over time (Morris and Horvitz, 2007). Task states, stored on the client or (preferably for accessibility) in the cloud, allow search services to use this information to cluster search histories, support the selective application of these histories for applications such as personalization, and to suggest future actions based on what the searcher has accomplished thus far.

Next-generation search systems will also become more personal and more ubiquitous. They will learn about people’s activity beyond information seeking, enabling richer models and personalized search experiences tailored to individuals’ preferences. Intelligent assistants – such as Siri, Microsoft Cortana, and Google Now – can communicate directly with searchers and adapt their offerings to their users via access to personal content such as calendars, e-mails, and contact information.

Slow search (Teevan et al., 2013) describes a situation where information retrieval is unconstrained by time. Searchers may have an expectation that search systems respond quickly to their requests. However, in many search scenarios (such as long-term tasks, missions – e.g., chronic medical conditions for the searcher or a family member such as that highlighted in the earlier example) the search task may span a significant amount of time and the search system may have longer to reply to the original request and alert searchers if new information is found (e.g., notify those affected by chronic medical conditions that there is a new treatment option or study finding). If there is this additional time, the system could use that to survey available information (Lieberman, 1995; Joachims et al., 1997) or find humans to help (Horowitz and Kamvar, 2010; White et al., 2011). Search is not performed in isolation, and the social context is important in analyzing the value of the answers found, but can also have direct utility in helping people reach their goals, for example by broadcasting questions to members of the searcher’s social network (Morris et al., 2010).

Wearable devices allow richer sensing of the searcher (e.g., their physiological functions; Healy and Picard, 1998; Feild et al., 2010) and contextual factors, including their intentions and even the physical environment within which they are searching. Better modeling of search situations helps search systems determine task-appropriate information so as to offer tailored support to searchers. More expressive interaction methods involving gesture, touch, and speech are becoming commonplace in computer systems generally,¹ and it is likely that there will be widespread adoption of such methods in search systems over the next decade and beyond. Richer interfaces to support exploration and learning will allow searchers to discover relationships between items and make more informed decisions about both search strategies and information

use. Coupling information visualization techniques with natural interaction methods will create compelling and engaging immersive search experiences to promote effective discovery and learning.

An important theme in this book is the development and application of methods to model and support searchers *at scale*. Historically, the focus in search interaction has been on sophisticated and dynamic direct manipulation interfaces applied *client-side* within a restricted domain such as movies (e.g., Williamson and Shneiderman, 1992; Card et al., 1999) or a searcher's personal data, or simple methods such as ranked lists of surrogates applicable *server-side* on the Web but with only crude interaction support for text queries and hyperlink selections. High data volumes and bandwidth concerns have meant that only limited information can be shared with remote systems, and massive processing capabilities are needed to handle the exabytes of data collected by online services. Improvements in network bandwidth and availability, reductions in the price of hardware (especially data storage), and algorithmic and technological advances in distributed computing are finally allowing providers of next-generation search systems to offer rich and engaging search experiences to billions of searchers.

Mining the behavior of large numbers of searchers has utility for both improving the system performance and yielding insights that benefit society (e.g., via initiatives in public health; White et al., 2013b, 2014a). The ability to monitor search interactions at scale enables the study of information-seeking behavior in naturalistic settings across a broad range of information needs. Query and click-through statistics can be used to identify relevant results based on dominant search intents mined from many searchers (Joachims, 2002; Agichtein et al., 2006; Anick, 2003), and also at the individual level for applications such as the personalization of search results or the generation of query suggestions (Teevan et al., 2011b, Shokouhi, 2013). Log analyses may facilitate more complete modeling of information behavior, but it often lacks data about the rationales behind the observed search activity (i.e., what actions were taken, but not why the search is necessary). Mixed-methods approaches are necessary to develop a more complete picture of people's intentions and situations when searching. There are also drawbacks and challenges associated with collecting and using these data (including the important issue of searcher privacy), and care needs to be taken in the design of logging schemes, but also in detecting and mitigating potential biases in the data collected (including demand characteristics associated with participant awareness of the experimental setting; Orne, 1962). As an alternative to tracking attention within existing search interfaces, others have deliberately designed search interfaces to facilitate the capture of fine-grained implicit feedback (White et al., 2005a). This can be advantageous since signals can be accurately interpreted as evidence for particular aspects of documents, but such interfaces may also be unfamiliar to searchers.

Beyond the search process, there are other important issues to consider. Search engines are often part of commercial enterprises that receive revenue from online advertising, primarily via advertisements juxtaposed with search results. This has implications for how search data are employed and shared. The data that they collect is valuable for improving search engine services and attracting searchers, and more searchers means more revenue. Privacy considerations and the high business value of the data mean that it is not distributed broadly. Limited access to large-scale behavioral data for the research community is also necessary given appropriate consideration of important issues such as privacy and consent to share data for research purposes.

Proposals are emerging to release samples of the queries and clicks collected by commercial search engines in a manner that rigorously guarantees searcher privacy (Korolova et al., 2009).

As people own a broader range of devices (e.g., desktop computers, smartphones, laptops, gaming consoles, etc.), support for searching across those devices is becoming more important. Research in cross-device search is in its infancy, and there is significant opportunity to help multi-device searchers be more productive by supporting information-seeking tasks that span devices. (e.g., using searcher behavior observed on one search engine vertical to personalize the experience for the same searcher on another vertical) and cross-application (e.g., applying models learned from a person's usage of one application to tailor their search experience on another) scenarios will also grow in importance (Elkahky et al., 2015). Interest profiles developed from people's interactions within one or more domains or applications can help those same searchers when searching in other domains or using other applications. A core component of this support will involve cloud-based user profiles, accessible at any time from any device or application. These profiles can contain search histories but also useful information from other sources, such as authored documents and social media activity (with searcher consent as necessary). This would directly support cross-device/domain/application search scenarios, including supporting task resumption by rehydrating prior context (Wang et al., 2013). People's tasks are not limited to device boundaries, and increasingly search engines will consider the device-appropriateness of the information that they surface. They will support search activities across devices, including predicting whether the searcher will resume a task on another device (Wang et al., 2013; Montañez et al., 2014). If a service could predict the next device that a person will use (e.g., it can be confident that they are going to transition from a laptop to a smartphone), how long it has until that transition occurs, and the topic of the next search, it can employ a range of services, including crowdsourcing, to leverage the downtime between devices and find the best information in advance of the searcher resuming their task on the destination device.

As the capabilities of search systems improve to support richer search tasks in a broader array of search settings, there is also a need to evaluate search systems from a range of different perspectives. Evaluation metrics and methodologies need to consider search *process* in addition to the outcomes such as relevance and search satisfaction. Process-based evaluation considers factors such as searcher learning (Jansen et al., 2009; Wilson and Wilson, 2013) and cognitive load (Gwizdka, 2010) as part of system evaluation. Monitoring these signals is important in measuring the performance of next-generation search systems and creating benchmarks such as TREC that can facilitate the comparison of search systems across experimental sites. To spur progress in the design of interactive search systems, other experimental methods such as searcher simulations (especially those derived directly from behavioral data; White et al., 2005b; White, 2010) and facility for large-scale flighting of new interface developments could well be made available to researchers in *both* industry and academia through partnerships with commercial search engines and related efforts.

Interactions with search systems play an important role in the daily activities of many people. From the brief overview offered in the book thus far, it is clear that there are many opportunities for next-generation search interaction. I will discuss many of these in more detail as we move together through the book.

1.3 Outline

The remainder of the book is divided into the following four parts:

1. **Part I (Background):** I set the scene for the topics covered in this book. In Chapter 2 I discuss how search interaction is represented and modeled presently in search systems and how that support could be expanded in the future with the advent of a broad range of new technologies such as gaze tracking, touch interaction, and affective feedback. Chapter 3 describes how searcher interests and intentions are modeled, including models of interaction with next-generation search systems and the components necessary to build and deploy such models in practice. Toward better understanding searcher needs, Chapter 4 describes a range of models and frameworks that have been proposed to characterize different aspects of information seeking (i.e., information exploration, seeking information, gathering and organizing information, and applying information in practice).
2. **Part II (System Support):** I describe various types of system support for search systems, including the design of interfaces to help people locate information items, but also help people explore, learn, and apply gained knowledge. Chapter 5 describes both existing and emerging technologies to help people search more effectively. Of particular interest, as mainstream search support expands beyond its existing boundaries, is support for exploration and discovery, including recommendations (Chapter 6) and learning and information use (Chapter 7). I describe how exploration and discovery can be modeled and supported by systems. As part of this, proactive search and recommendations by anticipatory services such as Google Now and Microsoft Cortana are useful to push relevant content to searchers at opportune moments. I would expect that such anticipatory services become more central, especially as they are integrated into the operating systems of desktop and mobile devices. Chapter 8 considers the role of others both in supporting an individual searcher and in searching collaboratively as a team. Other people are frequently stakeholders in search tasks, and system designers need to consider how to integrate the viewpoints and preferences into the search process, as well as how to support coordination of information seeking across individuals and teams, so as to reduce redundancy and improve overall search effectiveness. In addition, search tasks do not occur in isolation from searchers' experiences and the surrounding context. Personalization and contextualization are important in improving the quality of search engine responses to the current situation, but methods are also needed should personal content be unavailable (e.g., using data from searchers with similar interests). I expect to see significant advances given richer sensing through wearable devices and more complete context-modeling methods. Chapter 9 discusses research in these critical areas and how it can be applied to improve future search systems.
3. **Part III (Evaluation):** An important aspect of the search process is the evaluation of the search systems. Chapters 10 and 11 discuss the evaluation of these systems, in particular the formative and summative methods that can be employed to inform system design. Given the richness of next-generation search interactions, I pay particular attention to better evaluating the search *process*, and also highlight