Jonathan Schatz — University of British Columbia
Nadia Stennes-Spidahl — University of British Columbia
Samantha Mills — University of British Columbia
Aaron J. Loehrlein — University of British Columbia

Bibliographic Induction: How KO Systems Optimize Browsing by Supporting Library Users' Prior Knowledge

Abstract

We investigate category-based induction as an aspect of browsing a library collection. Category-based induction is one of the primary uses of categories that are stored in memory. Knowledge organizing systems represent concepts in broadly the same way as models of category-based induction. Accordingly, it is reasonable to suppose that knowledge organizing systems facilitate category-based inductions about the collections that they organize. The processes of familiarization and differentiation are key aspects of browsing (Ellis 1989). Intuitively, these approaches appear to involve category-based induction in a bibliographic context. By examining induction, we hope to shed new light on the role of knowledge organizing systems in shaping browsing behavior. We also seek to investigate the viability of using inductive confidence as a dependent variable in assessing the utility of a KOS. A system that supports induction is potentially of great benefit to people seeking to browse a collection, whether the collection exists virtually or is part of a library's physical stacks.

Introduction

This paper explores the theoretical viability of a cognitive process that users of knowledge organization systems (KOS) may employ during interaction in order to evaluate the KO systems themselves. Specifically, we investigate whether users of a KO system are able to transfer what they know about a given section of a KO system to other, associated sections of the system. For example, if a researcher learns through experience that books that occupy a particular class generally have a particular property in common (perhaps a certain tone, bias, or subtopic that frequently appears among these books), how likely is the researcher to conclude that the same property may be found in books that occupy similar classes? In cognitive psychology, this process is referred to as a *category-based inductive inference*, since it involves a person mentally transferring knowledge from one category to another. We refer to this process simply as *induction*. We argue that inductions help users optimize their browsing by allowing them to take shortcuts to relevant materials by leveraging their prior knowledge. We are interested in testing such judgments as a metric to help evaluate the usability of two KO systems: Library of Congress Subject Headings (LCSH) and Library of Congress Classification (LCC), as they are used to organize the collections of the University of British Columbia Libraries. These well-known systems are then compared against a third system that is not a KOS but is the predominant IR system encountered in digital library interfaces: the Google-like "single search box" provisioned by modern library OPACs for keyword searching.

We will begin by first reviewing the literature in cognitive science, providing a brief history of models of induction and the utility of inductive inferences, as well as situating our approach within an established line of inquiry. The literature on browsing behavior in library and information science (LIS) is explored, particularly regarding a model of browsing that closely corroborates the findings of our exploratory interviews.

This will lead into a brief review of research in knowledge organization that discusses how KO systems should be evaluated or better integrated with user behaviors such as browsing. Finally, we discuss hypothetical scenarios in which inductive inferences occur and provide an argument for the utility of inductive inferences for researchers in both physical and digital library settings.

Induction in Cognitive Science

According to Holland, Holyoak, Nisbett, and Thagard (1986, 1), the term *induction* refers to "all inferential processes that expand knowledge in the face of uncertainty." This paper is specifically concerned with category-based induction, in which an inference is made about a category by bringing to bear knowledge about other categories. The process is also known as *property transfer*, since it is the process of mentally transferring a property from one category to another. Rips (1975) conducted an investigation into category-based induction in which participants were provided with two categories: a *given category* and a *target category*, both of which have a common superordinate. Participants were told that a property was true of the given category and were asked to estimate the proportion of the target category for which the property was also true. For example, participants might be told that robins are susceptible to a new type of blood disease and asked to estimate the proportion of eagles that are susceptible to that disease. Rips found that, as the given category became more typical of the superordinate, the estimates of the proportion of the target category tended to be higher (see Figure 1).

Figure 1. Inductions between robin and eagle

Subsequent studies revised Rips' model by asking participants to estimate their *inductive confidence*, which is their confidence that a conclusion holds, given a stated premise (e.g., Coley, Medin, and Atran 1997). Osherson, Smith, Wilkie, López, and Shafir (1990) developed a model that explained category-based induction as being entirely a product of similarity. According to the model, a person's confidence that a property will transfer from the given category to the target category depends on the similarity between a) the given category and the target category and b) the given category and the superordinate category. For example, if a person were told that robins were susceptible to a blood disease and were asked to estimate their confidence that eagles were also susceptible to that disease, the person would consider two factors: the similarity between Robin and Eagle and the similarity between Robin and Bird.

Although this model is highly robust, there are many aspects of induction that it does not take into account. For example, an induction from Whale to Tuna will be made with

more confidence if the property if relevant to aquatic behavior than if the property is relevant to anatomy (Heit and Rubinstein 1994). Lassaline (1996) found situations in which it was possible to manipulate similarity without manipulating inductive confidence, and vice versa. Papadopoulos, Hayes, and Newell (2011) suggest that category-based inductions are made with less confidence when the categories lack internal coherence, when the person has little experience with the categories, or when exemplars of categories are difficult to produce.

Nevertheless, the basic premise of category-based induction, where a property is transferred from a given category to a target category, is widely accepted and continues to be investigated. Coley et al. (1997) used inductive inferences to identify basic-level categories across dissimilar cultures. Griffiths, Hayes, and Newell (2012) found that people were more likely to make category-based inductive inferences after they had been trained to think of items as being members of categories. Research in cognitive psychology has revealed category-based induction to be one of the major benefits of holding categories in memory. However, to date there has not been any research into category-based induction where the categories are maintained by KO systems.

It is important to note that few studies of inductive inferences specifically investigate their accuracy. For the most part, research into induction is concerned with the factors that influence whether inductive inferences are made with confidence, not whether those inferences are correct. As an illustrative example, a person might be given the information that tuna prefer to feed at night. That person might then make a category-based induction, albeit reasoning poorly, and conclude that sharks, being similar to tuna, are also likely to feed at night. As a result of this inference, the person goes swimming in the afternoon without taking precautions against sharks and gets attacked by a shark. Although the person has made an induction, the induction proved to be incorrect. By the same token, if knowledge organizing system facilitated inductive inferences that were incorrect, it would be of dubious benefit to the user. However, in order to determine whether inductive inferences are being made correctly, it is first necessary to determine the types of the inductive inferences that are being made, and the levels of confidence at which they are made.

Browsing in Library and Information Science
Studies in LIS have extensively investigated browsing behavior (Herner 1970; O'Connor 1993; Rice et al. 2001) but the concept appears highly variable and resistant to most specification. In a recent approach, Bates (2007) challenges the conception in Rice et al. and others that browsing is only visual scanning of materials with a purpose, arguing for more distinction in definition. She situates browsing as a strategy within a larger exploratory search episode that "evolves" (in terms of refining a query) as new information is learned.

Many browsing studies focus on browsing in digital interfaces which may or may not involve interaction with a KO system. A recent example of a study of decisions made during physical browsing is Chuttur (2011). Users browsed a video collection organized alphabetically (except for LCC-organized documentaries) and had difficulty determining relevant videos. Chuttur theorized this stemmed from limited information users had from video covers to determine relevance (e.g., a video cannot be thumbed through). Acknowledging that relevance is highly subjective to users, and with only

superficial, limited information to consider, inductive inferences can help support reasoning about relevance by lending further implicit information. A KO system that supports inductive inferences may allow users to quickly gain an impression of the work and its associative similarity to other works.

Browsing models in the LIS literature also tend to analyze the browsing process at a high level of abstraction. For our investigation, a more granular focus was needed on an element of browsing in which users may draw on prior knowledge to evaluate new materials. In a widely cited study of the browsing behaviors of university researchers, Ellis (1989) presented a model of a search episode. He characterizes the *Browsing* stage of his model as comprised of two elements: *familiarization* and *differentiation*. Familiarization, in Ellis' model, describes the process of getting to know an area of interest in order to evaluate and apply prior knowledge to new sources in that area. Differentiation describes the side-by-side evaluation of sources. Differentiation is defined as "...employing differences in the nature of the source materials to filter material" (177), to assess the "differing probability of their containing useful material" (190). Ellis' concept of differentiation in the research process is quite similar to induction in the context of a user's interaction with library KOS.

Differentiation and familiarization, as Ellis describes them, are useful in conceptualizing the at-a-glance evaluation element of inductive inference that we explore. Ellis makes it clear that prior experience, emerging from the interplay of factors such as the library material's level of technicality, approach to the topic, and intended audience, makes differentiation possible. Although he does not use the term, the act of using prior experience to fill in a gap in one's knowledge is essentially an inductive inference. Ellis' study is notable in being derived from researchers reporting their own behavior. Similarly, our work in induction seeks to simulate real research conditions as closely as possible. Not every aspect of Ellis's work is generalizable, but a behavioral model of browsing that accounts for differentiation is applicable to the design of IR systems in general, and KO systems in particular.

Ellis' differentiation process has been employed in more recent browsing studies. In 2008, Bilal, Sarangthem and Bachir proposed a model of children's browsing behavior which incorporated Ellis' Browsing stage. Notably, Bilal et al. disqualified several other elements of Ellis' model, but found differentiation applicable in this vastly different user group (Arabic-speaking children, cf. English-speaking social scientists). Bilal et al. focused on digital browsing, and the applicability of differentiation in their model supports our belief that induction and related processes occur similarly in both physical and digital browsing environments. Other studies include those concerned with building models of information-seeking and information literacy (Timmers and Glas 2010). However, while these may include descriptions of user evaluation of information sources or differentiation between them (Joseph, Debowski and Goldsmith 2013), little attention is directed to the process by which prior knowledge enters into evaluation. This is an area we seek to expand upon.

Shelf browsing continues to be an important element of library research—and thus continues to be a valuable process to support and improve. Whitmire's (2001) longitudinal study surveyed a cohort of university students regarding their library activities. "Found materials by browsing in stacks" (383) was an activity that showed consistent correlation between years of study, the strongest of which was between first

and third years. Extending the investigation of induction to digital browsing may well be possible, but even if we never add another book to the collection, library users still rely on extant physical stacks. Studying cognitive processes involved could lead to refinements of library organization systems in the future. For example, Clark (2012) reports on a successful program at Kent State, where print placeholders for electronic journals were shelved to aid browsing and enhance comprehension of the structure of electronic serials organization. This was to mitigate the confusion from duplicate entries for print and electronic serials and to facilitate "discovery," and was well-received by faculty and students.

Furthermore, Massis (2011) has reviewed an ongoing debate in academic libraries regarding physical and virtual browsing. Significant protests often occur when shelf space is compressed or removed, showing that users are attached to physical browsing and perhaps also to the possibility of serendipitous discovery can occur within it.. Massis adds to reports from academic library users who maintain that the increase in compact shelving interrupts the browsing process. Indeed this concept of serendipity is often cited as a reason to maintain some sort of surrogate for the physical shelf in new digital library systems. One way of looking at the cognitive investigation we aim to undertake is that we are trying to understand the serendipity that comes from browsing the shelves, in order to discover which knowledge organization conditions best support strong and useful serendipitous discoveries by bringing together clusters of useful resources, or "hot spots", in the library.

The phenomenon of users in academic libraries freely browsing the stacks is relatively young, historically, dating back only to the middle of the twentieth century (Barclay 2010; Massis 2011). While Massis acknowledges the perspectives of those who claim that browsing electronic surrogates or call number lists can be just as effective, he reports that "serendipitous browsing remains today, an integral practice among academic library faculty and researchers who frequently report having discovered important secondary titles on the shelves adjacent to those they were originally seeking. Such finds result in providing added value to their course content by virtue of serendipitous browsing" (180). These arguments reveal that LIS has not yet established how to achieve equally compelling and valuable experiences for users in digital browsing. We hope that our investigation will go some way toward uncovering how different KOS factor in this regard.

Evaluation of KO systems

We are not aware of studies that investigate ways of evaluating a KO system's latent ability to support inductions during the browsing process. Some traditional KO systems, such as LCSH, have had their structure visualized to better support navigation and been evaluated after enhancement (Julien et. al 2012). Andersen (2004) asks how aspects of IR systems should be evaluated. He recommends that researchers ought to reconsider Swanson's (1977) suggestion "that systems be evaluated according degree to which they facilitate trial-and-error nature of IR." That is, the flexibility of the system to support the multifarious character of user's search episodes could be a reasonable criterion. Because a KO system that supports confident inductive inferences may help users optimize their use of library IR systems *and* their interaction with a

given library KO system itself—in the mind, on the OPAC and at the shelf—our criterion for evaluation may be helpful in minimizing the error portion of IR use.

Hjørland (2013) has contested the notion that prevalent KO systems have been or should be designed from the ground up based on user data or cognitive models. We agree that cognitive models are not well suited to the construction of KO systems. However, this limitation may be due to the fact that cognitive models are typically designed to predict aspects of people's behavior. For that reason, cognitive models could be used to study how people use existing KO systems. Methods that measure the strength of inductive inferences made between sections of a KOS are one example of this approach. The qualities by which KO systems support inductive inferences are unclear. However, Jacob (2004) describes four general approaches to the organization and retrieval of items: free-text searching, classification, postcoordinate indexing, and precoordinate indexing. Classification schemes typically have a rigid hierarchy, which allows the relationships between classes to bear information. In contrast, precoordinate indexing (e.g., subject heading systems) is frequently "unprincipled, unsystematic and polyhierarchical" (536), while neither postcoordinate indexing nor free-text searching methods organize according to any principle other than query matching. Based on this analysis, it seems reasonable to conclude that classification schemes provide the most support for inductive inferences, while postcoordinate indexing and free-text searching approaches provide the least support.

Two KO systems widespread in North American academic libraries, and in use at UBC, Library of Congress Classification and Library of Congress Subject Headings are evaluated for their ability to support inductive inferences. While LCC prescribes a shelf order for related items, both LCSH and general keyword searches also bring together potentially relevant items as a results list generated by querying the library's catalogue. We regard the proprietary relevance ranking algorithm of the OPAC as another system to evaluate alongside LCC and LCSH for its capacity to facilitate inductive inferences between sets of items brought together. This will be referred to as "keyword search" henceforth. These KO systems, are experienced as overlapping retrieval conditions by most users (save for LCC unless a rare "call number browse" is initiated), vary in how they bring together sections in the library. Our hypothesis is that these systems vary in the degree to which they support inductions amongst groups of bibliographic items they organize. One system may prove demonstrably better at supporting inductive inferences on account of exhibiting more internal systematicity—linking items by certain consistent principles, as Jacob (2004) has outlined. We propose that if a particular retrieval condition better promotes inductive inference between its groups or classes, then that condition should be a better tool for assisting researchers in finding useful library materials.

Exploratory Interviews
In order to gain insight into how browsing the stacks fits into library research activities—and potentially situations where inductive inferences are employed—we conducted exploratory interviews with eleven graduate students at the University of British Columbia. Participants were selected based on their enrollment in masters and doctoral programs at UBC but-outside of library and information studies. Graduate students were sought to ensure that participants were likely to be performing

independent research. LIS students were excluded, since they are likely to be more familiar with library organization schemes than the general student population.

The participants were prompted to discuss their strategies for browsing the library's stacks as well as their general awareness of, and feelings about, the library's organizational systems. The interviews were transcribed to documents if they had been tape recorded. After compiling the interviews, coding was undertaken as a group to reveal any commonalities in the habits and notions related by participants. As expected, many students did not profess (or demonstrate) more than a passing understanding of the Library of Congress Classification system (LCC), although they generally understood that the library was organized by discipline and then by topics within disciplines. Most students were aware of subtopics in their field having a coherent organization or structure. For example, one student noted that the section for books about film theory was subdivided differently than the section for books about specific film directors. However, call numbers were not well understood by any of the students interviewed, and a few were skeptical that call marks were non-arbitrary in meaning. Through repeated library use, participants generally acquired the habit of visiting specific areas of their campus branch where books seemed to be most relevant to their research interests.

A concordance of techniques for locating research materials emerged from our reading of the interviews: students chiefly began at the UBC Library web site and used the search box on the front-page to perform mostly keyword searches via either the discovery service or OPAC. After further refinements or reformulations of their query, and upon retrieving the record of a promising book, they entered the library to locate it on the shelves. Many related that they engage in habitual shelf browsing, and that it is a reasonably successful activity for them, in accordance with Whitmire's (2001) findings. Most reported also taking the opportunity to browse in the same area for other related books, particularly the areas immediately adjacent to the initially located "anchor book." We concluded that this "anchor strategy"—beginning at the OPAC and ending with shelf browsing to ensure other relevant items are not missed—largely characterized the search episodes of graduate students seeking print materials at UBC. Seeking further insight, we also conducted brief interviews with two reference librarians at UBC. These librarians generally found this anchor book workflow sensible and in agreement with their experience instructing graduate students and researchers. The librarians also noted that certain disciplines tended towards distinctive patterns and strategies of collections use. For example, users from the sciences browsed the library's print collections much less frequently except when they were researching interdisciplinary topics that required print books, such as materials related to the history of science or forestry from an anthropological perspective.

Keyword search, subject headings, and physical proximity (i.e., bibliographic classes) are experienced as information retrieval conditions which overlap in a real-world search episode. For instance, keyword searches may retrieve records based partially on their subject headings, and a user may head to the shelves to find a particular book returned by such a search and find similarly useful materials in adjacent classes. For ease of analysis, we evaluate each KO system separately in this study. However, each them of may be encountered at any part of a search episode in naturalistic library settings. According to the anchor strategy discussed above, a user

beginning with either a keyword search or a LCSH search will, if these queries prove fruitful, enter the physical library seeking a known item on the shelf with the aim of retrieving or further assessing it. Depending on how much material they seek and how useful the known item seemed, they may also notice physically proximal items and begin browsing. At the very least, they will take note of the area where the relevant item was shelved. In this way, regardless of the strategies that the library users employed, interaction with LCC was a common feature of search episodes focused on print materials via the OPAC. We seek to investigate user perceptions of each of these conditions in isolation from each other, in order to assess how each supports inductive confidence during the searching and browsing process.

Use Case Scenarios

Patrons can be served more effectively with a more rigorous understanding of the decisions they make as they browse a collection. Because inductive inferences are cognitive processes that can help optimize decision-making during browsing, if a particular KOS better enables users to make inductive inferences, there is a clear rationale for extending the use of that system in both online and physical presentations of resources. In particular, OPAC functionality may be modified to better leverage inductive inferences, enabling remote patrons to better assess resources. The data collected from our interviews imply that library patrons already make inductive inferences in their browsing. The researchers we spoke with discussed the strategies by which they navigate areas in the library about which they had little prior experience. For example, one participant tried to apply his experience of how a section of books on dramaturgy at UBC Library was organised when he sought information in film sections regarding directors and their works. Although this approach was not immediately applicable or useful for him, it suggests commonplace situations where such inferences are attempted. Based on these findings, we considered hypothetical scenarios in which library users would interact with a KO system, begin browsing and be asked to make inductive inferences in order to efficiently conduct their research into specific topics.

For example, a researcher of economics interested in comparing economic development in two different Latin American countries (e.g., Brazil and Bolivia) could potentially transfer her knowledge of how books presented information about Brazilian economic development to a set of books about Bolivian economic development. This approach would help her spend more time finding useful information and less time learning how to navigate the resources. She is already using what she knows in her at-a-glance selection of various resources. For example, the researcher might remember where she went for the books about Brazilian economic development, and scan the shelves around it for Bolivian economic development. She will use the physical proximity of the books to each other on the library shelves to aid her inferences about the information contained therein.

Another researcher might be studying the school experiences of lesbian, gay, bisexual and transgender (LGBT) students. She is interested specifically in finding books on the experiences of gay and lesbian students in single-sex education environments (e.g., all girls schools). Using the OPAC at her university library and recalling a subject heading that had previously yielded results, the researcher tries to retrieve books from a more specific and relevant heading such as "Single-sex classes

(Education)—Gay and Lesbian." However, she finds that there is no such subject heading. Many books with the broader subject heading "Single-sex classes" are retrieved but it is not clear whether these books discuss gay and lesbian students. Upon further examination of the results, she notices that the books tend to be clustered in a few physical shelf locations throughout the library. The researcher recalls that, six months previously, she had browsed the books in one of those locations and found that they were indeed useful to her research. She hypothesizes that the books in the other clusters are also likely to be useful. In this manner, the researcher makes an inductive inference from one set of books to another as a shortcut.

In many cases, inductive inferences can be more implicit in the evaluative decisions users make. For example, another fictional researcher performs a keyword search for the term "cyberterrorism" in the library OPAC. He notices that many of the books come from a section he had browsed six months ago. He also notices that many of the books are in a section of the library he has not browsed before. The researcher locates this new section and begins browsing. He is able to quickly ascertain that the books in the new section conform to a general structure that is very similar to the books in the section that he is familiar with. In so doing, he has made an inductive inference. This inference allows him to bypass the process of learning to navigate the books' layouts, writing styles, and treatment of major subtopics, allowing him to delve into the deeper issues covered.

Conclusion

Following from our exploratory interview findings and these hypothetical situations, we are developing a methodology to test the strength of inductive confidence made in situations when differentiating between items by physical proximity, subject headings, or keyword searches. The KO systems that gather resources in OPACs currently use combinations of keyword relevance and subject headings. However, if one of these systems is more supportive of user's inductive inferences, OPACs could be modified to better support use of that system. For example, grouping resources by subject headings (and exposing these heading and their relationships to the user) before offering keyword search could potentially support user's judgments better, if we find that subject headings better support inductive inferences.

Future organization of digital libraries should attempt to recover and preserve useful affordances encoded into physical library organization, where possible. The findings of our study may prove useful for the design of systems that organize only digital resources and permit different modes of analogous browsing within the system's interface. Recently, Burrows (2012) argued that integration of the Linked Open Data framework may provide additional avenues for browsing to challenge the recent reliance on Google-like "single search" boxes in IR interfaces. We believe current KOS systems may be worthy candidates for providing stronger browsing alternatives. If certain KOS better support inductive inferences and thus aid users in browsing more efficiently, digital libraries and future IR systems can deliver more usable systems by designing features that afford interaction with this KOS. Understanding current optimal use of library KOS and browsing can provide a foundation for this approach.

References

Andersen, J. 2004. *Analyzing the role of knowledge organization in scholarly communication: An inquiry into the intellectual foundation of knowledge organization*. Copenhagen: Department of Information Studies, Royal School of Library and Information Science.

Barclay, D. A. 2010. The myth of browsing: Academic library space in the age of facebook. *American Libraries* 41(6−7): 52−54.

Bates, M. J. 2007. What is browsing-really? A model drawing from behavioural science research. *Information Research* 12(4), Paper 330.

Bilal, D., Sarangthem, S., and I. Bachir. 2008. Toward a model of children's information seeking behavior in using digital libraries. In *Proceedings of the second international symposium on Information interaction in context.* pp. 145−151. ACM.

Burrows, T. 2012. A machine for browsing: Beyond the "single search box". *OCLC Systems & Services* 28(1): 56−58.

Clark, J. C. 2012. Connecting print patrons to electronic holdings with shelf dummies. *Serials Review* 38(4): 250−253.

Chuttur, M. Y. 2011. Exploratory study of video collection browsing behaviour of patrons in an academic library. *Malaysian Journal of Library & Information Science* 16(1): 137−149.

Coley, J. D., Medin, D. L., and S. Atran. 1997. Does rank have its privilege? Inductive inferences within folkbiological taxonomies. *Cognition* 64(1): 73−112.

Ellis, D. 1989. A behavioural approach to information retrieval system design. *The Journal of Documentation* 45(3): 171−212.

Griffiths, O., Hayes, B. K., and B. R. Newell. 2012. Feature-based versus category-based induction with uncertain categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 38(3): 576−595.

Heit, E., and J. Rubinstein. 1994. Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20(2): 411−422.

Herner, S. 1970. Browsing. In *Encyclopedia of library and information science.* (1st ed.) Vol. 3, (pp. 408−415). New York, NY: Marcel Dekker.

Hjørland, B. 2013. User-based and cognitive approaches to knowledge organization: A theoretical analysis of the research literature. *Knowledge Organization* 40(1): 11−27.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., and P. R. Thagard. 1986. *Induction: Processes of inference, learning, and discovery.* Cambridge, MA: MIT Press.

Julien, C.-A., Tirilly, P., Leide, J. E. and C. Guastavino. 2012, August 6-9. Using the LCSH hierarchy to browse a collection. Paper presented at *Categories, Contexts and Relations in Knowledge Organization: Proceedings of the Twelfth International ISKO Conference,* Mysore, India, pp. 109−166.

Jacob, E. K. 2004. Classification and categorization: A difference that makes a difference. *Library Trends* 52(3): 515−540.

Jorgensen, C., and E. D. Liddy. 1996. Information access or information anxiety?—An exploratory evaluation of book index. *Indexer* 1(20): 64−68.

Joseph, P., Debowski, S., and P. Goldschmidt. 2013. Models of information search: A comparative analysis. *Information Research* 18(1). paper 562

Lassaline, M. E. 1996. Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22(3): 754−770.

Massis, B. E. 2011. ""Serendipitous" browsing versus library space", *New Library World* 112(3/4): 178─182.

O'Connor, B. 1993. Browsing: a framework for seeking functional information. *Knowledge: Creativity, Diffusion, Utilization* 15(2): 211─232.

Osherson, D. N., Smith, E. E., Wilkie, O., López, A., and E. Shafir. 1990. Category-based induction. *Psychological Review* 97(2): 185─200.

Papadopoulos, C., Hayes, B. K., and B. R. Newell. 2011. Noncategorical approaches to feature prediction with uncertain categories. *Memory & Cognition* 39*: 304─318.

Rice, R. E., McCreadie, M. and S. L. Chang. 2001. *Accessing and browsing information and communication*. Cambridge, MA: MIT Press.

Rips, L. J. 1975. Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior* 14(6): 665─681.

Timmers, C. F. and C. A. W. Glas. 2010. Developing scales for information-seeking behaviour. *Journal of Documentation* 66(1): 46─69.

Whitmire, E. 2001. A longitudinal study of undergraduates' academic library experiences. *Journal of Academic Librarianship* 27(5): 379─385.