CITY ST GEORGE'S UNIVERSITY OF LONDON

City Research Online

City St George's, University of London

Citation: Hoeber, O., Pirmoradi, A., Mon, E. E. & Makri, S. (2026). Don't Break the Flow: Supporting Deferral and Review of Potentially Serendipitous Information in Library Search. In: UNSPECIFIED (pp. 231-241). New York, United States: ACM. ISBN 9798400724145 doi: 10.1145/3786304.3788846

This is the published version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: https://openaccess.city.ac.uk/id/eprint/37154/

Link to published version: https://doi.org/10.1145/3786304.3788846

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to City Research Online policy.

City Research Online: http://openaccess.city.ac.uk/ publications@citystgeorges.ac.uk

Latest updates: https://dl.acm.org/doi/10.1145/3786304.3788846

PDF Download
3786304.3788846.pdf
23 March 2026
Total Citations: 0
Total Downloads: 0

Published: 22 March 2026

Citation in BibTeX format

CHIIR '26: 2026 ACM SIGIR Conference on Human Information Interaction and Retrieval
March 22 - 26, 2026
Seattle, USA

Conference Sponsors:
SIGIR

RESEARCH-ARTICLE

Don't Break the Flow: Supporting Deferral and Review of Potentially Serendipitous Information in Library Search

ORLAND HOEBER, University of Regina, Regina, SK, Canada

ABBAS PIRMORADI, University of Regina, Regina, SK, Canada

EI EI MON, University of Regina, Regina, SK, Canada

STEPHANN MAKRI, City St George's, University of London, London, U.K.

Open Access Support provided by:

University of Regina

City St George's, University of London

CHIIR '26: Proceedings of the 2026 Conference on Human Information Interaction and Retrieval (March 2026)
https://doi.org/10.1145/3786304.3788846
ISBN: 9798400724145

Don’t Break the Flow: Supporting Deferral and Review of Potentially Serendipitous Information in Library Search

Orland Hoeber
University of Regina
Regina, Canada
orland.hoeber@uregina.ca

Abbas Pirmoradi
University of Regina
Regina, Canada
abbaspirmorady@uregina.ca

Ei Ei Mon
University of Regina
Regina, Canada
eem243@uregina.ca

Stephann Makri
City St George’s University of London
London, United Kingdom
stephann@citystgeorges.ac.uk

Abstract

Online digital library searchers who find themselves engaged in multiple concurrent complex search tasks may encounter potentially serendipitous information for one task while pursuing another. This can propel them in unexpected and exciting directions that enhance their knowledge. However, existing search interfaces provide limited support for searchers to handle such serendipitous information encounters, forcing them into an often disruptive choice between continuing their current search or pursuing the new discovery. To address this issue, we introduce Revelio: an academic digital library search interface based on a ‘save now, organize later’ workflow. Revelio provides a low-effort mechanism for deferring potentially serendipitous information encounters, and a semantic similarity approach for subsequently reviewing and saving such search results within a multi-workspace structure. We evaluated Revelio in a between-subjects controlled study with 28 participants using a novel experimental design that creates the opportunity for serendipitous information encounters within a concurrent complex academic search task context. Compared to those who used a baseline search interface, participants who used Revelio were able to effectively defer and subsequently review encountered information, resulting in positive opinions about the interface, increased performance in both prescribed search tasks, and greater perceived knowledge gain. These findings demonstrate the value of search interfaces that explicitly support the deferral and subsequent review of potentially serendipitous information.

CCS Concepts

• **Human-centered computing** → **User interface design**; **User studies**; • **Information systems** → *Search interfaces*; *Digital libraries and archives.*

Keywords

Serendipity, Information encountering, Digital libraries, Experimental design, Exploratory search

This work is licensed under a Creative Commons Attribution 4.0 International License.
CHIIR ’26, Seattle, WA, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2414-5/26/03
https://doi.org/10.1145/3786304.3788846

ACM Reference Format:
Orland Hoeber, Abbas Pirmoradi, Ei Ei Mon, and Stephann Makri. 2026. Don’t Break the Flow: Supporting Deferral and Review of Potentially Serendipitous Information in Library Search. In *2026 ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR ’26), March 22–26, 2026, Seattle, WA, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3786304.3788846

1 Introduction

Information seeking in academic digital libraries is rarely a linear single-query process; rather, it is a complex exploratory activity where searchers manage multiple concurrent information needs over extended periods [46]. To cope with this complexity, searchers adopt sophisticated strategies, such as tab-based task management and using search interface workspace features to organize information for their different interests [14]. Such multi-task environments create fertile ground for serendipitous discoveries, where valuable information for one task may be unexpectedly encountered while pursuing another [13, 31].

Serendipity in the context of finding useful information unexpectedly is termed *information encountering* [13]. During search, information encounters often occur when looking for information for one task but finding information on a different (often partly-related) task. Initially, information encounters trigger the *potential* for serendipity; they are only considered serendipitous if an association is made to a different search task of interest and if the searcher is able to create a valuable outcome from the encounter (e.g., by saving then later reviewing the encountered information). Serendipitous information encounters present a critical dilemma: the searcher must either disrupt their current cognitive flow to pursue the new find, or risk losing the potentially valuable information by deferring its examination [31, 32]. This moment of choice imposes a significant cognitive and interactive burden on the searcher [1]. Designing search interfaces that can effectively mitigate this disruption to support, rather than hinder, serendipitous information encountering remains a key challenge.

The search interface has long been a critical component of modern digital libraries, serving as the primary gateway to vast information repositories [5]. The effectiveness of these interfaces, however, is often limited by a design paradigm inherited from general web search engines [21, 40]. This model, optimized for simple fact-finding queries, fails to support aspects of academic work, such as managing concurrent complex search tasks, following exploratory search processes, and capitalizing on serendipitous information encounters. Consequently, after noticing potentially useful information related to a different search task than the one underway, the minimalist design of these interfaces offer little support for helping the searcher manage the serendipity process.

To make matters worse, existing in-browser support for deferring the examination of encountered information is insufficient; bookmarks are rarely created and even more rarely re-visited [4]. Similarly, while new tabs are commonly used to 'park' information to read later [8, 11, 18], people can forget which pages they have already read and lose their place in the search process when relying on this approach to deferral [8]. This highlights the need to better support people in deferring their examination, and later review, of potentially serendipitous information encounters.

This paper presents the features and an empirical evaluation of Revelio, a novel academic search interface designed to support the handling of information encounters. To evaluate Revelio in a controlled laboratory setting, we developed a novel experimental design wherein search results relevant to one task are artificially inserted into the search results of another task, creating consistent and repeatable opportunities for serendipity. This research was guided by the following research questions:

RQ1: How can a search interface be designed to support deferral and subsequent review of serendipitously encountered information while minimizing disruption of the primary search task?

RQ2: Does the inclusion of explicit support for the deferral and subsequent review of serendipitously encountered information change how people handle this information? If so, how?

RQ3: How does the search interface influence subjective and objective performance measures when serendipitous information encounters occur whilst undertaking concurrent search tasks?

2 Background and Related Work

Serendipity in the context of information interaction and retrieval has been referred to as *information encountering* [13]. Although information encounters can happen when not consciously looking for information (e.g., noticing a billboard ad for an interesting product you did not know existed), they often occur while undertaking information seeking activities [13]. This usually involves unexpectedly finding useful information related to one search task when looking for information on another task. These tasks may appear to be unrelated, but usually have some degree of conceptual overlap [13]. Information found unexpectedly is only *potentially* serendipitous when it is first encountered; it requires the searcher to *follow-up* on the encountered information (e.g., by reviewing it and connecting it to other information needs). This follow-up facilitates *value creation* from the information encounter [25], a process known as *materialization* [7]. Only once the value of the information has materialized can an information encounter be considered serendipitous [25]. This highlights the importance of following-up on encountered information to cultivate serendipity, either at the time of the encounter or later.

When it happens, serendipity is a happy and surprising secondary outcome of a search activity; it can propel information-seekers in new and exciting directions [26, 31] which, in turn, can result in substantial knowledge enrichment for seemingly little effort [7]. Cultivating serendipity is therefore highly desirable in interactive information retrieval systems.

Existing models of the serendipity process (e.g., [7, 25, 31]) have adopted different terminology but have converged on key aspects of the process. Firstly, something in the environment, such as a search result, *triggers* serendipity and sparks a mental *association* (often a bisociation where two or more 'matrices of thought' combine to create a new idea or synthesis [20]). When information is unexpectedly encountered on one topic when looking for information on another, the association is made between the encountered information and an information need not currently being pursued. For example, for a student interested in 'human-centered AI', a search on 'bias in machine learning models' might uncover a novel recommender system interface aimed at reducing ML bias on e-commerce platforms. Next, the person who experiences the association *follows-up* on it, such as by examining the encountered information immediately, or saving the search result or document to examine later. Finally, once they have examined the encountered information, they *create value* from it. For example, the student might include information about the recommender interface they encountered in a literature review being prepared for another course they are taking.

The information encountering process [13] complements the generalized serendipity process and helps to explain how serendipity can occur in the specific situation of finding information on one topic while searching for information on another. When someone *notices* a promising informational trigger, they must decide whether or not to *stop* or *suspend* their current search task to *follow-up* on the information they noticed. If they do, this may involve *examining* the information in detail or *seeking more information* to assess its usefulness. If they determine it to be useful, they then *use* the encountered information (e.g., by incorporating it into their work). Afterwards, they might decide to *resume* the search task they were undertaking pre-encounter. However, they might find task resumption difficult due to the cognitive effort involved in getting 'back into the flow' of their previous search task [27], including paying cognitive penalties associated with cross-session searching [24].

Upon serendipitously encountering information during one search task that may be useful for another, the searcher is faced with a dilemma: Do they follow-up immediately on the new information they have just discovered but risk losing their place in the search they were conducting when they encountered it? Or do they postpone examining the new information until later (e.g., once they have finished their current search session) but risk forgetting to review it at all, resulting in a potentially missed opportunity for serendipity? [27]. This dilemma can be particularly difficult, as the cost-benefit of immediate follow-up is usually not knowable at the time of the encounter [27]. Therefore, it is often only possible to decide whether following up on encountered information was worthwhile *after* having followed-up.

Adding workspaces to a search interface holds potential for supporting serendipitous information encounters. Early approaches focused on providing a single workspace in which searchers could add notes to what they found [23] or interactively evaluate the relationships among saved search results in the workspace [16]. Others have studied multi-workspace environments to support cross-session searching [14] and organizing found information [35].

While such approaches provide opportunities for supporting the follow-up and value creation aspects of the serendipity process, navigating between different workspaces distracts from the primary task of searching, and may have a high context-switching cost that remains unexplored in the search interface design literature.

Compared to the extensive body of literature on the process of serendipity, few have studied the phenomenon from the perspective of designing to facilitate it within an information-seeking context. Early work suggested that visualizing a series of queries and associated search results has the potential to support serendipity [3]. More recent research suggested that providing diversification of search results across different sources can foster serendipity [43], as can surfacing community-generated information such as tags [38]. A common pattern among these approaches is that the focus is on enabling the serendipity *trigger* only. To our knowledge, no approaches have been proposed to explicitly support the searcher as they follow-up and create value from serendipitous information encounters, nor have any considered how the search interface might mitigate the disruption to the primary search task that an encounter might cause.

3 Search Interface Design (RQ1)

To address RQ1, we introduce the design of a novel search interface, Revelio, which was previously presented as a concept demonstration [32] and was refined and updated as part of this research. The central design goal of Revelio is to facilitate deferral of the follow-up and value creation stages of the serendipity process, by allowing searchers to save potentially serendipitous search results to a relevant workspace for later review. Its aim is therefore not to spark serendipitous information encounters during search, but to support searchers in postponing the detailed examination of such encountered information, so as not to disrupt their primary search task. After a searcher encounters potentially serendipitous information (in Revelio's case a search result), it allows them to act on the initial association with minimal disruption, deferring the rest of the serendipitous process to a later review stage. During that review stage, the system provides cognitive support to help the searcher to *follow-up* with the connection by matching the search result to their other search interests, thereby enabling a *valuable outcome* of the serendipitous information encounter. It achieves this by separating the immediate, low-cost act of preserving an interesting find from the more cognitively demanding task of reviewing it and deciding where within a structured workspace to save it. This is implemented through a workflow we describe as 'save now, organize later.'

Revelio is built within a multi-workspace academic digital library search framework [14]. Workspaces are used to help searchers keep their different search interests distinct from one another, enabling meta-cognitive planning of the search activity [23]. Searchers have autonomy to create and name as many workspaces as needed. By choosing a workspace before issuing a query, all search activity is mapped to that workspace. Searchers can easily access these saved search results, as well as switch to a different workspace if they decide to pursue a different search task.

The 'save now, organize later' workflow in Revelio is supported by two key features: The first is a dedicated *Read it Later (RIL) button* provided for each search result within the search engine results page (SERP) (see Figure 1). When a searcher views the SERP and encounters an unexpected but potentially useful search result for another search task, they can click this RIL button ❶ to capture this serendipity trigger and defer the search result for later review. This single-click action adds the document to a temporary holding area without requiring the searcher to switch their context away from their primary search activity or make any further decisions. Doing so allows the searcher to maintain their focus on the current task, knowing the encountered information has been safely set aside for future review.

The second key feature is the Read it Later *list page*, which holds all search results deferred for later review using the RIL button. Access to this page is provided at the top of the search interface ❷, including a counter for how many search results have been saved to this temporary holding area. Upon navigating to the RIL list page (Figure 2), each search result is displayed ❸ along with its originating query and when it was saved ❹. To enable the searcher to follow-up on the encountered search result, a list of all available workspaces and visual indicators of the saved search result's textual similarity to these workspaces is provided beside each search result ❺. For each of the searcher's workspaces, the text from the workspace title and the titles and abstracts of all search results within it are aggregated.

The cosine similarity between the BERT (Bidirectional Encoder Representations from Transformers) [10] embedding of the search result and the aggregated embedding of the task-based workspace is represented visually, helping the searcher to quickly and confidently determine the most appropriate workspace destination for the temporarily saved search result. The search result can be moved to the appropriate workspace with a single click. This streamlined process supports the final step of moving the resource into its correct workspace, thus completing the 'save now, organize later' workflow, and encouraging serendipity to materialize by *creating value* from the information encounter. Within the context of recently published design principles for exploratory search interfaces [15], these features are lightweight additions to an established design pattern for academic search interfaces, that are visual, scrutable, interactive, and persistent.

The benefit of Revelio over other typical academic digital library search interfaces is that it supports the capture of encountered information in a way that allows the deferral of the rest of the serendipity process to a later review stage. Relating the workflow to the serendipity process [13, 31], when a serendipity *trigger* is noticed and a mental *association* made (identifying a search result that may be relevant to another search task), the result can be saved to the RIL list page and the searcher can continue with their primary search activity with the confidence that the encountered information can be easily found and reviewed later.

During examination of what has been saved to the RIL list page, the visual similarity match to other workspaces and single-click mechanism for moving the result to an appropriate workspace encourages later *follow-up*. If this follow-up occurs and is fruitful (i.e., the encountered information is deemed useful), the searcher can *create value* from the encountered information (for example, by saving the information with other related work). Overall, the ability to manage serendipitous information encounters in this way

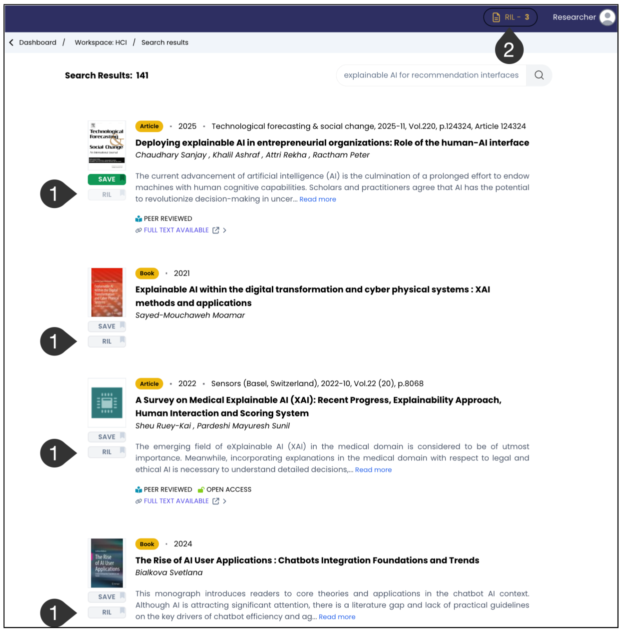

Figure 1: Revelio search results page. As well as providing a save button for each search result, a RIL button is provided to defer the search result for later review.

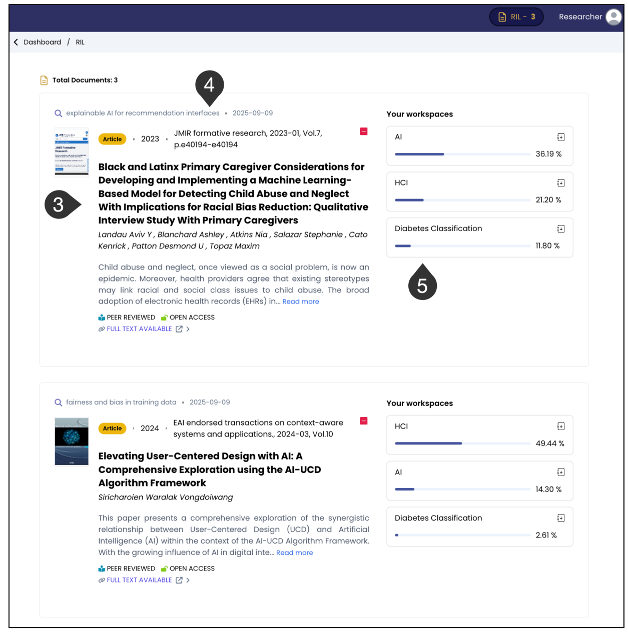

Figure 2: Revelio RIL list page. Each deferred search result is shown with visual similarity indicators to each workspace, providing a visual cue to help the searcher move the resource into the most relevant workspace.

can encourage materialization, without forcing the searcher to stop their current search activity at the time of the encounter.

Revelio was implemented as a web application using the MERN (MongoDB, Express.js, React, Node.js) technology stack. The client-side user interface is a single-page application built with ReactJS [17]. Server-side operations are managed by a custom REST API developed in Node.js, which acts as a proxy to our University's Ex Libris Primo academic digital library search API [44] for retrieving search results in JSON format. All user-specific data, including saved search results, search result metadata, and user-specific workspaces, are stored in a MongoDB NoSQL database [9].

4 User Study Methodology

To empirically evaluate how serendipitously encountered information is handled when searchers are provided with explicit support for deferring and subsequently reviewing such information (RQ2), and how this influences subjective and objective measures of performance (RQ3), we conducted a controlled laboratory experiment. The study was structured around the pursuit of two concurrent academic search tasks that are complex in nature (multi-faceted and open-ended).

We carefully considered the trade-offs between within-subjects and between-subjects designs, and chose the latter to avoid overloading participants with too many different tasks during the study. With this between-subjects design, half the participants used Revelio to conduct an assigned search activity; the other half conducted the same search activity using a multi-workspace baseline system that is visually and functionally the same as Revelio, but without the RIL buttons and access to the RIL list page. Participants were assigned to these two conditions in a round-robin fashion. To isolate the effect of our primary design intervention, other common digital library features such as faceted navigation and query auto-completion were omitted from both systems. This ensured that user behavior was focused on the core task of assessing and saving search results into the multi-workspace structure.

Although the original concept prototype included a feature to allow a searcher to create a new workspace for a search result on the RIL list page, this feature was hidden in Revelio to avoid potentially confounding factors during the user study. By restricting the support to match the serendipitously encountered information to existing workspaces, we limited the ability of searchers to manage natural information encounters that are not part of the study. This feature will be enabled for future naturalistic studies, where a searcher may serendipitously encounter information for an interest that does not yet have a workspace.

4.1 Concurrent Search Tasks

The two concurrent academic search tasks were framed within a typical academic scenario for computer science students (see Table 1). The task design followed established principles of incorporating uncertainty to encourage information seeking, using topics that were engaging but not overly familiar, and providing sufficient context for participants to understand the goals [22].

The task instructions were presented as simulated work tasks [6], with a focus on the exploratory browsing (discovery, learning, and investigation) aspect of the exploratory search model [46]. The activity was limited to gathering high-quality resources and organizing these into separate workspaces for later review. Participants

Table 1: The simulated work task scenario including the specific task topic areas and task descriptions presented to participants.

Simulated Work Task Scenario	
Imagine you are taking two Computer Science courses: Artificial Intelligence and Human-Computer Interaction. For your course projects, you need to gather academic articles on the two topics below. Your goal is to find a diverse set of high-quality resources for each project, saving them into separate workspaces to keep them organized.	
Task Topic Area	**Task Description**
1. Bias in Machine Learning Models	Explore the academic literature to understand the causes of bias in machine learning models and identify potential mitigation strategies.
2. User Interfaces for Recommender Systems	Identify research that discusses how Human-Computer Interaction (HCI) principles can be used to design more effective and transparent user interfaces for recommender systems.

had the freedom to start on whichever task they wished, and to switch between the two tasks as they saw fit, without a time limit.

4.2 Cultivating Serendipity in Search

When describing serendipity, Race and Makri noted that "It cannot be controlled but it can be influenced. It cannot be created on demand, but it can be cultivated" [39]. Motivated by this, we devised a method to cultivate opportunities for serendipitous information encounters in the context of pursuing concurrent search tasks. This was done by subtly manipulating the search results provided by the back-end search system, an approach that is commonly used to create desired experimental conditions within interactive information retrieval studies [2, 33, 36, 37, 42]. The general process is illustrated in Figure 3, and described below.

To start, we issued a variety of queries to collect a set of ten high-quality search results for each of the two search tasks. These were validated by the first two authors for relevance, and served as a pool of search results for creating potentially serendipitous encounters. Using the names of the workspaces as indicators of which search task the participant was pursuing, whenever a set of search results was generated for a given query, a randomly selected search result from the pool for the other task was artificially inserted into the search results list and the last item in the list was deleted (leaving 10 results in total). The insertion of a search result related to a 'background' task was inspired by a prior approach to artificially spark serendipity in a controlled study [12]; the conceptual overlap is aligned with prior studies of serendipity [13]. In order to avoid the potential of participants detecting a pattern, the insertion occurred randomly at the third, fifth, or seventh position in the list.

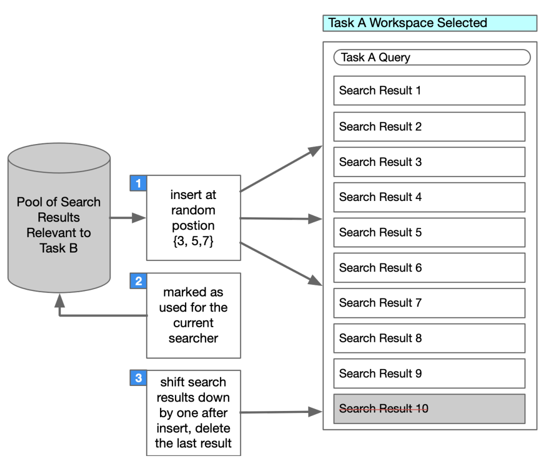

Figure 3: The process for creating a potential serendipitous information encounter by inserting search results relevant to Task B into Task A.

The search result selected from the pool was logged, ensuring that each result from the pool was used at most once for each participant. If the searcher moved to the next page of the results, this process was repeated for those results. Finally, a global caching mechanism was used to ensure that for any given query, the same set of search results would be presented (including the artificially inserted search result).

With this manipulation of the search results, every participant in the study had equal opportunity to encounter an unexpected yet potentially valuable search result for the other task. This consistent and repeatable mechanism makes it possible to compare multiple search interfaces in a controlled laboratory setting.

4.3 Data Collection

In order to measure search behavior while using the interfaces, all user interactions were captured using a client-side logging library [28]. This data was analyzed to detect specific patterns of use (e.g., changing workspaces, the use of the RIL buttons, saving search results).

To measure changes in perceived knowledge, participants were asked to rate their knowledge, certainty, and confidence regarding each task topic both before and after the search session. As these three items represent different facets of a single construct, their scores were aggregated to create a single Perceived Knowledge metric.

Following the completion of the search activity, participants' subjective experiences were captured using questionnaires. Usability measures of Ease of Use and Usefulness were captured with four statements each, adapted from the Technology Acceptance Model 2 (TAM2) [45]; Satisfaction was captured with four statements modeled after the TAM2 instrument [16]. User Engagement was assessed using the four User Engagement Scale Short Form (UES-SF) [34] sub-measures: Focused Attention, Perceived Usability, Aesthetic Appeal, and Reward. All responses were collected on a 5-point Likert scale ranging from 1 (strongly disagree) to 5 (strongly agree), with three items reverse coded as specified in the UES-SF instructions.

We did not attempt to measure participants' subjective experiences of serendipity. This was mainly because we did not want to reveal the study's specific focus on supporting aspects of serendipity until *after* capturing their experiences, so as not to bias the data collection. Also, while there are existing empirically-validated scales that measure individual propensity to experience serendipity in digital information spaces [19] and the extent to which digital information spaces support serendipity [29, 30], to our knowledge there are none that capture people's experiences of serendipity when interacting with digital spaces such as an academic digital library search interface.

All data was tested for normality and found to be within acceptable limits to justify the use of analysis of variance (ANOVA) for conducting statistical analysis, given its known robustness. The significance level is set at $p = 0.05$. Effect sizes are reported as η^2, where values less than 0.06 are considered small, up to 0.14 as medium, and greater than 0.14 as large.

4.4 Participants

We recruited 28 undergraduate students from the Computer Science programs at our university. To ensure participants had sufficient background to understand the academic search tasks, participation was limited to those in their second year of study or higher. The sample included 14 participants who self-identified as female, 13 as male, and one as non-binary, with a mean age of 22.64 years (SD=4.99). Participants were enrolled in several programs: BSc in Computer Science (n=17), BSc Honours in Computer Science (n=2), BSc in Software Systems Development (n=5), and BSc in Computer Science with a Business Concentration (n=4). The distribution by year of study was: 2nd year (n=4), 3rd year (n=9), and 4th year (n=15). A pre-study questionnaire confirmed that the median self-reported familiarity with academic digital library searching was "somewhat good". The courses specified in the simulated work task scenario were not being offered in the semester during which the data collection occurred; the specific tasks were not part of any previous offerings of these courses. This limited the ability for any participant to leverage current or previous coursework in conducting the search tasks.

4.5 Procedures

On arrival, participants provided informed consent and completed a pre-study questionnaire covering demographics and search experience. Each session began with training; participants watched an instructional video for their assigned interface and had the opportunity to familiarize themselves with its functionality for undertaking concurrent search tasks (saving search results into appropriate workspaces, switching between workspaces to pursue the other task, handling situations where a search result relevant to the other task is discovered). We did not explicitly indicate to participants that we were investigating serendipity or use terminology related to serendipitously encountering information during training; instead we described the study as one that focused on how people handle multiple concurrent search tasks with a search interface that supports multiple distinct workspaces.

Participants then proceeded to the main search activity. After presenting the simulated work task, a pre-task questionnaire was administered to gauge prior knowledge. The participants were then given access to the assigned search interface and asked to pursue the two concurrent search tasks. When they reported having completed the search activity, post-task questionnaires were administered. The total session duration ranged from 55-70 minutes.

A low-impact deception protocol was employed, which was reviewed and approved by our university's Research Ethics Board. Firstly, we did not reveal the true purpose of the study, as described above. Secondly, we informed participants that they would each receive $30 remuneration and that a $20 bonus would be awarded to whomever collected the most relevant resources for the two search tasks. This was done to incentivize participants to take the task seriously and do their best work. However, we gave all participants the full $50 compensation. Additionally, without participants' prior knowledge, the search results lists were manipulated to include search results relevant to the other task. At the end of the session, participants were debriefed. Participants were asked to consent to use their data with full knowledge of the study, and were offered to withdraw without penalty. All participants provided this secondary consent.

5 Results

Since the opportunity for potentially serendipitous information encounters was created artificially, we first verify that participants noticed some of the inserted search results relevant to the other task, and chose to save them. These are presented separately for each task and each experimental condition (search interface) in Table 2. A series of one-way analyses of variance (ANOVAs) were conducted to determine if there were differences between the experimental conditions with respect to whether these information encounters were pursued.

The interaction with each artificially inserted search results was classified as either *correctly saved* (i.e., the participant was able to save the search result in the correct workspace, facilitating follow-up and potential value creation from the encountered information), *incorrectly saved* (i.e., the participant saved the inserted search result, but to the current workspace instead of the correct one), or *missed* (i.e., either the participant did not consider the inserted search result, did not experience a serendipity trigger when considering the search result, or did not wish to exert the effort to put it into the correct workspace). Of note, the total number of these interactions varies between the tasks and experimental conditions, based on the number of queries issued and pages of search results viewed.

For both search tasks and and search interface conditions, participants identified some of the inserted search results that were relevant to the other task and chose to save them in their correct workspaces. While our training and incentives may have primed the participants to save more search results than they might have done in real life, these results show that the approach for cultivating serendipitous information encounters was effective.

5.1 Handling Encountered Information (RQ2)

Referring again to Table 2, there are clear differences in the success rate of identifying and saving the inserted search results relevant to the other task into the correct workspace, at statistically significant

Table 2: Mean (standard deviation) frequency of the result of encountering the inserted search results relevant to the other task.

Task	Encounter	Baseline	Revelio	ANOVA
Task 1 (Bias in ML)	Correctly Saved	1.43 (1.02)	2.29 (0.73)	**$F(1, 26) = 6.59$, $p < 0.05$, $\eta_p^2 = 0.20$**
	Incorrectly Saved	1.00 (0.88)	0.57 (0.65)	$F(1, 26) = 2.17$, $p = 0.15$
	Missed	1.57 (0.94)	0.64 (1.01)	**$F(1, 26) = 6.37$, $p < 0.05$, $\eta_p^2 = 0.20$**
Task 2 (UI for RecSys)	Correctly Saved	1.14 (1.10)	2.21 (0.70)	**$F(1, 26) = 9.47$, $p < 0.01$, $\eta_p^2 = 0.27$**
	Incorrectly Saved	1.07 (1.00)	0.36 (0.63)	**$F(1, 26) = 5.12$, $p < 0.05$, $\eta_p^2 = 0.16$**
	Missed	1.86 (1.03)	0.86 (0.86)	**$F(1, 26) = 7.77$, $p < 0.01$, $\eta_p^2 = 0.23$**

levels and with large effect sizes. For both tasks, the participants who used Revelio had more correctly saved serendipitous information encounters, and fewer missed opportunities. For Task 2, the participants who used the Baseline had more incorrectly saved encounters (saved into the wrong workspace), which is evidence of not wishing to incur the cost of changing workspaces in order to save the search result in the correct place. Given the similarity in the nature of how the serendipitous information encounters were handled, and the ability for the participants to freely switch between the workspaces and associated search tasks, from this point forward we report findings aggregated over the two tasks, but will return to considering the tasks separately when analyzing perceived knowledge gain.

An important element of studying concurrent search tasks is considering the cost of switching between tasks. To understand how each interface influenced this, we conducted a detailed analysis of the circumstances of each task switch using the interaction logs (see Table 3). The number of intentional workspace switches between the participants using the two interfaces was similar, which is expected given the common multi-workspace structure. However, when an serendipitous information encounter was triggered, there were statistically significant differences with large effect sizes in how participants dealt with these depending on which interface they used. For those in the Baseline condition, these resulted in a workspace switch and re-issuing the same query to find the search result that triggered the encounter. When this *serendipity-driven workspace switch* occurred, participants would sometimes remain in this workspace and continue searching for this task rather than returning to their previous workspace/task. Of note, the number of serendipity-driven workspace switches in the Baseline condition is greater than the number of correctly saved serendipitous information encounters ($3.57 > 1.43 + 1.14 = 2.57$). This difference represents the number of times the participants changed to the other workspace, re-issued the same query, but then had difficulty re-finding the search result that triggered the encounter. We consider these failures of the Baseline to adequately support the participants' handling of serendipitously encountered information.

When participants who used Revelio encountered a potentially serendipitous search result, they had two ways to handle this discovery: temporarily set it aside using the RIL button, or follow the same process as the participants who used the Baseline (switching workspaces and re-issuing the same query to re-find the search result). Although some participants defaulted to this workspace switching approach, it happened significantly fewer times, and the participants seldom remained in that workspace to continue searching. More commonly, Revelio participants used the RIL button to save the serendipitously encountered information for later review and continued with their primary search task. Here also, the number of serendipity-driven workspace switches and search results saved to the RIL list were greater than the number of correctly saved serendipitously encountered search results ($1.86 + 6.14 = 8.00 > 2.29 + 2.21 = 4.50$). While it is possible that some of the serendipity-driven workspace switches resulted in a failed re-finding similar to what was observed in the Baseline condition, there is a second factor at play; in some cases, participants used the RIL list not only to temporarily hold serendipitous discoveries, but also to hold other search results that they were not sure about and wished to evaluate for relevancy later. This highlights the value of the RIL button/list page beyond managing serendipitous information encounters.

5.2 Impact on Performance (RQ3)

5.2.1 Overall Usability and User Engagement. To evaluate participants' perceived value of each interface in relation to pursuing concurrent search tasks when serendipitous information encounters occur, we administered a series of questionnaires following the completion of the search activities that were focused on three aspects of overall usability and four aspects of user engagement. Statistical analyses were conducted on each of these concepts separately (see Table 4).

Statistically significant differences with large effect sizes were found for all three measures of overall usability. Although the feature-level differences between the two interfaces were small, the ability to set aside a potentially serendipitous search result for later review and matching to another workspace had a positive impact on Revelio's overall usability.

Considering user engagement, there were statistically significant differences with large effect sizes on three of four measures. Participants who used Revelio reported higher levels of focused attention, perceived usability, and reward factor. We attribute this to their ability to easily defer serendipitous information encounters when they occurred, stay focused on the current search activity, and later use the visual similarity indicators in the RIL list page to decide on an appropriate workspace for these search results. There was little difference with in aesthetic appeal measure, which is expected given that the interfaces were visually similar.

5.2.2 Search Task Performance. To assess search task performance, we analyzed three behavioral metrics derived from system logs: the search session duration, the number of saved search results in both workspaces, and the relevance (precision) of the saved search

Table 3: Mean (standard deviation) frequency of interactions associated with concurrent searching and information encounters.

Primary Interactions	Baseline	Revelio	ANOVA
Intentional workspace switch	3.86 (1.82)	3.57 (0.85)	$F(1, 26) = 0.64, p = 0.43$
Serendipity-driven workspace switch	3.57 (1.38)	1.86 (0.77)	$\mathbf{F(1, 26) = 21.27}$, $p < 10^{-4}$, η_p^2 = **0.45**
Remain in workspace after serendipity-driven switch	1.07 (0.62)	0.29 (0.47)	$\mathbf{F(1, 26) = 14.43}$, $p < 0.001$, η_p^2 = **0.36**
Search result saved using RIL button	-	6.14 (1.09)	
Access RIL list	-	5.43 (1.21)	

Table 4: Mean (standard deviation) of the 5-point Likert scores for overall usability and user engagement.

Measure	Baseline	Revelio	ANOVA Results
Overall Usability			
Ease of Use	3.73 (0.53)	4.25 (0.50)	**F(1, 26) = 7.04, p < 0.05, η_p^2 = 0.21**
Usefulness	3.32 (0.39)	4.12 (0.68)	**F(1, 26) = 14.88, p < 0.001, η_p^2 = 0.36**
Satisfaction	3.52 (0.56)	4.12 (0.66)	**F(1, 26) = 6.86, p < 0.05, η_p^2 = 0.21**
User Engagement			
Focused Attention	3.21 (0.89)	4.10 (0.66)	**F(1, 26) = 8.82, p < 0.01, η_p^2 = 0.25**
Perceived Usability	2.81 (0.96)	3.88 (0.61)	**F(1, 26) = 12.48, p < 0.01, η_p^2 = 0.32**
Aesthetic Appeal	3.52 (0.75)	3.90 (0.56)	F(1, 26) = 2.32, p = 0.14
Reward Factor	3.50 (0.96)	4.43 (0.40)	**F(1, 26) = 11.18, p < 0.01, η_p^2 = 0.30**

results with respect to the task associated with the workspace in which they were saved. The results are presented in Table 5.

Participants using Revelio were more efficient in completing the concurrent search tasks than those using the Baseline interface, at a statistically significant level with a large effect size. As the only difference between these interfaces was the RIL buttons and RIL list page, this suggests that deferring and later reviewing the serendipitous information encounter is more efficient than workspace switching. That is, instead of interrupting their primary task to switch workspaces, participants could use the low-effort RIL button to set aside the search result, keeping their momentum with respect to their current search task and leading to faster overall task completion.

In terms of the number of search results saved in the workspaces, there was little difference between the two conditions in this study. This is not unexpected, given that the search results lists were identical, as was the manner in which the search results were presented on the SERP.

To calculate the precision of the saved search results, it was necessary to assess the relevance of the search results in each workspace with respect to the task associated with that workspace. The first two authors served as independent judges, rating each search result saved on a three-point scale (0=not relevant, 1=partially relevant, 2=fully relevant). Inter-rater reliability was very high, confirmed by a Cohen's kappa of κ=0.90. In cases where there were disagreements, these were resolved through discussion. The ratings were converted to a binary scale for analysis, where both partially and fully relevant search results were considered relevant. Such an approach is common practice in the field [41], and is appropriate for the exploratory browsing aspect of the exploratory search model [16], where even partially relevant search results support learning and discovery.

The search results saved by the Revelio participants were statistically significantly more relevant to the search tasks than those saved by the Baseline participants, with a large effect size. This increased precision demonstrates two key benefits of the RIL features in Revelio. Firstly, by allowing users to easily save encountered information without having to make a decision on what to do with it until later, those who used Revelio were able to stay focused on their current search task and make better decisions about relevant search results. Secondly, when they returned to the RIL list page, they were able to give these search results their focused attention, and were supported by the similarity visualization with respect to each workspace. This 'save now, organize later' workflow resulted in a higher quality collection of saved search results for each task, ultimately increasing search performance.

Of note, there was also downward pressure on the precision of the saved search results among the participants in the Baseline condition, due to incorrectly saving inserted search results relevant to the other task into the current task's workspace. The decision to do so, presumably with the intention to put it in the correct workspace later, did not make the search activity any more efficient but did make the saved results less precise.

5.2.3 Perceived Knowledge Gain. To assess how each interface affected participants' perceived knowledge, we collected and analyzed pre- and post-task measurements, calculated the average knowledge gain per task, and conducted statistical analysis on this data (see Table 6). Overall, there was knowledge gain in both conditions and both tasks, as one would expect after performing a search activity that included external motivation for finding relevant resources. However, the increase in knowledge was only at a statistically significant level for the participants who used Revelio.

As the only difference was whether there was explicit support for managing potentially serendipitous information encounters, we conclude that the disruption of having to deal with encountered information as it is found (in the Baseline condition) had a negative effect on perceived knowledge gain. In contrast, those who used Revelio could remain focused on their primary search task, with

Table 5: Mean (standard deviation) values of the search performance metrics.

Measure	Baseline	Revelio	ANOVA Results
Time to Task Completion (seconds)	1291 (93)	1157 (172)	**F(1, 26) = 6.62, $p < 0.05$, $\eta_p^2 = 0.20$**
Number of Saved Search Results	9.07 (1.38)	10.07 (1.54)	F(1, 26) = 2.76, $p = 0.11$
Precision of Saved Search Results	64.4% (0.15)	80.5% (0.10)	**F(1, 26) = 10.51, $p < 0.01$, $\eta_p^2 = 0.26$**

Table 6: Mean (standard deviation) of perceived knowledge scores.

Task	Interface	Pre-Task Knowledge	Post-Task Knowledge	Δ	ANOVA Result
Task 1	Baseline	3.10 (0.76)	3.48 (0.73)	0.38	F(1,13) = 2.77, $p = 0.120$
(Bias in ML)	Revelio	3.21 (0.66)	4.07 (0.48)	0.86	**F(1,13) = 13.00, $p < 0.01$, , $\eta_p^2 = 0.26$**
Task 2	Baseline	3.24 (0.87)	3.67 (0.51)	0.43	F(1,13) = 3.39, $p = 0.089$
(UI for RecSys)	Revelio	3.19 (0.58)	4.24 (0.28)	1.05	**F(1,13) = 22.11, $p < 0.001$, , $\eta_p^2 = 0.43$**

a positive effect on their perceived knowledge at the end of the search activity.

6 Discussion

To date, most serendipity studies have examined the process from a human behavior perspective [7, 13, 25–27, 31]. Of the few studies that have investigated how serendipity can be supported in a search context, all have focused on supporting people in noticing serendipity triggers and making mental associations [3, 38, 43]. Our research is novel, as it focuses on supporting the later stages of the serendipity process, following up on potentially serendipitous information encounters in order to create value from them, in a way that minimizes the impact on the primary search task. It achieves this by using lightweight additions to a standard academic digital library search interface [15].

To the best of our knowledge, the existing literature has does not include reliable and repeatable methods for creating the opportunity for serendipitous information encounters within a controlled user study context. Drawing inspiration from prior attempts [12], we developed a novel approach that can easily be replicated in future studies. We hope that this will spark further research that incorporates controlled studies of how best to support aspects of serendipity in interactive information retrieval.

Although our study highlighted the benefits of being able to defer and subsequently follow-up on potentially serendipitous information, such deferral is not always beneficial. In some cases, immediate follow-up is desirable, such as when a searcher has become fatigued by the primary task and would welcome a change, when an association is triggered that sparks a moment of particular insight, clarity, or creativity that the searcher does not wish to lose, or when the task associated with the encountered information is perceived to be more important than the primary task. In such cases, the lightweight features of Revelio do not get in the way of such task switching; they simply provide an alternative pathway for dealing with serendipitous information encounters when the searcher would prefer to stay focused on their primary task.

Several limitations must be acknowledged. The primary limitation of Revelio's interface is a lack of scalability of the approach for matching the search results in the RIL list page to the searcher's existing workspace. If the searcher has a large number of workspaces, listing the top five and providing interactive access to the rest will keep this page compact. With respect to the study design, the simulated task environment created more opportunities for serendipity than might naturally occur; this was necessary to ensure that there were sufficient opportunities for serendipitous information encounters. Further study of Revelio in the wild is needed, including naturalistic observation, diary studies, and longitudinal logging. Such studies will allow for a more direct connection to be made between the features of the interface and the experiences of searchers as they encounter information.

7 Conclusion and Future Work

There are three novel contributions of this research. Firstly, it introduces a refined version of Revelio, an academic digital library search interface that supports the deferral and later review of potentially serendipitous information. This is one of just a few search user interfaces that have been designed to support serendipity and is a practical contribution. Secondly, it proposes and demonstrates the effectiveness of an approach for creating opportunities for serendipitous information encounters within a controlled laboratory setting. As prior approaches have been unsuccessful [12], this constitutes a methodological contribution. Thirdly, it provides empirical evidence of the value of Revelio over a baseline academic digital library search interface, showing that deferral and subsequent review of serendipitous information encounters increases the overall performance and subjective opinions of searchers.

Beyond conducting further studies of Revelio in naturalistic settings, there is broad scope for further research on how search interfaces can support and enhance serendipitous information encounters. For example, a slide-in sidebar could allow a potentially serendipitous information encounter to be immediately pursued to the point of determining it is likely to be valuable, allowing the searcher to save it in the appropriate workspace, then close the sidebar to return to their previous search. AI may also be leveraged in ways that further cultivate opportunities for serendipity, including triggering serendipity by suggesting potential connections between search results displayed in the SERP and existing workspaces, or across saved search results saved in different workspaces.

Further study of the 'save now, organize later' workflow in contexts beyond academic search are warranted, including within new search interaction paradigms such as conversational search and GenAI-based chat interfaces. More research into the circumstances when it is better to defer the serendipitously encountered information versus when it is better to immediately examine it are also needed. Integrated support for both circumstances could provide searchers with the agency to make the decision on how to pursue the information encounter with respect to their current situation and context. By supporting people in managing information they serendipitously encounter, we can provide an infrastructure to support insight creation, creativity, and fluid thinking.

Acknowledgments

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) through the Discovery Grant (RGPIN-2023-04070) held by the first author.

This research was conducted at the University of Regina, which is situated on Treaty 4 and Treaty 6 territories - the homelands of the nêhiyawak, Anihšināpēk, Dakota, Lakota, Nakoda peoples, and Michif/Métis nation.

References

[1] Paul André, m.c. schraefel, Jaime Teevan, and Susan T. Dumais. 2009. Discovery is never by chance: Designing for (un)serendipity. In *Proceedings of the ACM Conference on Creativity and Cognition.* Association for Computing Machinery, New York, NY, USA, 305–314. doi:10.1145/1640233.1640279

[2] Ioannis Arapakis, Joemon M. Jose, and Philip D. Gray. 2008. Affective feedback: An investigation into the role of emotions in the information seeking process. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval.* Association for Computing Machinery, New York, NY, USA, 395–402. doi:10.1145/1390334.1390403

[3] Russell Beale. 2007. Supporting serendipity: Using ambient intelligence to augment user exploration for data mining and web browsing. *International Journal of Human-Computer Studies* 65, 5 (May 2007), 421–433. doi:10.1016/j.ijhcs.2006.11.012

[4] Ofer Bergman, Steve Whittaker, and Joel Schooler. 2021. Out of sight and out of mind: Bookmarks are created but not used. *Journal of Librarianship and Information Science* 53, 2 (2021), 338–348. doi:10.1177/0961000620949652

[5] Christine L. Borgman. 1999. What are digital libraries? Competing visions. *Information Processing & Management* 35, 3 (1999), 227–243. doi:10.1016/S0306-4573(98)00059-4

[6] Pia Borlund. 2003. The IIR evaluation model: A framework for evaluation of interactive information retrieval systems. *Information Research* 8, 3 (2003), 3–8. http://informationr.net/ir/8-3/paper152.html

[7] Christian Busch. 2024. Towards a theory of serendipity: A systematic review and conceptualization. *Journal of Management Studies* 61, 3 (2024), 1110–1151. doi:10.1111/joms.12890

[8] Joseph Chee Chang, Nathan Hahn, Yongsung Kim, Julina Coupland, Bradley Breneisen, Hannah S. Kim, John Hwong, and Aniket Kittur. 2021. When the tab comes due: Challenges in the cost structure of browser tab usage. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21).* Association for Computing Machinery, New York, NY, USA, Article 148, 15 pages. doi:10.1145/3411764.3445585

[9] Kristina Chodorow. 2013. *MongoDB: The Definitive Guide.* O'Reilly Media, Inc., Sebastopol, CA.

[10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies.* Association for Computational Linguistics, Minneapolis, MN, USA, 4171–4186. doi:10.18653/v1/N19-1423

[11] Patrick Dubroy and Ravin Balakrishnan. 2010. A study of tabbed browsing among Mozilla Firefox users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10).* Association for Computing Machinery, New York, NY, USA, 673–682. doi:10.1145/1753326.1753426

[12] Sanda Erdelez. 2004. Investigation of information encountering in the controlled research environment. *Information Processing & Management* 40, 6 (2004), 1013–1025. doi:10.1016/j.ipm.2004.02.002

[13] Sanda Erdelez and Stephann Makri. 2020. Information encountering re-encountered: A conceptual re-examination of serendipity in the context of information acquisition. *Journal of Documentation* 76, 3 (2020), 731–751. doi:10.1108/JD-08-2019-0151

[14] Sebastian Gomes, Miriam Boon, and Orland Hoeber. 2022. A study of cross-session cross-device search within an academic digital library. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22).* Association for Computing Machinery, New York, NY, USA, 384–394. doi:10.1145/3477495.3531929

[15] Orland Hoeber. 2025. Design principles for exploratory search interfaces. In *Proceedings of the 2025 ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '25).* Association for Computing Machinery, New York, NY, USA, 12–22. doi:10.1145/3698204.3716443

[16] Orland Hoeber and Soumya Shukla. 2022. A study of visually linked keywords to support exploratory browsing in academic search. *Journal of the Association for Information Science and Technology* 73, 8 (2022), 1171–1191. doi:10.1002/asi.24623

[17] Shama Hoque. 2020. *Full-stack React Projects: Learn MERN Stack Development by Building Modern Web Apps Using MongoDB, Express, React, and Node.js, Second edition.* Packt Publishing, Birmingham, UK.

[18] Jeff Huang, Thomas Lin, and Ryen W. White. 2012. No search result left behind: Branching behavior with browser tabs. In *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining (WSDM '12).* Association for Computing Machinery, New York, NY, USA, 203–212. doi:10.1145/2124295.2124322

[19] Eunjin (Anna) Kim, Kevin Wise, Sanda Erdelez, and Yi-Hsuan Chiang. 2025. Development of a scale for measuring individual propensity for serendipitous information encounters in an online environment. *Journal of Information Science* 51, 2 (2025), 431–444. doi:10.1177/01655515221141041

[20] Arthur Koestler. 1964. *The Act of Creation.* Hutchinson, London, UK.

[21] Ralf Krestel, Gianluca Demartini, and Eelco Herder. 2011. Visual interfaces for stimulating exploratory search. In *Proceedings of the ACM/IEEE-CS Joint Conference on Digital Libraries.* Association for Computing Machinery, New York, NY, USA, 393–394. doi:10.1145/1998076.1998151

[22] Bill Kules and Robert Capra. 2009. Designing exploratory search tasks for user studies of information seeking support systems. In *Proceedings of the ACM/IEEE-CS Joint Conference on Digital Libraries.* Association for Computing Machinery, New York, NY, USA, 419–420. doi:10.1145/1555400.1555492

[23] Yuan Li, Anita Crescenzi, Austin R. Ward, and Rob Capra. 2023. Thinking inside the box: An evaluation of a novel search-assisting tool for supporting (meta)cognition during exploratory search. *Journal of the Association for Information Science and Technology* 74, 9 (2023), 1049–1066. doi:10.1002/asi.24801

[24] Yuan Li, Austin R. Ward, and Rob Capra. 2021. An analysis of information types and cognitive activities involved in cross-session search. In *Proceedings of ACM SIGIR Conference on Human Information Interaction and Retrieval.* Association for Computing Machinery, New York, NY, USA, 313–317. doi:10.1145/3406522.3446044

[25] Stephann Makri and Ann Blandford. 2012. Coming across information serendipitously – Part 1: A process model. *Journal of Documentation* 68, 5 (08 2012), 684–705. doi:10.1108/00220411211256030

[26] Stephann Makri, Ann Blandford, Mel Woods, Sarah Sharples, and Deborah Maxwell. 2014. "Making my own luck": Serendipity strategies and how to support them in digital information environments. *Journal of the Association for Information Science and Technology* 65, 11 (2014), 2179–2194. doi:10.1002/asi.23200

[27] Stephann Makri and Lily Buckley. 2020. Down the rabbit hole: Investigating disruption of the information encountering process. *Journal of the Association for Information Science and Technology* 71, 2 (2020), 127–142. doi:10.1002/asi.24233

[28] David Maxwell and Claudia Hauff. 2021. LogUI: Contemporary logging infrastructure for web-based experiments. In *Proceedings of the European Conference on Information Retrieval.* Springer, Cham, 525–530. doi:10.1007/978-3-030-72240-1_59

[29] Lori McCay-Peet and Elaine Toms. 2011. Measuring the dimensions of serendipity in digital environments. *Information Research* 16, 3 (2011). https://informationr.net/ir/16-3/paper483.html

[30] Lori McCay-Peet, Elaine Toms, and E. Kevin Kelloway. 2014. Development and assessment of the content validity of a scale to measure how well a digital environment facilitates serendipity. *Information Research* 19, 3 (2014). https://informationr.net/ir/19-3/paper630.html

[31] Lori McCay-Peet and Elaine G. Toms. 2015. Investigating serendipity: How it unfolds and what may influence it. *Journal of the Association for Information Science and Technology* 66, 7 (2015), 1463–1476. doi:10.1002/asi.23273

[32] Ei Ei Mon and Orland Hoeber. 2024. Enabling serendipity during digital library search. *Proceedings of the Association for Information Science and Technology* 61, 1 (2024), 1030–1032. doi:10.1002/pra2.1176

[33] Robert Moro and Maria Bielikova. 2020. Navigation leads for exploratory search and navigation in digital libraries. *Knowledge and Information Systems* 62, 7 (2020), 2739–2764. doi:10.1007/s10115-019-01434-2

[34] Heather L. O'Brien, Paul Cairns, and Mark Hall. 2018. A practical approach to measuring user engagement with the refined user engagement scale (UES) and new UES short form. *International Journal of Human-Computer Studies* 112 (2018), 28–39. doi:10.1016/j.ijhcs.2018.01.004

[35] Kimeya Orin and Orland Hoeber. 2025. Organizing found information in public digital library search. In *Proceedings of ACM SIGIR Conference on Human Information Interaction and Retrieval.* Association for Computing Machinery, New York, NY, USA, 358–366. doi:10.1145/3698204.3716471

[36] Abbas Pirmoradi and Orland Hoeber. 2025. Bridging in-task emotional responses with post-task evaluations in digital library search interface user studies. *Information Processing & Management* 62, 3 (2025), 25 pages. doi:10.1016/j.ipm.2025.104069
[37] Kim Plassmeier, Timo Borst, Christiane Behnert, and Dirk Lewandowski. 2015. Evaluating popularity data for relevance ranking in library information systems. *Proceedings of the Association for Information Science and Technology* 52, 1 (2015), 1–4. doi:10.1002/pra2.2015.1450520100125
[38] Chunxiu Qin, Yaxi Liu, Xubu Ma, Jiangping Chen, and Huigang Liang. 2022. Designing for serendipity in online knowledge communities: An investigation of tag presentation formats and openness to experience. *Journal of the Association for Information Science and Technology* 73, 10 (Oct. 2022), 1401–1417. doi:10.1002/asi.24640
[39] Tammera M. Race and Stephann Makri. 2016. *Accidental Information Discovery: Cultivating Serendipity in the Digital Age.* Chandos Publishing, Cambridge, MA, USA.
[40] Ian Ruthven. 2008. Interactive information retrieval. *Annual Review of Information Science and Technology* 42, 1 (2008), 43–91. doi:10.1002/aris.2008.1440420109
[41] Falk Scholer and Andrew Turpin. 2009. Metric and relevance mismatch in retrieval evaluation. In *Proceedings of the Asia Information Retrieval Symposium.* Springer, Berlin, Germany, 50–62. doi:10.1007/978-3-642-04769-5_5
[42] Yang Song, Xiaolin Shi, and Xin Fu. 2013. Evaluating and predicting user engagement change with degraded search relevance. In *Proceedings of the International Conference on World Wide Web.* Association for Computing Machinery, New York, NY, USA, 1213–1224. doi:10.1145/2488388.2488494
[43] Maria Taramigkou, Dimitris Apostolou, and Gregoris Mentzas. 2017. Supporting creativity through the interactive exploratory search paradigm. *International Journal of Human–Computer Interaction* 33, 2 (Feb. 2017), 94–114. doi:10.1080/10447318.2016.1220104
[44] Jason Vaughan. 2011. Ex Libris Primo Central. *Library Technology Reports* 47, 1 (2011), 39–47.
[45] Viswanath Venkatesh and Fred Davis. 2000. A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management Science* 46, 2 (2000), 186–204. doi:10.1287/mnsc.46.2.186.11926
[46] Ryen White and Resa Roth. 2009. *Exploratory Search: Beyond the Query-Response Paradigm.* Morgan & Claypool Publishers, San Rafael, CA. doi:10.2200/S00174ED1V01Y200901ICR003