

Understanding the “Pathway” Towards a Searcher’s Learning Objective

KELSEY URGO and JAIME ARGUELLO, School of Information and Library Science,
University of North Carolina at Chapel Hill

Search systems are often used to support learning-oriented goals. This trend has given rise to the “search-as-learning” movement, which proposes that search systems should be designed to support learning. To this end, an important research question is: How does a searcher’s *type* of learning objective (LO) influence their trajectory (or *pathway*) toward that objective? We report on a lab study ($N = 36$) in which participants gathered information to meet a specific type of LO. To characterize LOs *and pathways*, we leveraged Anderson and Krathwohl’s (A&K’s) taxonomy [3]. A&K’s taxonomy situates LOs at the intersection of two orthogonal dimensions: (1) *cognitive process* (CP) (remember, understand, apply, analyze, evaluate, and create) and (2) *knowledge type* (factual, conceptual, procedural, and metacognitive knowledge). Participants completed learning-oriented search tasks that varied along three CPs (apply, evaluate, and create) and three knowledge types (factual, conceptual, and procedural knowledge). A *pathway* is defined as a sequence of *learning instances* (e.g., subgoals) that were also each classified into cells from A&K’s taxonomy. Our study used a think-aloud protocol, and pathways were generated through a qualitative analysis of participants’ think-aloud comments and recorded screen activities. We investigate three research questions. First, in RQ1, we study the impact of the LO on pathway characteristics (e.g., pathway length). Second, in RQ2, we study the impact of the LO on the types of A&K cells traversed along the pathway. Third, in RQ3, we study common and uncommon *transitions* between A&K cells along pathways conditioned on the knowledge type of the objective. We discuss implications of our results for designing search systems to support learning.

CCS Concepts: • **Information systems** → **Users and interactive retrieval**;

Additional Key Words and Phrases: Search-as-learning, learning objectives, learning pathways

ACM Reference format:

Kelsey Urgo and Jaime Arguello. 2022. Understanding the “Pathway” Towards a Searcher’s Learning Objective. *ACM Trans. Inf. Syst.* 40, 4, Article 77 (January 2022), 42 pages.
<https://doi.org/10.1145/3495222>

1 INTRODUCTION

People often search for information in order to learn something new. While current search systems are effective in helping users complete simple look-up tasks (e.g., navigational or fact-finding

This work was supported by NSF grant IIS-1718295. Any opinions, findings, conclusions, and recommendations expressed in this article are those of the authors and do not necessarily reflect the views of the sponsor.

Authors’ address: K. Urgo and J. Arguello, School of Information and Library Science, University of North Carolina at Chapel Hill, Manning Hall, 216 Lenoir Drive, Chapel Hill, North Carolina 27599; emails: {kurgo, jarguello}@unc.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

1046-8188/2022/01-ART77 \$15.00

<https://doi.org/10.1145/3495222>

tasks), they provide less support for users working on complex tasks that involve learning. In recent years, the “search-as-learning” research community has argued that search systems should be better designed to support learning. Recent summits have taken place to develop research agendas in the area of search-as-learning [2, 9]. Participants at these summits proposed that future research should focus on: (1) understanding the contexts in which people search for information in order to learn, (2) understanding the cognitive biases promoted by existing search systems, (3) understanding search as a learning process, and (4) developing search interfaces and tools that encourage and support learning [2, 9]. Our research in this article focuses on understanding search as a learning process. Additionally, we discuss how our results have implications for designing novel search tools to encourage and support learning. Prior studies in the area of search-as-learning have investigated a wide range of research questions. Many studies have investigated how different factors can influence learning during search. Specifically, studies have investigated characteristics of the individual searcher [34, 36, 38, 49], characteristics of the search task [17, 22, 29, 30], and characteristics of the search system [6, 8, 14, 15, 18–20, 23, 35, 37, 39, 41, 42, 46, 50, 52]. Additionally, studies have investigated the relation between specific search behaviors and learning outcomes [1, 4, 8, 10, 16, 26, 28, 32, 35, 52, 53]. When people search to learn, they typically have a specific *learning objective* (LO) in mind—“I need to find information that enables me to do <learning objective>.” Our research in this article investigates an important question that has not been directly addressed in prior work:

How does the type of *learning objective* that a searcher is aiming to accomplish influence their trajectory or *pathway* toward that objective? In other words, how do searchers decompose a specific objective into a sequence of subgoals or learning instances?

To address this question, we conducted a lab study in which participants ($N = 36$) completed search tasks with different types of LOs. To gain insights about participants’ pathways (i.e., sequences of subgoals) toward an objective, the study used a think-aloud protocol. To manipulate LOs *and* to characterize pathways toward an objective, we leveraged **Anderson and Krathwohl’s (A&K’s)** taxonomy of learning (referred to as A&K’s taxonomy) [3].

1.1 A&K’s Taxonomy

In the field of education, A&K’s taxonomy was developed to help educators more precisely define LOs for students [3]. Additionally, it was developed to help educators *align* instructional exercises and assessments with the target LO. As illustrated in Table 1, A&K’s taxonomy situates LOs at the intersection of two orthogonal dimensions: (1) **cognitive process (CP)** and (2) knowledge type. Anderson and Krathwohl [3] argued that LOs can be viewed as a combination of a “verb” and a “noun” (e.g., recall factual knowledge). In this respect, the CP defines the “verb” and the knowledge type defines the “noun” of the LO.

The CP dimension defines the types of cognitive activities associated with the LO. In other words, it defines the types of mental activities learners should be able to perform once the objective is met. CPs range from simple to complex: remember, understand, apply, analyze, evaluate, and create. If a *remember* objective is met, it means that the learner will be able to recall or regurgitate information verbatim. If an *understand* objective is met, it means that the learner will be able to explain information in their own words or illustrate examples of a construct. If an *apply* objective is met, it means that the learner will be able to execute a process or use the acquired knowledge in a new scenario. If an *analyze* objective is met, it means that the learner will be able to explain relations between elements (e.g., similarities and differences). If an *evaluate* objective is met, it

Table 1. A&K’s Two-dimensional Taxonomy and Example Pathway Toward an Evaluate/procedural LO (Gray Cell)

Knowledge Type	Cognitive Process					
	Remember	Understand	Apply	Analyze	Evaluate	Create
Factual						
Conceptual		LI ₂				
Procedural		LI ₁ , LI ₃ , LI ₄	LI ₅	LI ₆ , LI ₇	LI ₈	
Metacognitive						

LI_{*i*} denotes the *i*th learning instance along the pathway.

means that the learner will be able to critique or prioritize elements. Finally, if a *create* objective is met, it means that the learner will be able to generate a new solution to a problem or organize information using a novel representation.

The *knowledge type* dimension defines the type of knowledge associated with the objective. A&K’s taxonomy defines four types of knowledge: factual, conceptual, procedural, and metacognitive knowledge. The first three knowledge types relate to external knowledge about the world: *factual knowledge* relates to self-contained, objective bits of information; *conceptual knowledge* relates to concepts, categories, theories, principles, schemas, and models; and *procedural knowledge* relates to knowledge about how to perform a task. Conversely, *metacognitive knowledge* looks inward, and relates to knowledge about one’s own cognition or cognition in general.

1.2 Characterizing Pathways

Our goal is to understand the pathways followed by searchers toward a specific LO. A *pathway* is defined as a sequence of *learning instances* toward an objective. A *learning instance* is defined as a point during the search and learning process in which the searcher either: (1) sets forth a new learning-oriented subgoal or (2) serendipitously learns something new and useful toward achieving the objective. The A&K taxonomy can be leveraged to classify LOs and learning instances along the pathway to the objective.

To better understand the concept of a pathway, consider the example in Table 1. Imagine a searcher who wants to determine which sorting algorithm is more efficient: quicksort or mergesort. This LO involves making a judgement (i.e., evaluate) between two algorithms (i.e., procedural knowledge). Therefore, based on A&K’s taxonomy, the objective can be classified as evaluate/procedural (gray cell in Table 1). Given this objective, a searcher may follow the pathway below, which includes **learning instances (LIs)** that are either planned or unplanned.

- **LI₁ (understand/procedural)**: First, the searcher may look for and review an example of mergesort to understand the steps.
- **LI₂ (understand/conceptual)**: Second, while pursuing LI₁, the searcher may encounter the concept of “divide and conquer” and look for a definition to understand this concept.
- **LI₃ (understand/procedural)**: Third, to deepen their understanding of mergesort, the searcher may review its pseudocode (i.e., a different representation of the procedure).
- **LI₄ (understand/procedural)**: Fourth, the searcher may look for and review an example of quicksort to understand the steps.
- **LI₅ (apply/procedural)**: Fifth, to deepen their understanding of how the algorithm works, the searcher may decide to sort a list of numbers using quicksort.
- **LI₆ (analyze/procedural)**: Sixth, upon realizing that mergesort and quicksort are *both* “divide and conquer” algorithms, the searcher may review an article that explains how both algorithms use “divide and conquer”.

- **LI₇ (analyze/procedural)**: Seventh, the searcher may read an article that explains the similarities and differences between mergesort and quicksort.
- **LI₈ (evaluate/conceptual)**: Finally, the searcher may read an article that explains why mergesort is better than quicksort for large arrays, and may use this information as a rationale to judge that mergesort is more efficient in real-world scenarios.

As illustrated in the example above, A&K's taxonomy can be leveraged to categorize LOs and learning instances traversed along the pathway toward the objective. LOs and instances along the pathway can all be assigned to cells in A&K's taxonomy.

1.3 Study Overview and Research Questions

Participants in the study completed three search tasks with LOs situated at the intersection of *three* CPs (apply, evaluate, and create) and *three* knowledge types (factual, conceptual, and procedural). To analyze the pathways taken by participants toward an objective, we performed a qualitative analysis of search sessions based on participants' think-aloud comments and recorded search and note-taking activities. The study investigated three main research questions (RQ1-RQ3), which all centered on the pathways participants followed toward an objective.

A preliminary analysis of pathways found that participants primarily stayed within the *same* knowledge type as the LO they were asked to accomplish. As illustrated in the example in Table 1, pathways toward a procedural objective mostly involved learning instances focused on procedural knowledge. Therefore, our analysis of pathways focused exclusively on the *CPs* traversed along the pathways.

We investigate the following research questions:

- **RQ1**: What is the effect of the LO (i.e., CP and knowledge type) on the characteristics of pathways toward the objective?
- **RQ2**: What is the effect of the LO (i.e., CP and knowledge type) on the CP traversed along pathways toward the objective?
- **RQ3**: What are common and uncommon *transitions* between CPs traversed along pathways toward the objective?

In RQ1, we investigate the effects of the LO on two pathway characteristics: (1) pathway length (i.e., number of learning instances) and (2) pathway diversity (i.e., number of distinct CPs traversed). We investigate these differences from two perspectives. First, we explore pathway differences by conditioning on the CP of the objective. For example, are pathways longer or more diverse when the objective is to create versus apply? Second, we explore pathway differences by conditioning on the knowledge type of the objective. For example, are pathways longer or more diverse during procedural versus factual objectives?

In RQ2, we investigate the effects of the LO on the CPs traversed along the pathways. As in RQ1, we investigate these differences from two perspectives. First, we explore pathway differences by conditioning on the CP of the objective. For example, do pathways involve more analyze learning instances when the objective is to evaluate versus apply? Second, we explore pathway differences by conditioning on the knowledge type of the objective. For example, do pathways involve more analyze learning instances when the objective involves conceptual versus factual knowledge?

In RQ3, we investigate the types of CP *transitions* along pathways toward an objective. Again, we investigate this question from two perspectives. First, we consider common and uncommon transitions *irrespective of the objective*. For example, regardless of the objective, are searchers more likely to transition from simple to complex processes (e.g., understand to analyze) or more likely to transition from complex to simple processes (e.g., analyze to understand)?

Second, we consider common and uncommon transitions by *conditioning on the objective*. As described below (Section 1.4), in a previous article [45], we reported on results derived from the same lab study. Specifically, we reported on the effects of the LO (i.e., CP and knowledge type) on participants’ pre- and post-task perceptions and search behaviors. Our results found that the knowledge type of the objective had a *much stronger* effect than the CP of the objective. Based on these results, we decided to focus our RQ3 analysis by conditioning only on the knowledge type of the objective. For example, are searchers more likely to transition from simple to complex CPs during objectives involving procedural versus conceptual knowledge?

1.4 Extension of Prior Work

Our research in this article is an extension of our own prior work. In Urgo et al. [45], we reported on results from the same study. Specifically, we reported on the effects of the LO (i.e., CP and knowledge type) on: (1) pre-task perceptions, (2) post-task perceptions, and (3) search behaviors. Interestingly, the objective’s CP (apply vs. evaluate vs. create) had *no significant effects*. Conversely, the objective’s knowledge type (factual vs. conceptual vs. procedural) had several significant effects. First, factual objectives were perceived to require *less* cognitive activity along processes more complex than “remembering”. Second, conceptual objectives were perceived to require more “understanding” and “analyzing”. Finally, procedural objectives were perceived to require more “applying”, “evaluating”, and “creating”. In this article, we extend this prior work by analyzing the pathways taken by participants toward a specific objective. In Section 5, we describe how our results relate with those reported in Urgo et al. [45].

2 RELATED WORK

Our research builds on two areas of prior work: (1) search-as-learning and (2) understanding how task characteristics can influence search behaviors and outcomes.

2.1 Search-as-Learning

Studies in the area of search-as-learning have investigated a wide range of research questions. Some studies have investigated factors that influence learning during search. Specifically, studies have focused on how learning is impacted by characteristics of the searcher, search task, and search system. Additionally, studies have investigated how learning outcomes are related to specific search behaviors. In the following sections, we review key insights gained from these prior studies in search-as-learning. We focus primarily on key takeaways with respect to learning outcomes.¹

2.1.1 The Effects of User Characteristics on Learning. Several studies have investigated the effects of domain knowledge on learning during search [34, 38, 49]. O’Brien et al. [34] measured learning by asking participants to produce knowledge summaries before and after completing

¹In our review, we focus primarily on key takeaways with respect to learning outcomes. As a side note, prior studies have used a *wide* variety of methods to measure learning. Some studies have measured learning by administering pre- and post-tests with pre-defined correct answers, including: (1) true-or-false [15, 16, 22, 33, 37, 53], (2) multiple-choice [13, 15, 18, 19, 22, 41, 46], and (3) short-answer tests [1, 6, 10, 13, 20, 38, 39]. Other studies have asked participants to complete more *open-ended* exercises. Specifically, studies have measured learning by asking participants to: (1) list relevant key phrases and facts [4, 23]; (2) create visual representations of a domain [29]; (3) enumerate arguments for and against a specific proposition [14]; and (4) summarize their knowledge of a topic [1, 10, 13, 22, 26, 28, 34–36, 40, 49]. To assess learning from open-ended responses, studies have adopted grading strategies that involve: (1) counting relevant concepts or facts [1, 4, 10, 23, 35, 49]; (2) counting relevant pro/con arguments [14]; and (3) counting statements that show evidence of generalization or critical thinking [1, 10, 28, 34, 35, 40]. Finally, studies have also considered self-reported *perceptions of learning* [10, 15, 17, 18, 23, 29] and *behavioral measures* that are *assumed* to provide evidence of learning [8].

three search tasks on the same general topic. Compared to domain experts, novices had slightly greater improvements in their summary scores. One explanation is that novices uncovered more *new* information while searching. Willoughby et al. [49] asked participants to produce knowledge summaries on domains where they had high and low prior knowledge. Additionally, one group of participants was instructed to search for 30 minutes before producing their summaries and another group produced summaries without searching. Participants in the *search condition* produced summaries with more accurate facts. Interestingly, however, this effect was only found when participants had high prior knowledge. The authors hypothesized that participants with higher prior knowledge were able to search more effectively. Roy et al. [38] investigated the role of domain knowledge on learning *during* the search session. To this end, participants completed quick vocabulary learning assessments at regular intervals during the session. Prior knowledge influenced *when* participants had the greatest knowledge gains—toward the start of the session for participants with low prior knowledge and toward the end of the session for participants with high prior knowledge.

To summarize, results suggest a complex relationship between domain knowledge and learning during search. Specifically, the effects of domain knowledge are likely to depend on *other* factors, such as the complexity of the task domain. For example, within simple domains, novices may learn *more* because they simply start with less prior knowledge. Conversely, within complex domains, novices may learn *less* because they lack the prerequisite knowledge to search effectively.

Beyond domain knowledge, prior work has also considered the impact of individual abilities on learning during search. Pardi et al. [36] considered the impact of working memory capacity and reading comprehension ability. Learning was measured based on the number of relevant concepts included in knowledge summaries produced by participants before and after searching. Both abilities had a positive effect on learning.

2.1.2 The Effects of Task Characteristics on Learning. Several studies have investigated how task characteristics influence learning during search. Similar to our work, studies have leveraged the A&K taxonomy to study learning-oriented search tasks involving different CPs. Ghosh et al. [17] had participants complete tasks associated with the CP of understand, apply, analyze, and evaluate. Participants self-reported significant knowledge gains across all tasks. Additionally, participants were asked to select “action verbs” describing their mental activities during each task. Participants selected different action verbs for each task type—“define” for remember, “demonstrate” for apply, and “relate” for analyze and evaluate tasks. Kalyani and Gadiraju [22] had participants complete tasks associated with all six CPs from the A&K taxonomy. Learning was measured using closed-ended tests for simple tasks and open-ended tests for complex tasks. Participants had lower knowledge gains for complex tasks (i.e., apply < evaluate). Liu et al. [29] had participants complete two tasks of varying cognitive complexity: a *receptive* (i.e., remember or understand) task and a *critical* (i.e., evaluate) task. To measure learning, participants constructed mind maps (i.e., graphical domain representations) before each task, and modified their mind maps throughout the search session. During receptive tasks, participants made structural changes to their mind maps throughout the whole session. Conversely, during critical tasks, participants made more structural changes toward the end of the session.

Beyond task complexity, research has also studied learning during multi-session search. Liu et al. [30] had participants complete three subtasks on the same general topic. In the dependent subtask condition, all three subtasks built on each other. Conversely, in the parallel subtask condition, the three subtasks were largely independent (i.e., could be hypothetically done in any order). To measure learning, participants rated their familiarity with the general topic after each subtask. As expected, participants reported greater topic familiarity after each subtask. Interestingly,

however, this increase in topic familiarity plateaued faster in the parallel (vs. dependent) subtask condition, suggesting that searchers benefit from subtasks that build on each other.

2.1.3 The Effects of System Characteristics on Learning. Studies have also investigated how search systems and features can impact learning. Studies have considered different system characteristics: (1) the type of device used to search, (2) the presence of novel interface features and tools, and (3) the underlying retrieval algorithm.

Devices: Demaree et al. [14] compared learning outcomes between participants searching on a smartphone versus laptop computer. Participants were asked to gather information on a controversial topic and write an argumentative essay. Learning was measured by counting the number of pro and con arguments in the essay. While participants issued more queries while searching on a laptop, their learning outcomes were not significantly different across devices.

Interface Features and Tools: Wilson et al. [50] evaluated different interfaces for browsing a music collection. To measure learning, participants enumerated facts learned about the items in the collection. Results found a positive effect on learning from an interface that highlighted item metadata. Kammerer et al. [23] evaluated a system that enabled users to use social tags to filter search results. To measure learning, participants completed tests that required them to summarize their knowledge and recall domain-relevant keywords. Participants scored higher on both tests with the experimental system versus a baseline system without social tags.

Freund et al. [15] investigated the impact of two factors on participants’ reading comprehension of pre-selected articles: (1) whether articles were displayed in plain text versus HTML, which included distracting elements (e.g., ads), and (2) whether participants could add “sticky notes” to articles. Without the “sticky notes” tool, participants had higher reading comprehension scores in the plain text versus HTML condition. Conversely, with the “sticky notes” tool, participants performed equally well in both conditions.

Syed et al. [42] evaluated an experimental system that combined eye-tracking and an automatic question-generation feature. The system was designed to ask automatically generated questions about passages read by the searcher during the session. To measure learning, participants completed short-answer tests before each search task, immediately after, and one week later (to measure retention). Participants using the experimental system achieved higher retention scores than participants in the control group.

Qiu et al. [37] investigated the impact of two factors on learning and retention: (1) traditional vs. conversation search and (2) note taking. To measure learning and retention, participants completed the same test before each search task, immediately after, and three days later. Participants achieved greater knowledge gains in the traditional (vs. conversational) search condition, and this effect was stronger when participants could take notes. Interestingly, while knowledge gains were lower in the conversational search condition, participants had greater *retention rates*.

Roy et al. [39] investigated the impact of two tools that allowed participants to highlight passages and take notes. To measure learning, participants wrote post-task knowledge summaries that were analyzed based on the number of facts included and subtopics covered. Results found benefits from each tool in isolation. Specifically, the note-taking tool enabled participants to write summaries with more facts and the highlighting tool enabled participants to write summaries covering more subtopics. Interestingly, participants did not exhibit greater knowledge gains when using *both* tools, possibly due to cognitive overload.

Câmara et al. [6] evaluated different interface features to support learning: (1) displaying subtopics in the task domain and (2) displaying the searcher’s level of coverage across subtopics during the session. Interestingly, these novel features did not significantly improve learning. Instead, they influenced participants to explore more subtopics *superficially*. As evidence, when given

feedback about their topical coverage, participants viewed more search results but had shorter dwell times. Importantly, this trend suggests that feedback features can have unintended effects—they can influence searchers to pursue strategies that undermine their *depth* of learning.

To summarize, prior work has investigated a wide range of tools to improve learning during search. In general, results have found benefits from interfaces that: (1) convey more information about the items in the collection [23, 50], (2) enable searchers to annotate documents [15, 37, 39], and (3) enable searchers to self-assess their understanding of material read during the session [42]. On the other hand, results also suggest that tools can have unintended effects. For example, they can lead to cognitive overload [39] and encourage searchers to cover more information *superficially* [6].

Retrieval Algorithms: Early work by Hersh et al. [20] evaluated two search systems (i.e., Boolean vs. TF.IDF retrieval) based on their ability to help medical students improve their performance on a short-answer test. Both systems performed equally well, suggesting that people can achieve comparable learning outcomes using systems that afford very different search strategies. In the context of vocabulary learning, Syed and Collins-Thompson [41] evaluated a retrieval algorithm that favored documents with a greater *density* of target vocabulary words. Participants had better learning outcomes with the experimental versus baseline system. Weingart and Eickhoff [46] explored the impact of several well-established retrieval techniques on learning. To measure learning, participants completed multiple-choice tests after each task. Query expansion had a negative effect on learning, possibly due to topic drift from the original query. On the other hand, passage (vs. document) retrieval had a positive effect on learning, possibly because passages have a higher *density* of query-related content than whole documents.

2.1.4 The Relation between Search Behaviors and Learning. Finally, studies have investigated how specific search behaviors relate to learning outcomes. Several studies have investigated how learning outcomes are related to behaviors that could be potentially captured by a search system [1, 4, 10, 16, 28, 32, 35, 53]. Studies have found that searchers with better learning outcomes have a tendency to: (1) spend more time reading documents [10, 16, 32, 53]; (2) issue queries with more advanced and uncommon vocabulary [4, 10, 16]; (3) issue more diverse queries within the session [35]; (4) click on search results with longer titles [53]; (5) review more search results that are relevant [10] and novel [1]; and (6) visit sources that are more suitable to the task, such as encyclopedic sources during *receptive* tasks and Q&A sources during *critical* tasks [28].

Other studies have considered search behaviors that are more difficult to capture within existing search environments but are nonetheless insightful. Using eye-tracking, Bhattacharya and Gwizdka [4] found that participants with better learning outcomes had fewer eye regressions (i.e., less re-reading of text). Lei et al. [26] examined the search behaviors of 5th graders in the context of a mock school assignment involving video search. An analysis of post-search interviews found that students with better learning outcomes engaged in more metacognitive planning (e.g., setting goals), monitoring (e.g., tracking progress), and evaluating (e.g., reconsidering strategies).

Prior work has also investigated learning outcomes during collaborative search. Several studies have compared learning outcomes from searchers working individually versus in pairs [8, 35]. Chi et al. [8] found that participants working in pairs issued more complex queries (i.e., not issued by other participants), which was interpreted as evidence of learning. Palani et al. [35] found that participants working in pairs did *not* achieve better outcomes. However, the authors noted that the time limit imposed on search tasks may have prevented pairs from overcoming the overhead of collaboration. Xu et al. [52] found that collaborators with better learning outcomes engaged in more “division of labor” strategies and communicated more frequently.

2.2 The Effects of Task Characteristics on Behaviors and Outcomes

In our study, participants completed learning-oriented search tasks with objectives that varied by CP (apply, evaluate, and create) and knowledge type (factual, conceptual, and procedural). In this respect, our research builds on prior work aimed at understanding how task characteristics can influence searchers.

Search tasks have been studied from many different perspectives. Based on an extensive literature review, Li and Belkin [27] proposed a unifying framework for characterizing search tasks in terms of *generic facets* and *common attributes*. Generic facets relate to external factors (e.g., self-imposed vs. assigned). On the other hand, common attributes relate to the task itself, and include subjective attributes (e.g., difficulty) and objective attributes (e.g., complexity).

Our research builds heavily on prior work on the effects of task complexity, which is defined as an inherent property of the task. Studies have investigated task complexity from different perspectives. Wildemuth et al. [48] conducted an extensive review of the different ways task complexity has been characterized in prior work. A few common themes emerged from this review. Specifically, complex tasks involve: (1) more subtasks and/or subtopics; (2) greater uncertainty about aspects of the task (e.g., inputs and outputs); and (3) more complex mental processes.

Closely related to our research, prior studies have investigated task complexity from the perspective of *cognitive* complexity, which relates to the types of mental processes associated with the task. To this end, studies have leveraged the A&K taxonomy [5, 7, 21, 24, 44, 51]. Importantly, however, these studies have leveraged the CP dimension and *ignored* the knowledge type dimension. Results from these studies have found that cognitively complex tasks are perceived to be more difficult [5, 7, 21, 24, 51], require more search activity [5, 7, 21, 24, 44, 51], and lead to more *divergent* strategies by searchers performing the same task [24].

2.3 Contributions of Our Research

Our research in this article extends prior work in three important ways. First, as illustrated in our review, most “search-as-learning” studies have focused on learning *outcomes*. In this article, we focused on the search and learning *process*—the pathways followed by searchers toward an objective. Understanding the search and learning process provides insights about what searchers do and why. Additionally, it provides insights about potential search tools to support learning.

Second, several studies have leveraged A&K’s taxonomy to investigate how LOs (i.e., the goal of the task) can impact perceptions and behaviors. Importantly, prior studies have only leveraged the CP dimension and ignored the knowledge type dimension. In our study, we manipulated objectives by varying both the CP (apply, evaluate, and create) and knowledge type (factual, conceptual, procedural) of the objective.

Finally, while A&K’s taxonomy has been used to manipulate and study search tasks, it has not been used to understand the search and learning process. Therefore, as a methodological contribution, our study shows how the taxonomy can be leveraged to study search sessions from a learning-oriented perspective.

3 METHODS

To investigate RQ1-RQ3, we conducted a laboratory study with 36 participants (25 female). Participants were recruited using an opt-in mailing list of employees from our university. Participants included 18 student employees and 18 non-student employees, and their ages ranged from 19 to 61 ($M = 32.61$, $S.D. = 12.82$).




Fig. 1. Study protocol.

3.1 Study Overview

Protocol: The study protocol is illustrated in Figure 1 and proceeded as follows. After providing informed consent, participants completed a demographics questionnaire. Then, participants completed *three* experimental tasks that followed the same sequence of steps.

First, after reading the task description, participants completed a pre-task questionnaire. Next, participants completed the *search phase* of the task. During the search phase, participants were given a learning-oriented search task (Section 3.2) and asked to use a custom-built search system to find information and take notes in an external electronic document. The search system was implemented using the Bing Web Search API.² In order to investigate participants' pathways toward the LO of each task, the study used a think-aloud protocol [11]. Participants' think-aloud comments and search activities were audio/video recorded and later analyzed using qualitative techniques to gain insights about participants' pathways (Section 3.3). Participants were given a maximum of 15 minutes to complete the search phase and were alerted by the moderator when they had 5 minutes remaining. After the search phase, participants were given 2 minutes to review their notes and then completed a 2-minute *video demonstration* phase. During the video demonstration phase, participants were asked to provide a verbal response to the task's main question. Responses were video recorded by the moderator. Finally, after the video demonstration phase, participants were asked to complete a post-task questionnaire. The study session lasted about 1.5 hours and participants were given US\$40 for participating.

In this article, we focus on understanding participants' pathways along A&K's taxonomy toward the LO of the task. Thus, responses to the pre-/post-task questionnaires were not analyzed as part of this article. In Urgo et al. [45], we report on the effects of the LO on participants pre-/post-task perceptions. In Section 5, we discuss how our RQ1-RQ3 results resonate with those reported in Urgo et al. [45].

Think-aloud: In order to investigate participants' pathways toward the LO of each task, the study used a think-aloud protocol [11]. That is, participants were asked to narrate their thoughts as they searched for information and took notes. The study moderator reminded participants to think aloud if they were silent for an extended period.

Video Demonstration Phase: During the video demonstration phase of each task, participants were instructed to demonstrate (to whatever extent possible) their achievement of the task's LO. Participants were asked to produce a 2-minute response to the task's main question. Responses were video recorded by the study moderator. For example, for Task 3 in Section 3.2, participants were asked to verbally demonstrate their novel method for finding the mathematical center of a circle. Participants' responses were largely verbal. However, participants were also provided

²Given a query, the system returned results from four different verticals in different tabs: web, images, news, and video. Each vertical tab retrieved the top-50 results. Except for the images tab, each tab included pagination controls at the bottom and displayed 10 results per page. The images tab displayed all 50 thumbnails in one SERP using a grid (vs. list) layout. The Bing API was configured to retrieve results for the US-EN market, and we enabled safe-mode to filter inappropriate results.

with tangible materials as supplemental support: letter-size article, a large notepad on easel, pens, markers, pencils, eraser, geometric compass, protractor, rulers, and calculator. Participants were informed of (and able to review) these supplemental materials before working on the first experimental task of the study session. The main objective of the video demonstration phase was to encourage participants to learn and discourage them from satisficing. We believe that asking participants to produce a live demonstration of what they learned achieved this objective.

3.2 Tasks

A total of 27 tasks were constructed across three topical domains (art, finance, and science). Each domain was associated with nine tasks that varied across three CPs (apply, evaluate, and create) and three knowledge types (factual, conceptual, and procedural) from A&K’s taxonomy [3]. To keep the study design manageable, we limited ourselves to three CPs and three knowledge types.

In terms of CPs, we decided to omit remember, understand, and analyze. Our goal was to investigate the learning process during complex objectives, which are likely to require pursuing and achieving multiple learning-oriented subgoals. For this reason, we omitted the CPs of remember and understand (i.e., the least complex). Additionally, we decided to omit the CP of analyze for two reasons. First, we wanted to include objectives that would lead to *divergent* pathways. To this end, we decided to include evaluate and exclude analyze. Evaluation is closely linked to analysis. As noted by Anderson and Krathwohl, critiquing elements usually requires first understanding their relations (e.g., similarities and differences) [3, p. 79]. Second, Anderson and Krathwohl argue that analyzing is rarely the *ultimate* LO in and of itself. Instead, analyzing is “probably more defensible [as an LO] as a prelude to evaluating or creating.” [3, p. 79].

In terms of knowledge types, we omitted metacognitive knowledge because it is very different from the other three. Metacognitive knowledge involves internal (rather than external) knowledge about one’s own cognition.³ Next, we provide three example tasks from the science domain: (1) apply/factual, (2) evaluate/conceptual, and (3) create/procedural.

- (1) **Apply/Factual—Scenario:** You recently watched a TV show about the deepest part of the ocean. The show mentioned the depth of the deepest part of the ocean in meters. However, this number (in meters) did not quite give you clear “perspective” on just how deep this is. You want to get a more “tangible” appreciation for the depth of the deepest part of the ocean. **Task:** Use the height of the world’s tallest building as a unit to measure the deepest part of the ocean.
- (2) **Evaluate/Conceptual—Scenario:** During a recent trip to the National Air and Space Museum with your extended family, your younger cousin, who is in high school, said she is interested in better understanding how planes are able to fly. You are not very familiar with the principles behind the notion of lift, so when you get home you decide to do some investigating. After some initial research you notice that there are two predominant explanations of lift, Bernoulli’s principle and Newton’s laws of motion. **Task:** Determine, which best explains the notion of lift and why: Bernoulli’s principle or Newton’s laws of motion? Provide a well-reasoned, logical argument to support your explanation.
- (3) **Create/Procedural—Scenario:** You are building a firepit in your backyard. You have constructed a large circle so that chairs can fit around the firepit. You have not yet dug the firepit because you want to be sure that it is positioned precisely in the center of the circle.

³We also consulted with a cognitive science researcher who advised us to exclude metacognitive knowledge. Understanding metacognitive activities during search-as-learning is an important area for future work.

Task: Explore different methods for finding the mathematical center of a circle, then create a novel method for finding the mathematical center of your firepit circle. The method can be completely different from those you find, a combination of methods, or a method you find with slight variations.

As shown, each task was situated in a background scenario that motivated the learning-oriented search task. Task 1 is apply/factual because it requires applying one fact (i.e., the height of the world’s tallest building) to gain appreciation of another fact (i.e., the depth of the deepest part of the ocean). Task 2 is evaluate/conceptual because it requires determining which concept (i.e., Bernoulli’s principle vs. Newton’s laws of motion) best explains a phenomenon (i.e., lift). Finally, Task 3 is create/procedural because it requires creating a new procedure (i.e., a new method for finding the mathematical center of a circle). Table 2 summarizes the main objectives of our 27 tasks. The full tasks are available online: <https://www.kelseyurgo.com/tois-pathways/>.

Task Assignment: Each participant completed three tasks from the *same* domain (i.e., arts, finance, or science). Domains were assigned to participants such that 12 participants completed tasks from each domain. Additionally, each participant completed tasks that varied across all three CPs (i.e., apply, evaluate, and create) and all three knowledge types (i.e., factual, conceptual, and procedural). For example, participant P1 completed the following tasks: (1) apply/factual/art, (2) evaluate/conceptual/art, and (3) create/procedural/art. The ordering of tasks was rotated such that every participant experienced our three CP and knowledge types in a different order (i.e., 6 CP orders $\times$ 6 KT orders = 36 participants).

3.3 Qualitative Coding of Pathways

To address RQ1–RQ3, we conducted a qualitative analysis of participants’ recorded search sessions, which included their think-aloud comments and search/note-taking activities. The pathway annotation task involved two separate processes.

Identifying Learning Instances (LIs): The first annotation process involved representing each search session as a sequence of *learning instances* (LIs). We define LIs as instances during the search session where the participant either: (1) set forth a new learning-oriented subgoal or (2) serendipitously learned something new and meaningful toward the task’s objective. This definition allowed us to account for instances where the participant mentioned pursuing a new specific subgoal (e.g., “First, I need to find the world’s tallest building.”), as well as moments when the participant mentioned serendipitously learning something new without planning (e.g., “Oh, this says that XYZ was the world’s tallest building in 2003. Maybe there’s a taller one in 2020.”). To represent each session as a sequence of LIs, one of the authors reviewed all recorded search sessions (i.e., screen activities and think-aloud comments) and marked the timestamp of each LI identified. While recording these timestamps, the question in the author’s mind was: At this point, has the searcher set forth a new subgoal or has the searcher learned something new and relevant to the task? Most LIs occurred during points where the participant issued a new query, made a new note in the external document, or engaged with information (e.g., a search result) while thinking aloud. Table 3 provides a full example pathway. Each row is a distinct LI. The motivations for marking a new LI are described in the “Action/Comment” column.

Annotating Learning Instances (LIs): The second annotation process involved classifying LIs into cells in A&K’s taxonomy—assigning each LI to a specific CP and knowledge type. To classify LIs, we developed a coding guide based on A&K’s book [3], which provides definitions and examples of LOs for each CP and knowledge type. The full coding guide is available in Appendix A.1.

The coding guide was developed iteratively using search sessions (i.e., LI sequences) from three participants. First, both of the authors worked together to derive an initial coding guide. Next,

Table 2. LOs for Our 27 Tasks

Knowledge Dimension	Cognitive Process Dimension		
	Apply	Evaluate	Create
Factual	S: Use the height of tallest building to measure the deepest part of the ocean.	S: Which best describes the deepest part of the ocean: The Mariana Trench or the Challenger deep?	S: Generate a table of the ocean’s 10 deepest parts including useful and interesting criteria.
	A: Use the price of the most expensive painting ever sold to express how many Honda civics you could purchase at the same price.	A: What is the primary reason for the high price of the most expensive painting ever sold?	A: Generate a table of the world’s 10 most expensive paintings including their most unique and interesting characteristics.
	F: Use the average retirement savings in the U.S. to calculate how many years one could live on a \$40,000 per year budget.	F: Argue whether or not average retirement savings in U.S. is sufficient for determining if citizens are well-prepared for retirement.	F: Generate a table of U.S. average retirement savings across demographics you find interesting.
Conceptual	S: Use Bernoulli’s principle to explain how an airplane flies.	S: Which best explains the notion of lift and why: Bernoulli principle or Newton’s laws of motion?	S: Construct a diagram to illustrate the differences between Bernoulli’s principle and Newton’s laws of motion applied to lift.
	A: Select which of the following paintings exemplify the “automatism” artistic style and note the characteristics indicative of “automatism”.	A: Which artistic movement is more closely related to “automatism” and why: Dadaism or Surrealism?	A: Construct a graphic that encompasses all artistic forms of “automatism” and expresses how each are similar and different.
	F: Use your knowledge of mortgages to label the following monthly mortgage payment breakdown.	F: Which component of a mortgage is most critical for first-time home buyers to negotiate and why: down payment, interest rate, or date of maturity?	F: Construct a sample mortgage that would be optimal for your monthly budget and overall savings.
Procedural	S: Find the mathematical center of the firepit circle below.	S: Find two methods for finding the mathematical center of a circle and choose the best for finding the center of your new firepit and explain why.	S: Create a novel method for finding a mathematical center of a firepit circle.
	A: Find a method for making a paper airplane and make that paper airplane.	A: Choose which method of paper airplane making would be most fun and interesting and explain why.	A: Create a novel method for making a paper airplane.
	F: Use the 50/30/20 budget to categorize the following sample of monthly expenses.	F: Explain which budget would be most effective for saving money for your trip and explain why.	F: Create a novel budgeting method that is most useful to you in saving money for your vacation.

Sub-rows correspond to the domains of science (S), arts (A), and finance (F), respectively.

both authors independently coded all three sessions from one participant, discussed disagreements, and refined the coding guide. After repeating this process with all three participants, the authors arrived at a final version of the coding guide.

As we developed the coding guide, we found it relatively straightforward to classify LIs into a knowledge type (i.e., factual, conceptual, and procedural). We found it more challenging to classify LIs into a CP (i.e., remember, understand, apply, analyze, evaluate, and create). To alleviate this challenge, the coding guide primarily focuses on criteria for classifying LIs by CP. Additionally, we found it helpful to use *different wording depending on the LI’s knowledge type*. To illustrate, for

Table 3. Examples of Learning Instances (LI) with Associated Cell Assignment and Justification

LI	Action/Comment	Cell Annotation	Justification
1	automatism (query)	Understand/Procedural	Searcher is gathering information to summarize (understand) automatism (conceptual).
2	"The first one that pops up mentions Surrealist automatism, so maybe it has more to do with Surrealism"	Analyze/Conceptual	Searcher is differentiating and structuring (analyze) automatism (conceptual) within the context of surrealism (conceptual).
3	"Maybe we should look at pictures"	Understand/Conceptual	Searcher is trying to exemplify (understand) automatism (conceptual).
4	"It says it was used to express the subconscious...hand is allowed to randomly move across the paper"	Remember/Conceptual	Searcher is reading a definition (remember) of automatism (conceptual).
5	subconscious expression (note)	Understand/Conceptual	Searcher is summarizing (understand) a characteristic of automatism (conceptual).
6	"Talks about taking material from the subconscious and putting it into art"	Understand/Conceptual	Searcher is summarizing (understand) a characteristic of automatism (conceptual).
7	"Talks about Freud"	Remember/Factual	Searcher is reading (remember) a fact (factual) related to automatism.
8	"Free association"	Remember/Conceptual	Searcher is reading (remember) about free association (conceptual).
9	"This also mentions Surrealism...so let's see what Surrealism and Dadaism are about"	Analyze/Conceptual	Searcher is differentiating (analyze) between surrealism (conceptual) and dadaism (conceptual).
10	dadaism art (query)	Understand/Conceptual	Searcher is exploring information to summarize (understand) Dada art (conceptual).
11	"An art movement formed during the first world war...a negative reaction to the horrors of the war"	Remember Conceptual	Searcher is reading an overview (remember) of Dadaism (conceptual).
12	"I can compare images of Dadaism"	Analyze/Conceptual	Searcher is differentiating (analyze) representations of Dadaism from Surrealism (conceptual).
13	surrealism art (query)	Analyze/Conceptual	Searcher is trying to differentiate (analyze) representations of Surrealist art from Dada art (conceptual).
14	"I really want to know how to tell the difference"	Analyze/Conceptual	Searcher is trying to differentiate (analyze) representations of Surrealist art from Dada art (conceptual).
15	automatism art (query)	Analyze/Conceptual	Searcher is trying to differentiate (analyze) representations of automatism from Surrealist art and Dada art (conceptual).
16	"Automatism doesn't really have like portraits or anything like the other ones"	Analyze/Conceptual	Searcher is differentiating (analyze) representations of automatism from Surrealist art and Dada art (conceptual).
17	surrealism art (query)	Understand/Conceptual	Searcher is trying to exemplify (understand) Surrealist art (conceptual).
18	"Going to look for more facts about Surrealism"	Remember Conceptual	Searcher is trying to read (remember) specific information about Surrealism (conceptual).

(Continued)

Table 3. Continued

LI	Action/Comment	Cell Annotation	Justification
19	unconscious mind (note)	Understand/Conceptual	Searcher is gathering information to summarize (understand) Surrealism (conceptual).
20	“Has to do with the unconscious mind which is very similar to automatism”	Analyze/Conceptual	Searcher is differentiating (analyze) between Surrealism and automatism (conceptual).
21	“Talks about free association as well”	Analyze/Conceptual	Searcher is differentiating (analyze) between Surrealism and automatism (conceptual).
22	“Not all of this art is abstract”	Understand/Conceptual	Searcher is identifying (understand) characteristics of Surrealist art (conceptual)
23	“Dadaism was reaction to the horrors of the war so maybe Dadaism is less abstract?”	Analyze/Conceptual	Searcher is differentiating (analyze) Dadaism from Surrealism (conceptual).
24	“Have to do with more what they were experiencing consciously rather than subconsciously”	Analyze/Conceptual	Searcher is differentiating (analyze) Dadaism from Surrealism (conceptual).
25	“I think that automatism and Surrealism have more in common”	Evaluate/Conceptual	Searcher is judging (evaluate) that automatism is more closely related to Surrealism than Dadaism (conceptual).
26	similarities between automatism and surrealism (query)	Analyze/Conceptual	Searcher is differentiating (analyze) between automatism and Surrealism (conceptual).
27	processes not under conscious control (note)	Understand/Conceptual	Searcher is summarizing (understand) characteristic of automatism (conceptual).
28	“I think that’s the biggest similarity”	Analyze/Conceptual	Searcher is differentiating (analyze) between automatism and Surrealism (conceptual).
29	interpretation of dreams (note)	Understand/Conceptual	Searcher is summarizing (understand) characteristic of automatism (conceptual).
30	automatism artists (query)	Understand/Factual	Searcher is exploring (understand) artists (factual) that used automatism.
31	technique used by surrealist painters (note)	Understand/Conceptual	Searcher is summarizing (understand) characteristic of automatism (conceptual).
32	“I don’t see how it’s not the same thing”	Analyze/Conceptual	Searcher is differentiating (analyze) between automatism and Surrealism (conceptual).
33	influenced by freud (note)	Understand/Factual	Searcher is summarizing (understand) isolated unit of information (factual) associated with automatism.
34	dadaism (query)	Understand/Conceptual	Searcher is gathering information to summarize (understand) Dadaism (conceptual).
35	“I don’t really know what to describe Dadaism as”	Understand/Conceptual	Searcher is gathering information to describe (understand) Dadaism (conceptual) in own words.

(Continued)

Table 3. Continued

LI	Action/Comment	Cell Annotation	Justification
36	"World War I, wasn't an artistic style"	Understand/Conceptual	Searcher is summarizing (understand) information about Dadaism (conceptual).

For this pathway, the task's LO (evaluate/conceptual) involved deciding which artistic movement is most closely related to Automatism: Dadaism or Surrealism?

classifying LIs as "understand", we included the following *analogous* criteria: (1) "Restates fact in own words." (factual); (2) "Summarizes definition of concept in own words." (conceptual); and (3) "Summarizes steps of procedure in own words." (procedural).

To validate the final version of our coding scheme, both authors independently coded all sessions (i.e., LI sequences) from an additional six participants (about 17% of the data). The Cohen's Kappa agreement was $\kappa = 0.809$, which is considered "almost perfect" agreement [25]. Given this high level of agreement, one of the authors (re-)coded all sessions from the remaining 30 participants.

Table 3 describes an example annotated pathway, including the action and/or comment that triggered the LI, the A&K cell associated with the LI, and an explanation/justification of the cell annotation.

4 RESULTS

In RQ1-RQ3, we analyze the pathways followed by participants toward a specific type of LO. A pathway is defined as a sequence of LIs that were each manually assigned to a cell from A&K's taxonomy (i.e., CP and knowledge type).

A preliminary analysis of pathways found that participants primarily stayed within the *same* knowledge type as the LO's knowledge type. For example, pathways toward factual objectives mostly involved factual LIs. The percentage of LIs associated with the *same* knowledge type as the objective was 97% for factual objectives, 85% for conceptual objectives, and 97% for procedural objectives. Therefore, in our analyses for RQ1-RQ3, we focus *exclusively* on the CPs associated with LIs along the pathway. In other words, while A&K's taxonomy involves *two* orthogonal dimensions (i.e., CP and knowledge type), pathways were analyzed from a *one*-dimensional perspective (i.e., CP).⁴

Pathways were analyzed from three perspectives: (RQ1) pathway length and number of distinct CPs traversed; (RQ2) types of CPs traversed; and (RQ3) types of transitions between CPs conditioned on the knowledge type of the LO. In RQ1 and RQ2, we investigate differences based on the LO's CP (apply vs. evaluate. vs. create) and knowledge type (factual vs. conceptual. vs. procedural). To test for significant differences, we used one-way ANOVAs with Bonferroni-corrected post-hoc comparisons.

4.1 RQ1: Pathway Characteristics

In RQ1, we investigate the characteristics of pathways *conditioned* on the LO's CP (Table 4) and knowledge type (Table 5). Tables 4 and 5 show the average pathway length (i.e., average number of LIs traversed along the pathway) and the average number of *distinct* CPs traversed.

⁴Our analysis of pathways found that participants primarily stayed within the same knowledge type as the given objective. It is important to emphasize that LIs were classified based on the *type of knowledge* participants were aiming to acquire and *not* solely the *type of information* participants engaged with in pursuit of the subgoal. To illustrate, imagine an LI in which a searcher is trying to get a basic understanding of Bernoulli's principle. Now, suppose the searcher encounters and struggles to internalize the following statement: "Pressure decreases when the speed of a fluid increases." This, of course, is a *factual* statement. However, we would classify this LI as "understand/conceptual" because the current subgoal is to understand the *concept* of Bernoulli's principle.

Table 4. The Effects of the LO’s- CP on Pathway Length and Distinct CPs Traversed (Means $\pm$ 95% CIs)

	Apply LO	Evaluate LO	Create LO
Average Length	15.97 $\pm$ 2.45	16.89 $\pm$ 2.52	21.17 $\pm$ 4.04
Average Distinct CPs	3.72 $\pm$ 0.28	3.75 $\pm$ 0.26	4.03 $\pm$ 0.34

Table 5. The Effects of the LO’s Knowledge Type on Pathway Length and Distinct CPs Traversed (Means $\pm$ 95% CIs)

	Factual LO	Conceptual LO	Procedural LO
Average Length	19.25 $\pm$ 3.97	19.64 $\pm$ 2.88	15.14 $\pm$ 2.22
Average Distinct CPs [‡] (C < P)	3.83 $\pm$ 0.24	3.53 $\pm$ 0.25	4.14 $\pm$ 0.35

[‡] denotes rows with significant differences across knowledge types: factual (F), conceptual (C), and procedural (P).

First, we compare pathway characteristics conditioned on the LO’s CP (Table 4). The LO’s CP did not have a significant effect on the pathway length. However, as might be expected, there is an upward trend as the LO increases in complexity from apply ($M = 15.97$) to evaluate ($M = 16.89$) to create ($M = 21.17$). In other words, more complex LOs had longer pathways (i.e., more learning instances). Similarly, the LO’s CP did not have a significant effect on the number of distinct CPs traversed along the pathway. However, there is a small upward trend as the complexity of the LO increases. Create objectives had pathways with slightly more distinct CPs than apply and evaluate objectives.

Next, we compare pathway characteristics conditioned on the LO’s knowledge type (Table 5). The LO’s knowledge type did not have a significant effect on the pathway length. Conversely, the LO’s knowledge type did have a significant effect on the number of distinct CPs traversed along the pathway ($F(2, 105) = 4.75, p < .05$). Procedural LOs had significantly more distinct CPs than conceptual LOs ($p < .01$). We discuss these trends in Section 5.

4.2 RQ2: Effects of LO on CPs

In RQ2, we investigate the effects of the LO on the types of CPs traversed along the pathway. For example, are apply LIs more common for some LOs than others? Similar to RQ1, we explore differences by conditioning on the LO’s CP (Table 6) and by conditioning on the LO’s knowledge type (Table 7). Tables 6 and 7 show the average number of LIs (per pathway) associated with each CP. For example, during apply LOs (Table 6), pathways had 4.58 remember LIs on average. During conceptual LOs (Table 7), pathways had 9.61 understand LIs on average.

We start by discussing our results conditioned on the LO’s CP (Table 6). Our results found five main trends. First, irrespective of the LO’s CP, remember and understand were the most frequent CPs traversed. Second, the LO’s CPs had a significant effect on the number of remember LIs traversed ($F(2, 105) = 5.66, p < .005$). Remember LIs were significantly more common during tasks with an objective to create versus apply ($p < .05$) or evaluate ($p < .01$). Third, the LO’s CP had a significant effect on the number of apply LIs traversed ($F(2, 105) = 9.47, p < .001$). Apply LIs were significantly more common during tasks with an objective to apply versus evaluate ($p < .001$) or create ($p < .005$). Fourth, the LO’s CP had a significant effect on the number of create LIs traversed ($F(2, 105) = 4.98, p < .01$). Create LIs were significantly more common during tasks with an objective to create versus apply ($p < .01$). Finally, while not statistically significant, analyze and evaluate LIs were more common during tasks with an objective to evaluate versus apply or create.

Next, we discuss our results conditioned on the LO’s knowledge type (Table 7). Our results found five main trends. First, irrespective of the LO’s knowledge type, remember and understand

Table 6. The Effects of the LO's CP on the Number of LIs (Per Pathway) Associated with Each CP (Means $\pm$ 95% CIs)

CP	Apply LO	Evaluate LO	Create LO
Remember [‡] (A, E < C)	4.58 $\pm$ 1.28	4.17 $\pm$ 1.26	8.33 $\pm$ 2.88
Understand	6.19 $\pm$ 1.36	6.19 $\pm$ 1.76	6.83 $\pm$ 2.35
Apply [‡] (A > E, C)	2.03 $\pm$ 0.89	0.47 $\pm$ 0.37	0.53 $\pm$ 0.28
Analyze	1.64 $\pm$ 0.76	3.14 $\pm$ 1.16	2.78 $\pm$ 0.84
Evaluate	1.58 $\pm$ 0.77	2.67 $\pm$ 0.55	1.92 $\pm$ 0.83
Create [‡] (A < C)	0.08 $\pm$ 0.13	0.25 $\pm$ 0.32	0.78 $\pm$ 0.46

[‡] denotes rows with significant differences across CPs: apply (A), evaluate (E), and create (C).

Table 7. The Effects of the LO's Knowledge Type on the Number of LIs (Per Pathway) Associated with Each CP (Means $\pm$ 95% CIs)

CP	Factual LO	Conceptual LO	Procedural LO
Remember [‡] (F > C, P)	9.22 $\pm$ 2.63	5.13 $\pm$ 1.56	2.72 $\pm$ 0.93
Understand [‡] (C > F, P)	4.03 $\pm$ 2.15	9.61 $\pm$ 1.81	5.44 $\pm$ 1.01
Apply	0.89 $\pm$ 0.56	1.39 $\pm$ 0.81	0.86 $\pm$ 0.45
Analyze	2.58 $\pm$ 0.95	2.50 $\pm$ 1.05	2.36 $\pm$ 0.88
Evaluate [‡] (C < F, P)	2.44 $\pm$ 0.82	0.89 $\pm$ 0.45	2.83 $\pm$ 0.74
Create [‡] (F, C < P)	0.08 $\pm$ 0.10	0.11 $\pm$ 0.18	0.92 $\pm$ 0.51

[‡] denotes rows with significant differences across knowledge types: factual (F), conceptual (C), and procedural (P).

were among the most frequent CPs traversed. Second, the LO's knowledge type had a significant effect on the number of remember LIs traversed ($F(2, 105) = 13.08, p < .001$). Remember LIs were significantly more common during tasks with an objective involving factual versus conceptual ($p < .01$) or procedural ($p < .001$) knowledge. Third, the LO's knowledge type had a significant effect on the number of understand LIs traversed ($F(2, 105) = 11.69, p < .001$). Understand LIs were significantly more common during tasks with an objective involving conceptual versus factual ($p < .001$) or procedural ($p < .005$) knowledge. Fourth, the LO's knowledge type had a significant effect on the number of evaluate LIs traversed ($F(2, 105) = 9.32, p < .001$). Evaluate LIs were significantly *less* common during tasks with an objective involving conceptual versus factual ($p < .005$) or procedural ($p < .001$) knowledge. Fifth, the LO's knowledge type had a significant effect on the number of create LIs traversed ($F(2, 105) = 9.10, p < .001$). Create LIs were significantly more common during tasks with an objective involving procedural versus factual ($p < .001$) or conceptual ($p < .005$) knowledge.

4.3 RQ3: Transitions Between CPs

Participants completed three search tasks, each with a LO associated with a specific knowledge type—factual, conceptual, or procedural knowledge. In RQ3, we investigate CP transitions *conditioned* on the knowledge type of the LO. We discuss our RQ3 results from two perspectives. First, we discuss common trends irrespective of the objective's knowledge type. Second, we discuss trends that are *unique* to LOs involving a specific knowledge type (factual vs. conceptual vs. procedural knowledge).

Table 8(a)–(c) show the transition probabilities between CPs along pathways toward a LO involving factual knowledge (Table 8(a)), conceptual knowledge (Table 8(b)), and procedural knowledge (Table 8(c)). Table 8(a)–(c) are Markov matrices. The values along each row correspond to

Table 8. Each Markov Matrix Shows the Transition Probabilities between CPs in Learning Pathways

CP	Rem.	Und.	Apply	Ana.	Eval.	Create	End	Count
Start	0.47	0.44*	0.00	0.06	0.03	0.00	0.00	36
Rem.	0.63*	0.11	0.03	0.13	0.08	0.00	0.02	332
Und.	0.30	0.35*	0.03	0.15*	0.12	0.01	0.04	145
Apply	0.25	0.00	0.39*	0.07	0.11	0.00	0.18	28
Ana.	0.37	0.19*	0.01	0.21*	0.14	0.01	0.07	97
Eval.	0.20	0.26*	0.03	0.10	0.27*	0.00	0.13	88
Create	0.33	0.33*	0.00	0.00	0.00	0.00	0.33	3

(a) Transition probabilities between CPs across *factual* pathways (i.e., conditioned on the knowledge type of the LO).

CP	Rem.	Und.	Apply	Ana.	Eval.	Create	End	Count
Start	0.06	0.86*	0.03	0.06	0.00	0.00	0.00	36
Rem.	0.38*	0.44	0.06	0.05	0.02	0.01	0.04	185
Und.	0.25	0.46*	0.07	0.14*	0.04	0.01	0.04	346
Apply	0.18	0.50*	0.22*	0.04	0.02	0.00	0.04	50
Ana.	0.17	0.43*	0.03	0.21*	0.08	0.00	0.08	90
Eval.	0.06	0.22*	0.03	0.34	0.19*	0.00	0.16	32
Create	0.00	0.75*	0.00	0.00	0.00	0.00	0.25	4

(b) Transition probabilities between CPs across *conceptual* pathways (i.e., conditioned on the knowledge type of the LO).

CP	Rem.	Und.	Apply	Ana.	Eval.	Create	End	Count
Start	0.00	0.97*	0.00	0.00	0.03	0.00	0.00	36
Rem.	0.22*	0.31	0.07	0.15	0.15	0.05	0.04	98
Und.	0.20	0.31*	0.05	0.17*	0.19	0.05	0.03	196
Apply	0.13	0.26*	0.19*	0.03	0.19	0.06	0.13	31
Ana.	0.15	0.22*	0.00	0.25*	0.24	0.05	0.09	85
Eval.	0.15	0.32*	0.06	0.13	0.21*	0.06	0.08	102
Create	0.12	0.33*	0.06	0.06	0.03	0.18	0.21	33

(c) Transition probabilities between CPs across *procedural* pathways (i.e., conditioned on the knowledge type of the LO).

Transition probabilities marked with * are *common* across LOs. Transition probabilities in bold are *unique* to the particular LO shown in that sub-table (each sub-table corresponds to a particular objective knowledge type).

the probabilities of transitioning from one type of LI to another and therefore sum to 1. A “start” row and “end” column have been added as additional states. The values along the start row correspond to the probabilities of starting the pathway with a specific CP. The values in the end column correspond to the probabilities of ending the pathway with a specific CP. The last column shows the raw counts for each state across pathways conditioned on the knowledge type of the objective. Cells along the diagonal (shown in grey) are transitions to the same CP. Cells above the diagonal, and underneath the grey line, are transitions to a *more* complex CP. We refer to these as transition “upshifts”. Cells below the diagonal are transitions to a *less* complex CP. We refer to these as transition “downshifts”.

In our discussion of trends, we focus on high versus low transition probabilities. Thus, an important question is: What is a high transition probability? In other words, what is a logical threshold to distinguish between high and low transition probabilities? As shown in Table 8(a)–(c), each state can transition to seven states (i.e., six CP states + the end state). If all transitions were equally likely, then all transition probabilities would be 0.14 (i.e., 1/7). Therefore, we consider transition

probabilities ≥ 0.14 as “high” transition probabilities and transition probabilities < 0.14 as “low” transition probabilities.

Common Trends: First, we discuss common trends irrespective of the objective’s knowledge type. In Table 8(a)–(c), common trends are marked with an asterisk (*). Our results found four common trends. First, understand is a common starting point. All “Start” rows have high transition probabilities to understand. Second, downshifts to understand (under the diagonal in the understand column) are generally common. As shown in Figure 8(a), the single exception is for LOs involving factual knowledge. Factual LOs had a low transition probability from apply to understand. Later, we discuss possible explanations for this lower transition probability. Third, most CPs had high transition probabilities back to themselves. These probabilities are shown along the diagonal of each matrix, highlighted in grey. As the exception, create-to-create transitions were only common for procedural LOs (i.e., not factual nor conceptual objectives). Finally, transitions from understand to analyze were generally common.

Next, we describe trends unique to factual, conceptual, and procedural LOs. In Tables 8(a)–(c), unique trends are shown in bold.

Trends Unique to Factual LOs: Our results found four trends unique to factual LOs (Table 9(a)). First, starting with remember was likely only for factual LOs. Second, downshifts to remember from *all* CPs (remember column in Table 8(a)) were likely only for factual LOs. In other words, for conceptual and procedural LOs, downshifts to remember were likely from *some* CPs but not others. Conversely, downshifts to remember were likely from *all* CPs during factual LOs. Third, transitions from remember to understand were common for conceptual and procedural LOs, but *uncommon* for factual LOs. Fourth, ending with apply was likely only for factual LOs.

Trends Unique to Conceptual LOs: Our results found three trends unique to conceptual LOs (Table 8(b)). First, transitions from evaluate to analyze were likely only for conceptual LOs. Second, transitions from analyze to evaluate were common for factual and procedural LOs, but *uncommon* for conceptual LOs. Third, ending with evaluate was likely only for conceptual LOs.

Trends Unique to Procedural LOs: Our results found three trends unique to procedural LOs (Table 8(c)). First, upshifts were much more common for procedural LOs. Procedural LOs had seven likely upshifts compared to only two for factual and two for conceptual LOs. Four of these seven upshifts are unique to procedural LOs: (1) remember to analyze, (2) remember to evaluate, (3) understand to evaluate, and (4) apply to evaluate. Second, transitions from create to create were likely only for procedural LOs. Finally, ending with create was more common for procedural LOs. Factual and conceptual LOs had high transition probabilities from create to the end state. However, in both cases, these high transition probabilities are based on a single occurrence of create at the end of the pathway (see “Count” column in Table 8(a) and (b)). Therefore, we do not consider this trend as “common” to all LOs, but rather a likely transition specific to procedural LOs.

5 DISCUSSION

In this section, we summarize our findings and compare them to findings from prior work.

5.1 RQ1: Effects on Pathway Length and Diversity of CPs Traversed

In RQ1, we investigate the effects of the LO (i.e., CP and knowledge type) on the pathway length and distinct CPs traversed.

Effects of the LO’s CP: The LO’s CP did *not* have a significant effect on the pathway length (number of LIs traversed) nor the number of distinct CPs traversed along the pathway (Table 4). In Urgo et al. [45], we reported on results from the same study. Specifically, we reported on the effects of the LO (i.e., CP and knowledge type) on participants’ pre-/post-task perceptions and search behaviors. The objective’s CP did not significantly affect any outcome related to

participants’ perceptions nor behaviors. Thus, our RQ1 results in the current article are consistent with those reported in Urgo et al. [45].

A reasonable follow-up question is: Why did the objective’s CP not have a significant effect on the pathway length nor diversity? One possible explanation stems from our choice of LOs considered in the study. To keep the study design manageable, we considered LOs associated with the CPs of apply, evaluate, and create, which have mid-to-high levels of complexity. We might have observed greater (and significant) differences on the pathway length and diversity had we *also* considered objectives associated with low-complexity processes (i.e., remember and understand).

While the trend was not significant, our RQ1 results found that more cognitively complex objectives had longer (and slightly more diverse) pathways (apply < evaluate < create). In general, this trend resonates with prior studies, which have consistently found that more cognitively complex objectives require more search activity [5, 7, 21, 24, 44, 51]. Thus, our RQ1 results suggest that some of this increase in search activity may be due to the objective requiring more learning instances along the pathway.

Effects of the LO’s Knowledge Type: The LO’s knowledge type had a significant effect on the number of distinct CPs traversed along the pathway (Table 5). Specifically, procedural objectives had pathways involving a more diverse set of CPs. Based on our RQ2 results (discussed in Section 5.2), this trend is due to procedural objectives having more create LIs. As shown in the last row in Table 7, procedural objectives had about 1.0 create LIs on average. In contrast, factual and conceptual objectives had about 0.10 create LIs on average.

In our qualitative coding of pathways, we observed that procedural objectives involved more creative subgoals. Specifically, during procedural objectives, create LIs included instances of the participant: (1) simplifying a procedure by skipping steps; (2) modifying steps to fit the given scenario; (3) changing the implementation of a procedure by using different materials (i.e., those readily available); (4) combining steps from multiple procedures to develop a new procedure; (5) using concepts as inspiration to develop a new procedure; and (6) using innovative techniques to improve the accuracy of a procedure. The presence of more create LIs during procedural objectives is a trend that we also observed in our RQ3 results (discussed in Section 5.3).

In general, there are several possible reasons for why procedural objectives had more create LIs. First, procedural knowledge is knowledge about how to perform a specific task. Tasks can often be accomplished in many different ways. For example, there are many ways for finding the center of a circle, making a paper airplane, or budgeting expenses. Therefore, procedural objectives may often provide searchers with the *flexibility* to engage in creative learning processes (e.g., modify procedures). Second, searchers may often have individual preferences that influence them to engage in creative processes. For example, searchers may be inclined to skip or modify steps based on their individual skills or prior knowledge. Finally, searchers may have situational constraints that require them to engage in creative processes. Such constraints may include temporal constraints, available resources, and success criteria (e.g., efficiency vs. accuracy).

5.2 RQ2: Effects on CPs Traversed

In RQ2, we investigate the effects of the LO (i.e., CP and knowledge type) on the types of CPs traversed along the pathway.

Effects of the LO’s CP: We begin by discussing our results conditioned on the objective’s CP (Table 6). Our results found four main trends.

First, irrespective of the objective’s CP, remember and understand LIs (i.e., the simplest CPs) were the most common. This indicates that, regardless of the LO, remembering and understanding are frequent and important activities that support more complex learning subgoals. For example, before *analyzing* the relations between concepts A and B, it is necessary to first *understand* A and B

in isolation. This process may involve several iterations—exploring definitions of A, examples of A, definitions of B, examples of B, and so on. Additionally, understanding A and B may involve exploring definitions and examples from *different* perspectives. In terms of definitions, searchers may engage with textual overviews, visual representations, and even mathematical formulas. In terms of examples, searchers may engage with different types of examples in order to infer common themes.

Second, remember LIs were significantly more common during create LOs. We believe this may be an unintended effect of our three create/factual tasks (see Table 2), which asked participants to create a table of facts (i.e., a novel *representation* of factual knowledge). This particular task characteristic may have resulted in more remember LIs during create objectives.

Third, apply LIs were significantly more common during apply LOs and, similarly, create LIs were significantly more common during create LOs. One possible explanation (supported by Table 6) is that apply and create LIs are *generally uncommon*. In this respect, it is reasonable that apply and create LIs are *more common* when the ultimate objective is to apply and create, respectively. Apply and create processes are generally not supported by existing search environments. Apply processes include *executing* and *implementing* (e.g., using facts to perform a calculation, using a concept to explain a phenomenon, or using a procedure to perform a task). Apply processes often require a “tangible” application of knowledge that is difficult to carry out within the search environment. Similarly, create processes are also tangible in nature, involving “the construction of an original product” [3, p. 85]. Create processes require tools that support structuring and synthesis in order to generate something new that can be observed (e.g., creating a new table of facts, a concept map, or a step-by-step procedure).

Finally, analyze and evaluate LIs were most frequent during evaluate LOs. It is quite natural for evaluate objectives to involve more evaluate LIs. It is *also* natural for evaluate objectives to involve more analyze LIs. Anderson and Krathwohl support this connection between analyze and evaluate, noting that analysis is often a *prelude* to evaluation [3]. Analysis precedes evaluation because it provides the *evidence* needed for a logical evaluation. For example, in order to evaluate which concept best explains a phenomenon, it is necessary to first analyze the relationships between the different concepts under consideration. This process involves decomposing, comparing, and contrasting (i.e., analyze-level processes).

Effects of the LO’s Knowledge Type: Next, we discuss our results by conditioning on the LO’s knowledge type (Table 7). Our results found five main trends.

First, irrespective of the objective’s knowledge type, remember and understand LIs were the most common. As previously discussed, this result suggests that remembering and understanding are important processes regardless of the end goal.

Second, remember LIs were most frequent when the objective involved factual knowledge. Factual knowledge is made up of bits of information that tend to be concrete (vs. abstract) and objective (vs. subjective). Unlike other types of knowledge, facts are often immediately comprehended and do not benefit from multiple rounds of summarization and exemplification (i.e., understand-level processes). Additionally, facts tend to be self-evident (e.g., “Currently, the world’s tallest building is the Burj Khalifa in Dubai.”). Thus, comprehending a fact may not require multiple rounds of decomposition and comparison (i.e., analyze-level processes). Because of these characteristics, we argue that remember is the most common CP for acquiring factual knowledge. This trend resonates with results reported in Urgo et al. [45]. During factual LOs, participants perceived the task to require *more* memorization and *less* activity across CPs more complex than remember.

Third, understand LIs were most frequent when the objective involved conceptual knowledge. We believe this is due to the inherent nature of conceptual knowledge. Conceptual knowledge is knowledge about concepts, categories, theories, principles, schemas, and models. In this respect, concepts can be highly abstract (e.g., laws of physics) and even subjective (e.g., artistic movements).

The amorphous nature of conceptual knowledge may have required more summarization and exemplification (i.e., understand-level processes) to help delineate its boundaries. Furthermore, in our qualitative coding of pathways, we noticed that a single understand LI was often insufficient for acquiring conceptual knowledge. More often, participants iterated on different understand-level activities (e.g., summarizing a definition, exploring an example, sharpening a definition with new information). This trend also resonates with results reported in Urgo et al. [45]. During conceptual LOs, participants perceived the task to require more understanding.

Fourth, evaluate LIs were the *least* frequent when the objective involved conceptual knowledge. This trend is consistent with the previous trend—conceptual objectives had the most understand LIs. As shown in Table 7, during conceptual objectives, participants spent more LIs iterating on CPs less complex than evaluate (especially understand). In other words, our results suggest that in order to engage in evaluate-level processes during conceptual objectives, participants had to spend many LIs engaging with *lower*-level processes.

Finally, create LIs were most frequent when the objective involved procedural knowledge. This may be explained by procedural LOs allowing for the opportunity to create. As previously discussed, during procedural objectives, participants tended to explore and then modify procedures in order to fit their personal preferences and constraints (e.g., available tools and materials). This trend also resonates with results reported in Urgo et al. [45]. During procedural objectives, participants perceived the task to require more creating.

5.3 RQ3: Transitions Between CPs Traversed

In RQ3, we investigate how participants *transitioned* between CPs in their pathways. Our RQ3 results considered trends from two perspectives. First, we considered common transitions irrespective of the objective. Second, we considered trends specific to the objective’s knowledge type (i.e., factual vs. conceptual vs. procedural). Tables 9–13 summarize the main trends observed in our RQ3 results.



In the following sections, we provide example pathways that demonstrate high-probability transitions. Our goal is to give the reader a sense of why these transitions happened frequently. In each example, we provide: (1) an overview of the LO; (2) a narrative of the LIs shown in the example; and (3) an explanation for why the transition(s) shown in the example may be common during learning-oriented search tasks.

5.3.1 High-probability Transitions Across All LOs. As illustrated in Table 9, our results found four high probability transitions irrespective of the LO: (1) starting with understand; (2) downshifting to understand; (3) upshifting from understand to analyze; and (4) transitioning to the same CP. The following pathway excerpt demonstrates all four common LI transitions.

Learning Objective: The task required the participant to create a diagram of Bernoulli’s principle applied to the notion of lift. (create/conceptual)

- LI₁ (**understand**): Enters query to be able to summarize Bernoulli’s principle.
- ...
- LI₁₇ (**understand**): Reads through example diagram of forces involved in Bernoulli’s principle applied to a wing during lift.
- LI₁₈ (**analyze**): Differentiates those forces from those associated with Newton’s Laws applied to lift.
- LI₁₉ (**understand**): Decides they need more information on Newton’s laws of motion and next looks for examples of newton’s laws of motion and lift.
- ...
- LI₂₁ (**understand**): Verbally summarizes the movement of air above and below a wing in a diagram.
- LI₂₂ (**understand**): Summarizes in notes relevant forces at play in the diagram.
- ...

Table 9. Transitions between CPs That Were Common Across All LOs

Transition Diagram	Explanation
 Starting with Understand	Learner issues general query for overview of fact, concept, or procedure. E.g., “expensive painting”, “automatism”, or “paper airplane.”
 Downshift to Understand	Learner revisits definition or explores new examples to clarify and deepen understanding. Additionally, learner gathers more evidence to increase confidence when choosing between options (i.e., evaluate to understand).
 Back to Self	Learner iterates on the same cognitive process to deepen understanding. E.g., understand to understand occurs while exploring different characteristics of a concept.
 Understand to Analyze	After acquiring basic or deeper understanding, learner can differentiate from related facts, concepts, and procedures.

In the transition diagrams, R = remember, U = understand, P = apply, A = analyze, E = evaluate, and C = create.



First, starting with understand (summarizing or exemplifying) is a logical first step in learning. Participants often began the search task with little prior knowledge, leading them to start by issuing a general query with the intent to understand information (e.g., be able to summarize information in their own words).

Second, downshifts to understand are common because summarizing and exemplifying are basic processes that help people construct more precise and nuanced knowledge, which is needed to support more complex processes. In the above example, the participant downshifts from analyze to understand. This downshift occurs because the participant attempts to differentiate between Newton’s laws and Bernoulli’s principle (LI₁₈). In doing so, the participant realizes that they do not comprehend Newton’s laws well enough to analyze how they relate to Bernoulli’s principle in explaining lift. Downshifting to understand (LI₁₉) allowed the participant to review more examples of Newton’s laws in order to develop a more nuanced understanding of these concepts.

Third, transitions from understand to analyze are common because better understanding a fact, concept, or procedure enables a learner to analyze how it relates to *other* facts, concepts, or procedures. In the example above, LI₁₇ involved clearly defining the boundaries of Bernoulli’s principle. Subsequently, LI₁₇ enabled LI₁₈, which involved analyzing the relationships between Bernoulli’s principle and Newton’s laws in explaining lift.

Finally, transitions to the same CP were generally common. The example above shows the particularly ubiquitous understand-to-understand transition. These transitions were common because gaining a deeper understanding of a specific piece of knowledge often involved *iterating* over different definitions, examples, or perspectives. For example, in the above excerpt, LI₂₁ shows the participant summarizing the movement of different components in a diagram—the movement of air above and below a airplane’s wing during lift. Subsequently, LI₂₂ shows the participant making

Table 10. Transitions between CPs That Were Common to Factual Learning Objectives

Transition Diagram	Explanation
 Starting with Remember	Learner issues precise, well-defined query of specific fact. E.g., “name of most expensive painting” or “world’s tallest building.”
 Downshift to Remember	Learner gathers additional facts. E.g., learner assesses validity of fact and continues to collect <i>other</i> facts.
 Apply to End	Learner combines two facts to generate a third fact. Because facts are well-defined, learners are more likely to be confident in the answer and stop searching.

In the transition diagrams, R = remember, U = understand, P = apply, A = analyze, E = evaluate, and C = create.

notes about the *forces* acting on the wing during lift (i.e., a different *perspective* on the information depicted in the diagram).

5.3.2 High-probability Transitions Unique to Factual LOs. As illustrated in Table 10, our results found three high probability transitions unique to factual LOs: (1) starting with remember; (2) downshifting to remember; and (3) ending with apply. The following pathway excerpt demonstrates these three common LI transitions unique to factual objectives.



Learning Objective: The task required the participant to use the world’s tallest building as a unit to better appreciate the depth of the deepest point in the ocean. (apply/factual)

- LI₁ (**remember**): Enters a query to find specifically the world’s tallest building.
- ...
- LI₅ (**understand**): Summarizes fact in notes about the world’s tallest building.
- LI₆ (**remember**): Enters query to be able to recall the deepest point in the ocean.
- ...
- LI₂₀ (**apply**): Calculates the deepest point of the ocean in terms of the world’s tallest building.

First, remember was a common starting point during factual objectives. Compared to conceptual and procedural knowledge, factual knowledge tends to be concrete, well-defined, and self-contained. Anderson and Krathwohl noted that factual knowledge relates to bits of information that tend to have “a low level of abstraction” and “value in and of themselves” [3, p. 42]. We believe this helps explain why remember was a more common starting point during factual objectives. Participants were often able to gain factual knowledge by simply memorizing or copy/pasting information (versus summarizing or exemplifying). In the above excerpt, in LI₁, the participant starts the task by simply searching for the name of the world’s tallest building.

Second, downshifting to remember was a more common transition during factual objectives. This typically happened when participants sought a new piece of factual knowledge required by the task. In the above excerpt, the participant downshifts from understand (LI₅) to remember (LI₆) because they needed a new piece of factual knowledge required by the task (i.e., the name of the deepest point in the ocean).

Table 11. Transitions between CPs That Were Common to Conceptual Learning Objectives

Transition Diagram	Explanation
 Evaluate to Analyze	Learner gathers more evidence to bolster argument for judgement of concept. The amorphous, broad nature of concepts often make learners less confident when choosing which concept is most suited to a particular scenario (e.g., explaining a natural phenomenon).
 End with Evaluate	Learners engaged in fewer instances of evaluate overall in conceptual learning objectives. When it did occur, evaluate was more likely to be the final step after much understanding and analyzing of concepts throughout the pathway.

In the transition diagrams, R = remember, U = understand, P = apply, A = analyze, E = evaluate, and C = create.

Finally, apply was a more common end point during factual objectives. This common transition is also likely due to factual knowledge being concrete and well-defined. Applying factual knowledge involves less uncertainty than applying conceptual or procedural knowledge. In the above excerpt, in LI₂₀, the participant is able to apply two bits of factual knowledge to generate a new fact. Presumably, the participant felt confident enough in their application of this factual knowledge to end the task with apply. This was less common during conceptual or procedural objectives. When applying conceptual or procedural knowledge, participants often *downshifted* to lower-complexity processes (e.g., revisiting definitions, summaries, and examples) to *verify* whether they applied the knowledge correctly.


5.3.3 High-probability Transitions Unique to Conceptual LOs. As illustrated in Table 11, our results found two high probability transitions unique to conceptual LOs: (1) downshifting from evaluate to analyze and (2) ending with evaluate. The following pathway excerpt demonstrates these two common LI transitions unique to conceptual objectives.

Learning Objective: The task required the participant to determine which art movement most closely aligns with automatism: surrealism or Dadaism. (evaluate/conceptual)

- ...
- LI₃₂ (**evaluate**): Based on characteristics of surrealism and Dadaism, they judge that surrealism seems to be more closely aligned with automatism.
- LI₃₃ (**analyze**): In order to further clarify this distinction, they organize and differentiate surrealism, Dadaism, and automatism in relation to one another.
- ...
- LI₃₅ (**evaluate**): Using evidence that shows “surrealist automatism” as a term, they make a final judgment that surrealism is more closely aligned with automatism than Dadaism.

First, downshifting from evaluate to analyze was common during conceptual objectives. These downshifts mostly occurred when participants felt they needed to reexamine the relations between concepts in order to clarify or strengthen a previously made argument. In the excerpt above, the participant transitions from LI₃₂ (evaluate) to LI₃₃ (analyze) because they needed to revisit the relations between the three concepts associated with the task in order to make a clearer argument.

Table 12. Transitions between CPs That Were Common to Procedural LOs

Transition Diagram	Explanation
 Upshifts to more complex cognitive processes	Learner is able to upshift to analyzing or evaluating without iterating on less complex processes. This may be explained by target procedures being generally well-defined and concrete.
 Back to Create	After creating initial modification, learner is inspired to make additional modifications to procedure (e.g., skip, reorder, modify steps).
 Create to End	Learner creates novel procedure and ends search to test procedure with tangible materials.

In the transition diagrams, R = remember, U = understand, P = apply, A = analyze, E = evaluate, and C = create.

Second, evaluate was a common endpoint during conceptual objectives. This is likely due to the amorphous nature of conceptual knowledge. In contrast to facts and procedures, concepts (e.g., art movements) can be highly abstract and subjective. Based on our coding of pathways, we noticed that participants did not feel confident enough to evaluate conceptual knowledge until going through many iterations of understand (e.g., understanding individual concepts) and analyze (e.g., analyzing the relations between concepts). Our RQ2 results found that evaluate LIs were generally *rare* during conceptual objectives. Our RQ3 results suggest that participants tended to evaluate toward the *end* of the search session (i.e., the final LI in the pathway).

5.3.4 High-probability Transitions Unique to Procedural LOs. As illustrated in Table 12, our results found three types of high probability transitions unique to procedural LOs: (1) upshifting to more complex processes, (2) transitioning from create to create, and (3) ending with create.

The following pathway excerpt demonstrates three common transitions unique to procedural objectives: (1) remember to evaluate, (2) understand to evaluate, and (3) remember to analyze. These are the types of transition *upshifts* that were more common during procedural objectives.

Learning Objective: The task required the participant to find the mathematical center of a firepit circle. (apply/procedural)

- ...
- LI₄ (**remember**): Reads word-for-word through a particular method of finding the center of a circle.
- LI₅ (**evaluate**): Judges that the method looks too difficult.
- LI₆ (**remember**): Reads word-for-word through a different method.
- LI₇ (**understand**): Summarizes a list of necessary materials from the steps they just read.
- LI₈ (**evaluate**): Judges that this method looks even more complicated.
- LI₉ (**remember**): Reads word-for-word through an additional method.
- LI₁₀ (**analyze**): Relates method to previous method, concluding that the methods are similar.
- ...

First, remember to evaluate was a common upshift during procedural objectives. This suggests that participants were often able to make judgements about procedures using superficial evidence or heuristics (e.g., number of steps, number of tools or materials involved, and complexity of steps depicted visually). In the excerpt above, the participant reads the steps of a procedure in LI₄, and quickly judges that the procedure is too complex in LI₅ without trying to fully comprehend it.

Second, understand to evaluate was a common upshift during procedural objectives. Similar to the previous example, in many cases, participants were able to judge a procedure after understanding some aspect of it. In the excerpt above, the participant summarizes the materials involved in a procedure in LI₇, and subsequently judges that the procedure is too complex in LI₈.

Finally, remember to analyze upshifts were common during procedural objectives. This trend is likely due to the concrete nature of procedures, especially those involving physical steps and tangible materials. During procedural objectives, participants were often able to draw comparisons between procedures (an analyze process) without first iterating over multiple rounds of understanding.

The next pathway excerpt demonstrates two common transitions unique to procedural LOs: (1) create to create and (2) ending with create.

Learning Objective: The task required the participant to find different methods for finding the mathematical center of a circle and then develop their own method. (create/procedural)

- ...
- **LI₆ (analyze):** Explores alternatives that can act as a long straightedge in a particular scenario for finding the center of a firepit circle.
- **LI₇ (create):** Designs own alternative tool, explaining that they could use string instead of a straightedge and, additionally, use the string as a way to find the shortest distance between two points.
- **LI₈ (create):** Develops method of finding the absolute diameter, rather than simply a chord, by initially estimating the diameter then moving a string back and forth to make sure that there is not a point where the string is longer.
- ...
- **LI₁₆ (create):** Develops an additional procedure to find the center of the circle using the “reverse” of the circumference formula.

First, create-to-create transitions were more common during procedural objectives. There are two possible reasons for this trend. First, procedures are made up of steps that can be skipped, modified, re-ordered, or combined. Second, procedures often need to be modified or created according to individual preferences or constraints (e.g., prior knowledge, available tools/materials, and temporal constraints). As such, procedural objectives may have had more create-to-create transitions because one modification often leads to additional modifications. In the excerpt above, the participant designs a new tool in LI₇ and subsequently modifies a step to use this new tool in LI₈. Finally, create was a common endpoint during procedural objectives. This trend may simply be due to the prevalence of create LIs during procedural objectives. After modifying a procedure or creating a new one, it was not necessary for participants to downshift to lower-complexity processes (e.g., understand or analyze) because the new procedure was already fully understood. Thus, participants were more often able to end with create.

5.3.5 Low-probability Transitions Unique to a Knowledge Type. Finally, as illustrated in Table 13, there were two transitions that were *uncommon* to a particular knowledge type. First, remember-to-understand transitions were uncommon during factual objectives. Most remember transitions simply iterated back to remember. As noted above, this is likely due to factual knowledge being concrete, well-defined, and self-contained. We believe these attributes of factual knowledge rendered

Table 13. Transitions between CPs That Were Uncommon

Transition Diagram	Explanation
 Remember to Understand (<i>Factual</i>)	Facts often did not require understand (e.g., summarization, exemplification) for learner to internalize
 Analyze to Evaluate (<i>Conceptual</i>)	Learners had more difficulty in increasing in complexity when learning concepts

In the transition diagrams, R = remember, U = understand, P = apply, A = analyze, E = evaluate, and C = create.

more complex CPs (e.g., summarizing, exemplifying) unnecessary for comprehension. Instead, participants tended to iterate from remember to remember LIs as they gathered multiple distinct facts in quick succession.

Second, analyze-to-evaluate upshifts were uncommon during conceptual objectives. Most analyze transitions iterated back to analyze or downshifted to understand. This trend is likely due to conceptual knowledge being abstract, subjective, and interconnected. Evaluating conceptual knowledge required more iterations of understand (i.e., understanding individual concepts) and analyze (i.e., analyzing the relations between concepts). Participants often iterated over these lower-complexity processes before being able to confidently evaluate conceptual knowledge. This trend resonates with results from Urgo et al. [45], which found that participants perceived conceptual objectives to be more difficult than factual and procedural objectives. Participants’ inability to move up in complexity from analyze to evaluate may have contributed to their perceptions of task difficulty.

6 IMPLICATIONS FOR DESIGNING TOOLS TO SUPPORT LEARNING

Our results found that the CP, and to a greater extent, the knowledge type of the LO impacted: (RQ1) the length and diversity of pathways; (RQ2) the types of CPs traversed along the pathways; and (RQ3) the types of transitions between CPs. Our results have implications for designing search environments that support learning.

In this section, we propose different tools and features to encourage and support learning during search. Table 14 provides a summary of our proposed tools. In Table 14, tools are grouped by type of LO. First, we discuss tools to support conceptual objectives. Then, we discuss tools to support procedural objectives. Finally, we discuss tools to support *all* objectives. Tools to support factual objectives (e.g., note-taking tools that enable searchers to save information) are also likely to support conceptual and procedural objectives. Therefore, in Table 14, we did not create a separate section for tools to uniquely support factual objectives. Instead, tools to support factual objectives are discussed under “All LOs”. In Table 14, we describe each proposed tool, list which of our results motivate the proposed tool, and acknowledge prior work that has investigated similar tools.

6.1 Tools to Support Conceptual LOs

First, in Table 14, we discuss three potential tools and features to support searchers with conceptual LOs.

Our results found that conceptual LOs had the most understand LIs (RQ2) and frequent downshifts to understand (RQ3). As previously noted, during conceptual objectives, participants often felt unsure if they understood concepts well enough to support more complex subgoals (e.g.,

Table 14. Implications for Search System Modifications Based on Findings (Section 4)

Learning Objective Type	Potential Modification/Tool	Results Supported	Related Work
<i>Conceptual Learning Objectives</i>	Organize results by definitions, overviews, and examples	• Conceptual LOs involved more understand (RQ2)	
	Organize examples by type	• Conceptual LOs involve downshifts to understand (RQ3)	
	Questions posed to test understanding; e.g., “Which of these images exemplify Bernoulli’s principle?”	• Conceptual LOs involve downshifts to understand (RQ3)	• Question generation [42]
<i>Procedural Learning Objectives</i>	Linking procedures through “querying-by-example” and related procedures side-by-side	• Procedural LOs involved more creating (RQ1, RQ2) • Procedural LOs involve upshifts (RQ3)	
	Substitute or functionally equivalent materials	• Procedural LOs involved more creating (RQ1, RQ2)	
	Relevant conceptual knowledge to inspire the creation of novel procedures	• Procedural LOs involved more creating (RQ1, RQ2)	
<i>All Learning Objectives</i>	Interactive space to copy, summarize, organize, connect, and modify evolving knowledge	• Downshifts to understand common in all LOs (RQ3) • Factual LOs involved more remember (RQ2) • Procedural LOs involved more creating (RQ1, RQ2)	• Sticky-note tool [15] • Note-taking [37] • Note-taking and highlighting [39]
	Trails of searcher’s own search history	• Downshifts to understand common in all LOs (RQ3)	• Search trails of other searchers [7, 47]
	Filter search results across dimensions such as complexity, perspective, and originating source	• Transitions back to the same CP common in all LOs (RQ3)	• Perspective-based search system [43]
	Related facts, concepts, and procedures; highlight differences/similarities	• Understand to analyze common in all LOs (RQ3) • Factual LOs involve downshifting to remember (RQ3) • Conceptual LOs involve downshifting from evaluate to analyze (RQ3)	

analyze the relations between concepts or evaluate the relevance of concepts in a given scenario). Participants often iterated over multiple understand LIs (i.e., understand-to-understand transitions) by reviewing definitions, summaries, and examples from different perspectives. We propose three tools and features to support such understand-level processes during conceptual objectives.

First, given a conceptual knowledge query (e.g., “cubism”), systems could diversify and organize search results by definitions, overviews, and examples. Commercial search systems already do this to some extent (albeit inconsistently). For example, in response to the query “cubism”, Google displays dictionary definitions, encyclopedic articles, and images of cubist paintings. However, systems could organize these different types of results more clearly with explicitly marked sections on the SERP.

Second, to deepen their understanding, participants often explored examples of a concept from different perspectives. To support these activities, systems could present and organize examples of a concept by *type* within the SERP. To illustrate, Bernoulli’s principle (B) can be demonstrated in multiple ways (e.g., B and sail boats, B and wings, B and Venturi tubes, B and chimneys, B and shower curtains, B and curveballs in baseball, and B and topspin in tennis). Grouping examples by *type* may encourage searchers to explore examples from *different* perspectives and learn about *common* themes, deepening their understanding of a concept.

Third, and perhaps more ambitiously, systems could enable searchers to test their own understanding of a concept. In prior search-as-learning studies, understand-level learning has often been measured using closed-ended tests (e.g., fill-in-the-blank, true/false, and multiple-choice tests). Such tests could potentially be automated by a search system. For instance, a system could show a searcher a set of images and ask them to select which ones *exemplify* a given concept (e.g., “Which of these images exemplify Bernoulli’s principle?”). Answering such automatically generated questions may help searchers identify knowledge gaps and assess whether they understand a concept well enough to pursue more complex subgoals (e.g., analyze or evaluate). Prior research has investigated systems that enable searchers to assess their own knowledge using automatically generated questions [18, 19, 42]. For example, Syed et al. [42] evaluated a system that prompted participants to answer automatically generated factoid questions to assess their comprehension of passages read during the session (i.e., remember/factual learning). Future research could extend such systems by asking more complex questions to support a wider range of LOs during search.

6.2 Tools to Support Procedural LOs

Next, in Table 14, we discuss four potential tools and features to support searchers during procedural LOs.

Our results found that procedural objectives involved more create LIs (RQ2) and create-to-create transitions (RQ3). As previously noted, during procedural objectives, participants often simplified, modified, combined, and created new procedures based on their unique preferences (e.g., familiar techniques) and constraints (e.g., available tools and materials). We propose three tools to support such create-level processes during procedural objectives.

First, systems should prioritize *linking* procedures through “querying-by-example”, enabling searchers to efficiently find alternative procedures for the same task. Additionally, systems could display related procedures side-by-side and emphasize their similarities and differences. Helping searchers find and compare related procedures may enable them to: (1) gauge the range of alternative approaches to the task at hand; (2) identify steps that are skippable or modifiable; and (3) discover different ways to execute a step. Ultimately, such tools could help searchers *combine* ideas from related procedures to fit their preferences and constraints.

Second, during procedural objectives, create LIs often occurred when participants considered alternative techniques, tools, and materials to implement a specific procedure. To support this

process, search systems could automatically identify and suggest alternative techniques, tools, and materials when a searcher is reviewing a specific procedure. For example, if a recipe calls for heavy cream, the system could automatically point out that half-and-half is a common replacement.⁵

Third, participants sometimes used concepts (and their definitions) as inspiration to create new procedures. For example, one participant used the concept of a diameter to develop their own approach for finding the center of a circle.⁶ To better support this process, search systems could display related concepts in response to a procedural knowledge query.

Our results also found that procedural objectives had more upshift transitions to more complex processes (RQ3). To support these upshifts, systems could enable searchers to “query-by-example” based on explicitly stated goals. In other words, the system could enable searchers to submit a specific procedure as a “query” and explicitly request either: (1) background information (i.e., to support upshifts to understand); (2) example videos of the procedure being executed (i.e., to support upshifts to apply); (3) alternative procedures with the same objectives (i.e., to support upshifts to analyze); (4) pros and cons of the procedure (i.e., to support upshifts to evaluate); or (3) potential modifications (i.e., to support upshifts to create).

6.3 Tools to Support All LOs

Finally, in Table 14, we discuss four potential tools and features to support searchers with LOs of any knowledge type.

First, our results found that remember and understand were frequent LIs regardless of the objective (RQ2). In other words, as might be expected, remember- and understand-level processes seem to be foundational in support of more complex subgoals. To support searchers with remember- and understand-level processes, systems should provide interactive spaces for searchers to copy/paste, summarize, organize, and annotate information as they search. Prior studies have found that such tools can provide learning benefits [15, 37, 39]. As an implication for future research, we believe that such tools should be designed to support and encourage both remember- *and* understand-level processes. Remember-level processes can be easily supported by allowing searchers to copy/paste and save information. More importantly, to support understand-level processes, tools should encourage searchers to summarize information in their own words and organize information using their own knowledge representations (e.g., labeled clusters).

Second, downshifts to understand were common regardless of the LO (RQ3). To support these downshifts, systems could provide users with search trails of their own search history. In our study, participants often downshifted to understand by revisiting pages to reread content or reissuing queries to examine alternative search results. Providing users with access to their own search trails could streamline this process. Prior studies have investigated the benefits of search trails from *other* searchers who performed a similar task [7, 47]. Our results suggest that users may also benefit from having access to their *own* search trails during a learning-oriented session.

Third, transitions back to the same CP were frequent regardless of the LO (RQ3). For example, if a searcher is evaluating, it is likely that they will continue to evaluate in the next subgoal. We noticed this happening when participants wanted: (1) simpler or more complex content to better *integrate* new information with their existing knowledge; (2) content from a different source to *verify* newly acquired knowledge; or (3) content from a different perspective to *deepen* their knowledge. This

⁵Procedural knowledge sites such as Wikihow often include user-generated comments that mention functionally equivalent techniques, tools, and materials. A system could mine such resources to identify interchangeable techniques, tools, and materials.

⁶Use a string to find the longest “chord” of the circle (i.e., the diameter), and then fold the string in half to find the center of the circle.

trend is consistent with results from Liu et al. [31], which found that learning-oriented searches often involve repeated iterations of the same overall intent. To support these iterative transitions, systems should enable users to filter search results along dimensions such as complexity (or target audience), originating source, and perspective. Developing systems that can filter results by complexity and source seem relatively straightforward. Developing systems that can filter results by perspective seem more challenging and are area of ongoing research. Tabrizi and Shakery [43] describe the different hypothetical components of a perspective-based search system.

Finally, upshifts from understand to analyze were frequent regardless of the LO (RQ3). This means that understanding a fact, concept, or procedure, typically leads to analyzing its relation (e.g., similarities and differences) to other facts, concepts, and procedures. In other words, understanding something in isolation typically leads to understanding it in a greater framework. To support these understand-to-analyze transitions, systems should enable searchers to explore related facts, concepts, and procedures in response to an understand-level query. Commercial systems already do this to some extent (albeit inconsistently). For example, in response to the query “burj khalifa skyscraper height”, Google also shows results for other skyscrapers in the “People also search for” section. By making this related information more accessible, searchers may be able to situate newly acquired knowledge in a greater framework. For example, the Burj Khalifa is 2,717 feet tall and (more impressively) more than twice as tall as the Empire State Building.

As an implication for future work, systems should enable searchers to explore related facts, concepts, and procedures in a more *consistent* and *self-directed* manner (e.g., “See related [facts | concepts | procedures]”). Enabling searchers to explore related facts, concepts, and procedures may encourage searchers to transition from understand- to analyze-level processes and may provide different benefits. As in the example above, exploring related facts may help searchers gain a deeper appreciation of a fact. Exploring related concepts may help searchers understand the unique aspects of a concept in relation to similar concepts. Exploring related procedures (i.e., procedures to accomplish a different but related task) may help searchers deepen their understanding of important concepts and techniques used in related tasks. Ultimately, such tools could help searchers situate newly acquired knowledge in a greater framework (i.e., an analyze-level process).

7 CAVEATS AND OPPORTUNITIES FOR FUTURE WORK

Our study and results have a few caveats worth noting.

First, participants completed tasks with LOs associated with three (out of six) CPs from A&K’s taxonomy. As previously mentioned, we omitted the CPs of remember and understand because they are the least complex, and we omitted analyze because analyzing is a necessary component of evaluating. Future work might consider the full range of CPs as LOs.

Second, participants were given a maximum of 15 minutes to complete the search phase of each task. Naturally, time constraints can influence search behaviors and outcomes [12]. In our case, the 15-minute time limit may have influenced the pathways taken by participants to achieve the given objective. Future work is needed to investigate how time constraints may influence the pathways taken by searchers toward an objective. In our study, the 15-minute time limit was imposed to keep the study session under 1.5 hours. Additionally, after several rounds of pilot testing, we determined that 15 minutes was enough time for participants to complete our tasks. Ultimately, participants spent about 10 minutes searching on average ($M = 9.79$, $S.D. = 5.39$).⁷

⁷In Urgo et al. [45], we reported on the effects of the objective on task completion time. The cognitive process of the objective had no significant effects. Conversely, participants took longer to complete tasks with conceptual objectives (11.65 minutes) versus procedural objectives (8.43 minutes).

Finally, in our study, we did not explicitly consider knowledge gain as a main dependent variable. In other words, participants did not complete pre- and post-tests to measure learning during the search process. Instead, participants were asked to demonstrate their achievement of the LO in a 2-minute video recorded by the moderator. We believe that the video demonstration phase of each task encouraged participants to achieve the given objective. Participants had to explain their main solution to the task “live” and in front of the moderator.

In terms of learning outcomes, we see important opportunities for future work. In our study, we used qualitative techniques to investigate the learning *process* during search. Future studies could adopt similar techniques to investigate how the learning process might influence learning *outcomes*.⁸ An important open question is: What are characteristics of pathways that lead to greater knowledge gains? For example, are greater knowledge gains achieved when pathways are longer, more diverse, or have more upshift transitions? Our study focused on how objectives influence pathways. Future work should also consider how pathways impact learning.

8 CONCLUSION

In this article, we have introduced the notion of a pathway—a sequence of learning instances (or subgoals) followed by a searcher toward a specific LO. We leveraged A&K’s taxonomy [3] to characterize both LOs and pathways. We studied the impact of a objective’s knowledge type and CP on the pathway length, diversity, and CPs traversed. Additionally, we analyzed the transition probabilities between CPs conditioned on the knowledge type of the objective.

Our research makes several important contributions. First, from a methodological perspective, we present a method for analyzing search sessions to gain insights about the learning *process* during search. In this article, we analyzed how the LO (i.e., the end goal) can influence the learning process. Future studies could use our methodology to more closely investigate how pathway characteristics influence learning outcomes.

Second, our results found that the objective can influence the pathways followed by searchers. Importantly, the knowledge type of the objective had a much stronger effect than the CP. For example, factual objectives involved more *remembering* (e.g., memorizing), conceptual objectives involved more *understanding* (e.g., summarizing and exemplifying), and procedural objectives involved more *creating* (e.g., simplifying, modifying, and combining). Prior studies have also leveraged A&K’s taxonomy to investigate how objectives impact search behaviors and outcomes. However, studies have primarily leveraged the CP dimension and ignored the knowledge type dimension. Our results indicate that searchers may need different types of support during objectives involving factual, conceptual, and procedural knowledge.

Finally, our analysis of CP transitions revealed several important trends. Our results found several likely transitions irrespective of the objective. Participants were likely to start with understand, iterate on the same CP, transition from high- to low-complexity processes, and transition from understand to analyze. Additionally, our results found several transitions to be likely *depending* on the knowledge type of the objective. For example, during factual objectives, participants were more likely to start with remember and transition to remember from more complex processes. During conceptual objectives, participants were more likely to start with understand and transition *back* to understand from more complex processes. During procedural objectives, participants were the most likely to transition *upwards* from low- to high-complexity processes. We have discussed implications from our results for designing search environments (i.e., search features and auxiliary tools) to support learning. In future work, we plan to explore the impact of *scaffolding* tools that support the development and externalization of subgoals based on the pathways observed in our study.

⁸In Section 2, we briefly summarize different methods used to quantify knowledge gains during the search process.

A APPENDIX

A.1 Coding Guide

The coding guide was developed based on A&K’s book [3], situating LOs at the intersection of two orthogonal dimensions: the CP dimension and the knowledge type dimension. This coding guide associates LIs with a particular CP *and* a particular knowledge type. Because metacognitive knowledge was not part of our analysis, the guide includes three knowledge types (factual, conceptual, and procedural) and six CPs (remember, understand, apply, analyze, evaluate, and create). Thus, there are 18 possible coding combinations. Each of these combinations (e.g., Factual/Remember) are outlined below.

The LI units consisted of recorded screen activities (e.g., queries) and think-aloud comments. When considering how to categorize an LI, the coder asked themselves “What is the intention of the learner at this moment?” Then, the coder used the following criteria to determine the most appropriate category. Criteria for these categories include particular actions and associated examples.

Factual

Factual/Remember

- Reads a fact (e.g., person, date, and place)
- Copy/pastes a fact without additional verbal or written elaboration
- Memorizes a fact
- Issues a query using specific language of fact
 - e.g., “tallest building”
- Notes having seen or heard of fact before

Factual/Understand

- Finds example of fact in a visualization
- Restates fact in own words
- Makes an estimation of a fact
- Writes down a fact with additional verbal elaboration
- Acknowledges or includes source of fact
- Makes a query using general language of fact
 - e.g., “tall buildings”
- Investigates a particular fact beyond just noting/reading the fact
 - e.g., “I don’t know Abraham Lincoln. Let’s look into him a bit more.”
 - e.g., “This artist keeps coming up, so I’ll look into them a bit more.”
- Makes some inference about fact
 - e.g., “This mentions that Americans have an average of 2.5 doctors for every 1,000 people while Europeans have an average of 3.5, we Americans must generally be less healthy.”

Factual/Apply

- Uses one fact with second fact to find third fact

Factual/Analyze

- Compares and differentiates facts that clarify or explain fact
- Compares two facts to arrive at most reasonable average or estimation

- Differentiates between categories of facts
 - e.g., “Oh I see this is the number of years a cat lives, but this is the number of human years a cat lives.”
- Compares facts to find correlations or themes among the facts or generalizations
 - e.g., “Most of the deep parts of the ocean are in the Pacific Ocean.”
- Lists or identifies correlations (or lack of correlation) or themes among the facts
- Identifies an outlier among correlations or themes
 - e.g., “Most of these are in the Pacific Ocean, but I see this point is in the Atlantic.”
- Disambiguates two facts
 - e.g., “The Challenger Deep exists within the Mariana Trench, they are not separate locations.”

Factual/Evaluate

The existence of a “judgment” word or phrase does not in and of itself elicit an evaluate code (e.g., “that’s interesting”, “so that’s good”, or “that’s not what we need” are not coded as evaluate). It is important that the judgment is made with respect to the learning subgoal or overall objective (e.g., “I can see that most of these paintings are from post-1900, so date of creation is a good potential explanation for the high cost of the painting.”)

- Judges, critiques, or questions validity of a fact
- Judges or critiques importance or validity of a fact that explains other fact
- Judges, critiques, or questions resource validity of fact
- Verifies a previously known fact.
- Making a judgment about a fact being good/bad at explaining something.
- States why a fact is good/bad at explaining something.
- Judges usefulness of criteria for assessing reasoning behind fact
 - e.g., “This cat breed chart’s variables are a good way to assess why cats are such popular pets.”
- Selects a reason for explanation of fact
 - e.g., “I think this cat is so expensive because the breed is so rare.”
- Compares lists of facts to find agreement
 - e.g., “I’m actually not finding a lot of overlap between the two lists of most expensive paintings.”

Factual/Create

- Generates new logical reason substantiating a fact that is not explicitly or implicitly cited in resource

Conceptual

Conceptual/Remember

- Reads through definition of concept
- Memorizes definition of concept
- Writes down or copy/pastes definition of concept without additional verbal or written elaboration
- Reads through (notes) techniques or concepts associated with main concept without additional verbal or written elaboration

- Reads through (notes) definitions of techniques or concepts associated with main concept without additional verbal or written elaboration
- Reads or notes (not queries) name of concept without additional verbal or written elaboration
- Recalls definition of concept

Note: If examples are involved in *any* way then move up (at least) to understand

Conceptual/Understand

- Reviews example of concept (e.g., reads the labeled parts of an example of concept, looks at paintings that exemplify an artistic style)
- Summarizes definition of concept in own words
- Writes down concept and elaborates verbally
 - e.g., notes concept and states that they do not understand the concept
 - e.g., notes concept and states they will look into it later
- Reviews examples of techniques or concepts associated with main concept
- Summarizes definition of techniques or concepts associated with main concept
- Reviews example or representative work of concept
- Reviews example or representative work of technique associated with main concept
- States that they are attempting to “understand” a technique or concept
- Makes a guess at the definition of a concept in own words
 - e.g., “So it seems like surrealism might be art that is made from an artist’s perception of dreams or other alternate realities.”
- Makes some inference about concept
 - e.g., “It says that insurance can be included in a mortgage, so it’s probably not mandatory.”
- Issues exploratory query of concept
 - e.g., “I really don’t know what this is, but I’ll start with this.’ Bernoulli’s principle”
- Reads, writes down, or notes example of concept
 - e.g., “So water through a hose with a thumb partially covering the hole is an example of Bernoulli’s principle.”
- Recognizes explanations of phenomena that use different concepts
 - e.g., “So there are two different explanations for why the sky is blue, Tyndall Effect and Rayleigh Scattering.”
- Uncovers conflicting or disputed information
 - e.g., “So it looks like there is disagreement on which is the appropriate theory to explain why the sky is blue.”

Conceptual/Apply

- Talks through application of concept (e.g., expresses how essential components of concept function or are represented in a given example)
- Questions whether example is instance of concept by citing definition or characteristics of concept
- Applies essential components of concept to specific novel example
- Revisits example to apply participant’s newly acquired characteristics of concept to example
 - e.g., “Ah yes, I see now that this is made up of corners consisting of perpendicular lines as I learned from the definition of a square.”
- Investigates a particular application of a concept to a scenario
 - e.g., “I understand the general idea of Newton’s second law, but I am really interested in how this applies to a kayak moving through the water.”
- Applies concept to example from memory (example is not explicitly stated in site)

- e.g., “oh so the Tyndall Effect is like when you see blue smoke from motorcycle exhaust.” (when motorcycle example had not ever been mentioned during search session)

Conceptual/Analyze

- Differentiates between two or more concepts
- Connects the dots among examples, extrapolates, generalizes concept
 - e.g., “So it looks like all of these are abstract and involve lions.”
 - e.g., “The image search is not really bringing up what I expected...it’s almost like it’s [concept is] broader than I assumed.”
- Differentiates between forms, representations, or types of concept
 - e.g., “If this needs to include all systems of democracy I should include parliamentary and presidential.”
- Groups examples by form, representation or type of concept
- Realizes that two different concepts together best explain a phenomena
 - e.g., “Ok, so I guess you need to know about the continuity equation and Bernoulli’s equation to analyze the flow of a fluid.”

Conceptual/Evaluate

The existence of a “judgment” word or phrase does not in and of itself elicit an evaluate code (e.g., “that’s interesting”, “so that’s good”, or “that’s not what we need” are not coded as evaluate). It is important that the judgment is made with respect to the learning subgoal or overall objective (e.g., “I think Bernoulli’s principle is easier to understand than Newton’s laws of motion when explaining lift.”)

- Judges or critiques effectiveness, appropriateness, and/or validity of concept to explain a phenomenon
- Judges or critiques a concept given some characteristics (e.g., efficiency, effectiveness, context, complexity, ease, and appropriateness)
 - e.g., “So that makes sense and it’s a little bit easier than the general definition of Bernoulli’s principle.”
- Judges or critiques usage of concept
- Judges, critiques, or questions resource validity of concept
 - e.g., “In order to be super thorough I am going to visit another link just to make sure the information about the components of a mortgage are consistent across sources.”

Conceptual/Create

- Generates new representation of concept
- Draws connections among concepts beyond those explicitly or implicitly stated in resource

Procedural

Procedural/Remember

- Reads or notes steps of procedure
- Rotely memorizes steps of procedure
- Copy/pastes steps of procedure

Procedural/Understand

- Summarizes steps of procedure
- Reviews example of procedure with verbal or written elaboration

- Reviews diagram of procedure’s end product with verbal or written elaboration
- Issues exploratory query of procedure
 - e.g., ways of paper airplane making
- Notes some characteristic of procedure
 - e.g., “So that said it would be x number of steps but it is really y number of steps.”

Procedural/Apply

- Executes/narrates/visualizes a procedure

Procedural/Analyze

- Explores options (e.g., materials, steps, and tools) that meet criteria for procedure
- Compares and contrasts two or more procedures
 - e.g., “This one has a lot of steps, this one has fewer steps. It seems like there is a wide range of steps that it might take to fold an airplane.”
- Selects essential components of procedure
- Explores tradeoffs of a particular method or methods
- Explores suggestions or techniques for improving execution of procedure
- Compares own version of steps of procedure with procedure on site
 - e.g., “I’m just going to make sure my steps that I’ve written down match with what I see.”
- Identifies important techniques in executing procedure that make the end product most effective
 - e.g., “I think it is important to hold the string very taught to make sure it is a straight line, and it is also important to make sure both lengths of string are measured exactly or it won’t be accurate.”
- Assessing which of the steps need to be executed for personal subjective context
 - e.g., “I already have a starter for my bread dough, so I won’t need to do the first three steps; I can start with step four.”

Procedural/Evaluate

The existence of a “judgment” word or phrase does not in and of itself elicit an evaluate code (e.g., “that’s interesting”, “so that’s good”, or “that’s not what we need” are not coded as evaluate). It is important that the judgment is made with respect to the learning subgoal or overall objective (e.g., “I think the Kite would be a better choice for my nephew because it can fly farther and do more tricks than the Hammer.”)

- Judges or critiques a procedure given some characteristics (e.g., efficiency, effectiveness, context, complexity, expertise of executor, ease, enjoyment, attractiveness, and appropriateness)
- Checks that procedure was executed correctly
- Outlines criteria used to judge methods
 - e.g., “So I am looking for methods that are challenging, interesting, and have many steps.”
- Makes judgement that a procedure is difficult or too difficult/not worth the effort
- Makes judgement that a procedure is interesting or fun
- Judges steps of procedure to be (in)correct or expressed (in)correctly
 - e.g., “It doesn’t make sense to do these steps in this order, I wouldn’t arrive at the end product.”
- Judges, critiques, or questions resource validity of procedure

Procedural/Create

- Finds useful/appropriate replacement materials for a procedure
- Modifies step(s) of a procedure in order to suit a particular application
- Combines steps of procedures into new procedure
- Generates a novel procedure

REFERENCES

- [1] Mustafa Abualsaud. 2017. *Learning Factors and Determining Document-level Satisfaction In Search-as-Learning*. Master's Thesis. University of Waterloo, Waterloo, Ontario, Canada.
- [2] James Allan, Bruce Croft, Alistair Moffat, and Mark Sanderson. 2012. Frontiers, challenges, and opportunities for information retrieval: Report from SWIRL 2012 the second strategic workshop on information retrieval in Lorne. *Association for Computing Machinery SIGIR Forum* 46, 1 (May 2012), 2. DOI : <https://doi.org/10.1145/2215676.2215678>
- [3] Lorin W. Anderson, David R. Krathwohl, Peter W. Airasian, Kathleen A. Cruikshank, Richard E Mayer, Paul R. Pintrich, James Raths, and Merlin C. Wittrock. 2001. *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives, Complete Edition*. Addison Wesley Longman.
- [4] Nilavra Bhattacharya and Jacek Gwizdka. 2019. Measuring learning during search: Differences in interactions, eye-gaze, and semantic similarity to expert knowledge. In *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval*. Association for Computing Machinery, Glasgow Scotland UK, 63–71. DOI : <https://doi.org/10.1145/3295750.3298926>
- [5] Kathy Brennan, Diane Kelly, and Jaime Arguello. 2014. The effect of cognitive abilities on information search for tasks of varying levels of complexity. In *Proceedings of the 5th Information Interaction in Context Symposium*. Association for Computing Machinery, New York, NY, 165–174. DOI : <https://doi.org/10.1145/2637002.2637022>
- [6] Arthur Câmara, Nirmal Roy, David Maxwell, and Claudia Hauff. 2021. Searching to learn with instructional scaffolding. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*. Association for Computing Machinery, New York, NY, 209–218. DOI : <https://doi.org/10.1145/3406522.3446012>
- [7] Robert Capra, Jaime Arguello, Anita Crescenzi, and Emily Vardell. 2015. Differences in the use of search assistance for tasks of varying complexity. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. Association for Computing Machinery, New York, NY, 23–32. DOI : <https://doi.org/10.1145/2766462.2767741>
- [8] Yu Chi, Shuguang Han, Daqing He, and Rui Meng. 2016 Exploring knowledge learning in collaborative information seeking process. *SAL @ SIGIR*.
- [9] Kevyn Collins-Thompson, Preben Hansen, and Claudia Hauff. 2017. Search as learning (dagstuhl seminar 17092). *Dagstuhl Reports* 7, 2 (2017), 135–162. DOI : <https://doi.org/10.4230/DagRep.7.2.135>
- [10] Kevyn Collins-Thompson, Soo Young Rieh, Carl C. Haynes, and Rohail Syed. 2016. Assessing learning outcomes in web search: A comparison of tasks and query strategies. In *Proceedings of the 2016 ACM on Conference on Human Information Interaction and Retrieval*. Association for Computing Machinery, New York, NY, 163–172. DOI : <https://doi.org/10.1145/2854946.2854972>
- [11] Lynne Cooke. 2010. Assessing concurrent think-aloud protocol as a usability test method: A technical communication approach. *IEEE Transactions on Professional Communication* 53, 3 (Sept. 2010), 202–215. DOI : <https://doi.org/10.1109/TPC.2010.2052859>
- [12] Anita Crescenzi, Rob Capra, Bogeum Choi, and Yuan Li. 2021. Adaptation in information search and decision-making under time constraints. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*. Association for Computing Machinery, New York, NY, 95–105. DOI : <https://doi.org/10.1145/3406522.3446030>
- [13] Sarah Davies, Kirsten Butcher, and Corey Stevens. 2013. Self-regulated learning with graphical overviews: When spatial information detracts from learning. *Proceedings of the Annual Meeting of the Cognitive Science Society* 35, 35 (2013), 2136–2141. Retrieved from <https://escholarship.org/uc/item/5kp8b0n9>.
- [14] Diego Demaree, Halszka Jarodzka, Saskia Brand-Gruwel, and Yvonne Kammerer. 2020. The influence of device type on querying behavior and learning outcomes in a searching as learning task with a laptop or smartphone. In *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval*. Association for Computing Machinery, New York, NY, 373–377. DOI : <https://doi.org/10.1145/3343413.3378000>
- [15] Luanne Freund, Rick Kopak, and Heather O'Brien. 2016. The effects of textual environment on reading comprehension: Implications for searching as learning. *Journal of Information Science* 42, 1 (Feb. 2016), 79–93. DOI : <https://doi.org/10.1177/0165551515614472>
- [16] Ujwal Gadiraju, Ran Yu, Stefan Dietze, and Peter Holtz. 2018. Analyzing knowledge gain of users in informational search sessions on the web. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval*. Association for Computing Machinery, New York, NY, 2–11. DOI : <https://doi.org/10.1145/3176349.3176381>

- [17] Souvick Ghosh, Manasa Rath, and Chirag Shah. 2018. Searching as learning: Exploring search behavior and learning outcomes in learning-related tasks. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval*. Association for Computing Machinery, New York, NY, 22–31. DOI : <https://doi.org/10.1145/3176349.3176386>
- [18] Michael Heilman, Kevyn Collins-Thompson, Jamie Callan, Maxine Eskenazi, Alan Juffs, and Lois Wilson. 2010. Personalization of reading passages improves vocabulary acquisition. *International Journal of Artificial Intelligence in Education* 20, 1 (Jan. 2010), 73–98. DOI : <https://doi.org/10.3233/JAI-2010-0003> Publisher: IOS Press.
- [19] Michael Heilman and Maxine Eskenazi. 2006. Language learning: Challenges for intelligent tutoring systems. In *Proceedings of the Workshop on Intelligent Tutoring Systems for Ill-Defined Domains*.
- [20] William R. Hersh, Diane L. Elliot, David H. Hickam, Stephanie L. Wolf, and Anna Molnar. 1995. Towards new measures of information retrieval evaluation. In *Proceedings of the 18th Annual International Association for Computing Machinery SIGIR Conference on Research and Development in Information Retrieval - SIGIR’95*. Association for Computing Machinery Press, Seattle, Washington, 164–170. DOI : <https://doi.org/10.1145/215206.215355>
- [21] Xiao Hu and Noriko Kando. 2017. Task complexity and difficulty in music information retrieval. *Journal of the Association for Information Science and Technology* 68, 7 (July 2017), 1711–1723. DOI : <https://doi.org/10.1002/asi.23803>
- [22] Rishita Kalyani and Ujwal Gadiraju. 2019. Understanding user search behavior across varying cognitive levels. In *Proceedings of the 30th ACM Conference on Hypertext and Social Media*. Association for Computing Machinery, New York, NY, 123–132. DOI : <https://doi.org/10.1145/3342220.3343643>
- [23] Yvonne Kammerer, Rowan Nairn, Peter Pirolli, and Ed H. Chi. 2009. Signpost from the masses: learning effects in an exploratory social tag search browser. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, Boston, MA, 625–634. DOI : <https://doi.org/10.1145/1518701.1518797>
- [24] Diane Kelly, Jaime Arguello, Ashlee Edwards, and Wan-ching Wu. 2015. Development and evaluation of search tasks for IIR experiments using a cognitive complexity framework. In *Proceedings of the 2015 International Conference on The Theory of Information Retrieval*. Association for Computing Machinery, New York, NY, 101–110. DOI : <https://doi.org/10.1145/2808194.2809465>
- [25] J. Richard Landis and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics* 33, 1 (1977), 159–174. DOI : <https://doi.org/10.2307/2529310>
- [26] Pei-Lan Lei, Chuen-Tsai Sun, Sunny S. J. Lin, and Tsung-Kuan Huang. 2015. Effect of metacognitive strategies and verbal-imagery cognitive style on biology-based video search and learning performance. *Computers & Education* 87 (Sept. 2015), 326–339. DOI : <https://doi.org/10.1016/j.compedu.2015.07.004>
- [27] Yuelin Li and Nicholas J. Belkin. 2008. A faceted approach to conceptualizing tasks in information seeking. *Information Processing & Management* 44, 6 (Nov. 2008), 1822–1837. DOI : <https://doi.org/10.1016/j.ipm.2008.07.005>
- [28] Chang Liu and Xiaoxuan Song. 2018. How do information source selection strategies influence users’ learning outcomes’. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval*. Association for Computing Machinery, New York, NY, 257–260. DOI : <https://doi.org/10.1145/3176349.3176876>
- [29] Hanrui Liu, Chang Liu, and Nicholas J. Belkin. 2019. Investigation of users’ knowledge change process in learning-related search tasks. *Proceedings of the Association for Information Science and Technology* 56, 1 (2019), 166–175. DOI : <https://doi.org/10.1002/pra2.63> eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/pra2.63>.
- [30] Jingjing Liu, Nicholas J. Belkin, Xiangmin Zhang, and Xiaojun Yuan. 2013. Examining users’ knowledge change in the task completion process. *Information Processing & Management* 49, 5 (Sept. 2013), 1058–1074. DOI : <https://doi.org/10.1016/j.ipm.2012.08.006>
- [31] Jiqun Liu, Shawon Sarkar, and Chirag Shah. 2020. Identifying and predicting the states of complex search tasks. In *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval*. Association for Computing Machinery, New York, NY, 193–202. DOI : <https://doi.org/10.1145/3343413.3377976>
- [32] Yihan Lu and I-Han Hsiao. 2017. Personalized information seeking assistant (PiSA): From programming information seeking to learning. *Information Retrieval Journal* 20, 5 (Oct. 2017), 433–455. DOI : <https://doi.org/10.1007/s10791-017-9305-y>
- [33] Les Nelson, Christoph Held, Peter Pirolli, Lichan Hong, Diane Schiano, and Ed H. Chi. 2009. With a little help from my friends: Examining the impact of social annotations in sensemaking tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1795–1798. DOI : <https://doi.org/10.1145/1518701.1518977>
- [34] Heather L. O’Brien, Andrea Kampen, Amelia W. Cole, and Kathleen Brennan. 2020. The role of domain knowledge in search as learning. In *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval*. Association for Computing Machinery, Vancouver BC, Canada, 313–317. DOI : <https://doi.org/10.1145/3343413.3377989>
- [35] Srishti Palani, Zijian Ding, Stephen MacNeil, and Steven P. Dow. 2021. The “Active Search” hypothesis: How search strategies relate to creative learning. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*. Association for Computing Machinery, New York, NY, 325–329. DOI : <https://doi.org/10.1145/3406522.3446046>
- [36] Georg Pardi, Johannes von Hoyer, Peter Holtz, and Yvonne Kammerer. 2020. The role of cognitive abilities and time spent on texts and videos in a multimodal searching as learning task. In *Proceedings of the 2020 Conference on*

- Human Information Interaction and Retrieval*. Association for Computing Machinery, Vancouver BC Canada, 378–382. DOI: <https://doi.org/10.1145/3343413.3378001>
- [37] Sihang Qiu, Ujwal Gadiraju, and Alessandro Bozzon. 2020. Towards memorable information retrieval. In *Proceedings of the 2020 ACM SIGIR on International Conference on Theory of Information Retrieval*. Association for Computing Machinery, New York, NY, 69–76. DOI: <https://doi.org/10.1145/3409256.3409830>
 - [38] Nirmal Roy, Felipe Moraes, and Claudia Hauff. 2020. Exploring users' learning gains within search sessions. In *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval*. Association for Computing Machinery, Vancouver BC, Canada, 432–436. DOI: <https://doi.org/10.1145/3343413.3378012>
 - [39] Nirmal Roy, Manuel Valle Torre, Ujwal Gadiraju, David Maxwell, and Claudia Hauff. 2021. Note the highlight: Incorporating active reading tools in a search as learning environment. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*. Association for Computing Machinery, New York, NY, 229–238. DOI: <https://doi.org/10.1145/3406522.3446025>
 - [40] Ladislao Salmerón, Pablo Delgado, and Lucia Mason. 2020. Using eye-movement modelling examples to improve critical reading of multiple webpages on a conflicting topic. *Journal of Computer Assisted Learning* 36, 6 (2020), 1038–1051. DOI: <https://doi.org/10.1111/jcal.12458>
 - [41] Rohail Syed and Kevyn Collins-Thompson. 2017. Optimizing search results for human learning goals. *Information Retrieval Journal* 20, 5 (Oct. 2017), 506–523. DOI: <https://doi.org/10.1007/s10791-017-9303-0>
 - [42] Rohail Syed, Kevyn Collins-Thompson, Paul N. Bennett, Mengqiu Teng, Shane Williams, Dr. Wendy W. Tay, and Shamsi Iqbal. 2020. Improving learning outcomes with gaze tracking and automatic question generation. In *Proceedings of The Web Conference 2020*. Association for Computing Machinery, New York, NY, 1693–1703. DOI: <https://doi.org/10.1145/3366423.3380240>
 - [43] Shayan A. Tabrizi and Azadeh Shakery. 2019. Perspective-based search: A new paradigm for bursting the information bubble. *FACETS* 4, 1 (2019), 350–388. DOI: <https://doi.org/10.1139/facets-2019-0002>
 - [44] Paul Thomas, Peter Bailey, Alistair Moffat, and Falk Scholer. 2015. Towards nuanced system evaluation based on implicit user expectations. In *Proceedings of the Information Retrieval Technology (Lecture Notes in Computer Science)*, Guido Zuccon, Shlomo Geva, Hideo Joho, Falk Scholer, Aixin Sun, and Peng Zhang (Eds.), Springer International Publishing, Cham, 332–344. DOI: https://doi.org/10.1007/978-3-319-28940-3_26
 - [45] Kelsey Urgo, Jaime Arguello, and Robert Capra. 2020. The effects of learning objectives on searchers' perceptions and behaviors. In *Proceedings of the 2020 ACM SIGIR on International Conference on Theory of Information Retrieval*. Association for Computing Machinery, Virtual Event Norway, 77–84. DOI: <https://doi.org/10.1145/3409256.3409815>
 - [46] Nino Weingart and Carsten Eickhoff. 2016. *Retrieval techniques for contextual learning*. SAL @ SIGIR
 - [47] Ryen W. White and Jeff Huang. 2010. Assessing the scenic route: Measuring the value of search trails in web logs. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. Association for Computing Machinery, New York, NY, 587–594. DOI: <https://doi.org/10.1145/1835449.1835548>
 - [48] Barbara Wildemuth, Elaine G. Toms, and Luanne Freund. 2014. Untangling search task complexity and difficulty in the context of interactive information retrieval studies. *Journal of Documentation* 70, 6 (Oct. 2014), 1118–1140. DOI: <https://doi.org/10.1108/JD-03-2014-0056>
 - [49] Teena Willoughby, S. Alexandria Anderson, Eileen Wood, Julie Mueller, and Craig Ross. 2009. Fast searching for information on the Internet to use in a learning context: The impact of domain knowledge. *Computers & Education* 52, 3 (April 2009), 640–648. DOI: <https://doi.org/10.1016/j.compedu.2008.11.009>
 - [50] Max L. Wilson, Paul André, and mc schraefel. 2008. Backward highlighting: Enhancing faceted search. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, 235–238. DOI: <https://doi.org/10.1145/1449715.1449754>
 - [51] Wan-Ching Wu, Diane Kelly, Ashlee Edwards, and Jaime Arguello. 2012. Grannies, tanning beds, tattoos and NASCAR: Evaluation of search tasks with varying levels of cognitive complexity. In *Proceedings of the 4th Information Interaction in Context Symposium (IIIX'12)*. Association for Computing Machinery, New York, NY, 254–257. DOI: <https://doi.org/10.1145/2362724.2362768>
 - [52] Luyan Xu, Xuan Zhou, and Ujwal Gadiraju. 2020. How does team composition affect knowledge gain of users in collaborative web search?. In *Proceedings of the 31st Association for Computing Machinery Conference on Hypertext and Social Media*. Association for Computing Machinery, New York, NY, 91–100. DOI: <https://doi.org/10.1145/3372923.3404784>
 - [53] Ran Yu, Ujwal Gadiraju, Peter Holtz, Markus Rokicki, Philipp Kemkes, and Stefan Dietze. 2018. Predicting user knowledge gain in informational search sessions. In *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. Association for Computing Machinery, New York, NY, 75–84. DOI: <https://doi.org/10.1145/3209978.3210064>

Received January 2021; revised August 2021; accepted November 2021