ClassSearch: Facilitating the Development of Web Search Skills through Social Learning

Neema Moraveji[1], Meredith Ringel Morris[2], Dan Morris[2], Mary Czerwinski[2], Nathalie Riche[2]
[1]Stanford University, Stanford, CA, USA
[2]Microsoft Research, Redmond, WA, USA
neema@moraveji.org, {merrie, dan, marycz, nath}@microsoft.com

ABSTRACT
We explore the use of social learning – improving knowledge skills by observing peer behavior – in the domain of Web search skill acquisition, focusing specifically on co-located classroom scenarios. Through a series of interviews, pilot studies, and classroom deployments, we conclude that a peripheral display of Web search activity within a classroom facilitates both social learning and teacher-led discourse. We present the ClassSearch system for shared awareness of Web search activity, which embodies principles gleaned from our iterative design process, and show results from a ClassSearch deployment in twelve middle-school classroom sessions. Finally, we highlight design suggestions and opportunities for future work while taxonomizing the space of co-located search pedagogies.

Author Keywords
Web search, education, search expertise.

ACM Classification Keywords
H.3.3. Information Systems: Search Process.

General Terms
Human factors

INTRODUCTION

Expertise in Web-based information retrieval has been studied in great detail recently [2, 32], owing largely to the increased relevance of Web search to many information needs and tasks. However, little consensus has been reached regarding how to incorporate search skills education into school curricula, and schools often have limited resources to devote specifically to Web search education.

We hypothesize that allowing expertise to spread socially among searchers will be an engaging and effective way to integrate search skills instructions into traditional classrooms. Leveraging social learning goes beyond benefits of scale: social learning theorists have long extolled the benefits of designing a sociocultural context to motivate learning. A 'community of practice' [19] provides avenues for advancing in levels of expertise, but does not exist in the domain of search. Furthermore, social learning also allows learning to occur implicitly, without courses structured around teaching search. That is, search can be learned in an authentic context [7] by building on the research process already incorporated into traditional school subjects (e.g., history, literature).

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

Figure 1: The ClassSearch system (shown here in a middle-school classroom deployment) provides a shared-awareness display of Web search activity within a classroom, facilitating social learning and teacher-led discourse about Web search skills. Students search using individual computers and results are shown on an instructor-controlled shared display.

In this paper, we present the design and evaluation of ClassSearch, a system to support these insights. We do so through an iterative, user-centered design process informed by feedback from 8 instructors and 8 students. We then present the results of an in situ study with 160 middle school students who used the system. The results motivate future work on systems to improve opportunities for social learning of search expertise with co-located users.

This work makes four primary contributions: (1) We introduce the use of social learning as a mechanism for Web search skill development; (2) We present results from the iterative design of a shared display for search data in classroom environments; (3) We describe the results of a deployment of ClassSearch in 12 middle school classroom sessions; (4) We discuss the design implications of our observations, and taxonomize the space of co-located search tasks and associated pedagogies.

RELATED WORK

Understanding search expertise

A variety of measures of search performance are used in information retrieval research and agreement exists that

certain factors are relevant: domain expertise [2, 31], knowledge of the search engine's features [14], general literacy [17, 21], and knowledge of Web resources [31]. Condensing these disparate areas of expertise into measurable definitions of search skills, however, remains elusive. Although qualitative metrics for information literacy have been proposed [1], causes of differentiated expertise among searchers are not well-understood.

Formulating queries is particularly complex task due to its reliance upon abstraction, vocabulary, domain knowledge, and grammar. Despite this complexity, query formulation is often the focus of explicit instruction because it is so central to Web search practice [9,12].

Prior work studying the habits and behaviors of students of various ages using the Internet in directed or self-directed tasks have associated novice searchers with the inability to identify credible sources [10], issuing vague queries [4], focusing on surface features of websites [5], and limited confidence about and awareness of Web resources [8].

Supporting searchers with instruction & tools

Web search curricula have been proposed (e.g., [5, 18, 30]), though none are yet established as standards. Games [11] and inquiry-based learning [15] have been explored as methods of making search skill learning "authentic" (i.e., allowing the learner to address an actual need during instruction). Instructional guidelines have been proposed for teaching search skills [1, 9, 20, 29]. These approaches have focused on procedural learning, not social learning.

Several novel interfaces have been proposed to ease query formulation [16, 27] or visualize search result sets [26, 33]. Lin et al. [22] used visualization to give researchers deeper insight into search strategies. Navigation Flow Maps [23] enabled a categorization of search strategies among university students. PadPrints [13] used a similar visualization to facilitate within-browser reflection on search history. SearchTogether [25] and CoSearch [2] employed multi-user search visualization to facilitate collaborative search tasks. These approaches do not attempt to impart search expertise on searchers, but rather ease the process of searching itself. They are thus a complement to the present work, which explores the *transfer of expertise* among searchers.

Social learning of search expertise

The concept of imitating a skilled practitioner is perhaps the oldest and most natural form of social learning [24]. However, simply imitating the queries used or links clicked by an expert does not directly enhance a searcher's skill set; it is the mental processes underlying search that are important to imitate. This is because information needs vary from moment to moment, making the conceptual model of Web search and navigation paramount. In postulating that these processes may be socially transferred, we draw upon two theories: *legitimate peripheral participation* [19], which describes how novices in a community develop expertise, and *cognitive apprenticeship* [6], which posits that learners can imitate experts by deriving components of their thought processes, not just their behavior.

Motivation for present work

The existing literature on search expertise and social learning, coupled with existing systems for exposing Web search strategies for non-pedagogical reasons, led us to hypothesize that social learning is particularly applicable in the domain of Web search education. The subject matter in this domain varies constantly as search engines, interfaces, and corpora evolve, and limited curriculum time is presently devoted specifically to search education. These factors motivated our team to develop a system whereby search processes can be observed and discussed casually, in classrooms still focused on traditional school subjects.

INTERVIEWS WITH SEARCH INSTRUCTORS

In order to inform the development of a system to support social learning of Web search skills in classrooms, we conducted six phone interviews with teachers and librarians for one hour each. The instructors came from public schools in moderately affluent areas of the United States. This section summarizes insights drawn from these interviews.

Search curriculum

Instructors indicated that sixth grade (10-12 years old) is the earliest point at which they began teaching students about search as a skill, beginning with very focused (often single-query) tasks. As the students grow in age, they are given progressively more open-ended tasks.

One of the interviewee's schools explicitly taught search as a skill during a 7-10 day "mini-class" as part of a larger class on computer usage, but in general search skills was restricted to a single lecture or less. Of all the possible subtopics to cover in such instruction, the instructors said they focus most on query formulation and refinement, choosing a search result, and discerning credibility.

There are many different search curricula and pedagogies. One teacher remarked that teaching search is like being "an elementary school art teacher – there is no standard curriculum. The teacher decides what to do, with varying amounts of success."

The most common pedagogy was to have the instructor demonstrate search practices on a projected display at the front of the room. This might be followed by one or more discursive practices around the search task such as brainstorming possible query words as a class and listing criteria for determining a site's credibility. The students would then be "released" to do their own investigations while the instructor roams to address individual problems.

Another common teaching method was the use of side-by-side query comparisons: typically, two students working in a computer lab would execute different queries, and the results would be compared aloud.

Determining which part of a Web page is relevant to the existing query is another skill deemed important to teach. Though the use of browser shortcut keys to find keywords

within a page was often mentioned to students, this functionality has low visibility and is often forgotten.

Challenges faced by instructors in class

According to interview participants, aggregate data about the class's search behavior is difficult or impossible for the instructor to glean. However, instructors told us there is a large range of search expertise in middle school classrooms. The majority of instructors we spoke with did not see themselves as experienced search experts (and students often do not always see instructors as experts), which creates an atypical dynamic in the classroom for this topic.

When an instructor approaches a student to aid in a Web search task, a form of over-the-shoulder learning [28], it is difficult and/or inefficient for the instructor to discover how the student arrived at the current navigation state (e.g., what queries they had tried, what pages they had seen). This makes search instruction very different from math, where the instructor can immediately "see the student's work" and troubleshoot individual steps. Without this affordance, feedback to students about search is often disconnected from the individual steps taken in the search process. One instructor stated: "I'd like to see what page the student is on and how they got there." Further, this over-the-shoulder approach does not scale well to typical class sizes (about 30 students), due to its inefficiency.

Teachers faced several common hurdles during search instruction. The most complex of these was communicating the algorithm for search result ranking. Furthermore, interview participants discussed students becoming preoccupied with the estimated number of results returned by a search; this is also a phenomenon we noted in our own classroom observations. This, we believe, is due to the common practice of instructors attempting to use result count as part of an easily-understood metric for search quality. This idea sometimes led to students being preoccupied with getting the "right" number of results or claiming some searches are "better" or "worse" than others based on this metric.

Another hurdle instructors faced was encouraging students to use domain-specific vocabulary, rather than colloquial terms, in their queries. Advanced operators are rarely used without a reminder from instructors.

Instructors also discussed students' perceived association between site popularity and site credibility; students often assume that large portals such as Wikipedia can be relied on for credible information. By keeping peer investigations largely invisible, this practice is reinforced. Instructors reported providing a common metric for reliability: one can judge a site in part by its domain suffix: ".edu is better than .org which is better than .com."

Logistical limitations were mentioned as well: comparing sets of search results is difficult with current tools. Results must be read aloud to the class, or students need to bring individual computers to the front of the class to attach to a projected display, disrupting class flow and prohibiting retrospective analysis of queries. In either case, it is difficult for *other* students to investigate more deeply, as only two students are participating in the exercise.

Existing tools for supporting instructors

We asked instructors about the tools they currently use during search instruction. Projectors are often used to show example searches at the front of the room. Whiteboards and chalkboards are used to enumerate query suggestions, useful sites, or related information for later reference. Teachers also mentioned monitoring software generally targeted at keeping individuals on-task and not specific to search instruction.

netTrekker [nettrekker.com] is a popular search engine to which schools purchase access; our interview participants generally worked in schools where netTrekker was the preferred search engine. The results of netTrekker are curated by a team of paid instructors, and the results can be filtered to target certain age groups. However, while instructors appreciated the high-quality information on curated sites such as netTrekker, they expressed concern that such tools do not expose students to the issues they will need to address when on the "open" Web.

THE CLASSSEARCH SYSTEM

ClassSearch architecture

Inspired by these interviews, we created ClassSearch, a system that makes aggregate (and, optionally, individual) student search data from the class visible via a large display. Instead of showing such data to the instructor alone to influence instruction, this approach makes data available to students, enabling conversations that use the data as a co-reference. ClassSearch is intended for use in a setting where the instructor has a projected display and each student has a computer workstation (such as a library, a school's computer lab, or a laptop-augmented classroom) (Figure 1).

We considered and prototyped designs that provided a separate, private display to the instructor, but discussions with instructors suggested that this would be impractical, given the degree to which instructors walk around the room during research sessions. Furthermore, instructors' lack of familiarity with multi-display systems rendered this approach prohibitively complex.

One common practice instructors described was demonstrating a search strategy on a projected display and allowing that strategy to propagate through the class. The idea that students emulate practices seen at the front of the room also informed our decision to use a shared display: in this environment, students may emulate not only instructor behavior, but peer contributions as well. This approach also follows directly from our interviewees' self-reported perceptions that they were not necessarily more skilled in search than some students.

Instructors frequently visited students at their computers, and desired a mechanism to quickly assess the search pattern that brought a student to the current browser state.

Consequently, we chose to include a component in our system that made recent query history (as opposed to complete navigation history) visible on a student's monitor, primarily to allow quick over-the-shoulder assessment by an instructor.

The ClassSearch system thus contains two components:

1) A browser add-on running on each student's computer that captures searches and Web browsing history, forwards that data to the server, and displays recent query history in a small, adjacent window.

2) A server computer, controlled by the instructor, that projects data (via the add-on) from student search behavior at the front of the room.

In this section, we describe the major features available on each system component, which were designed primarily to meet the needs identified in our preliminary interviews. For brevity, we describe only the feature set available in the final version of the software; in the following section ("Pilot Studies"), we will highlight features that were added or changed as a direct result of teacher and student pilot studies conducted around our first prototype.

Shared display

We prototyped several designs that attempted to surface complete query and navigation trails for all students, but we found that this approach did not scale well to classes at our target size (around 30 students). We experimented with abstractions that did allow this volume of data to be aggregated (such as iconic views of navigation patterns), but teachers perceived these as too high-level to support discussion around low-level search skills (though this approach will surface later in our "search strategies" view for quickly assessing individual student behavior). Through discussions with students and instructors, we identified *query terms* and *sites visited* as the most important points for discussion, leading to two of the data representations we used on our shared display.

In this section, we describe the four components available

Figure 2: The shared display in its default configuration.

in various layouts in ClassSearch: *query cloud*, *site cloud*, *search strategies*, and *individual student*. When it first starts up, the ClassSearch shared display shows the query cloud and site cloud views, along with a list of students and a Web browser (Figure 2). The student list allows quick access to the search strategies view or to the individual student view (by clicking on a student's name). The Web browser supports the existing practice of instructor-led demonstration. Individual components can be maximized to occupy the entire display.

The overview window also includes a configuration dialog, which allows the instructor to save and load the state of the entire display (for resuming class sessions on subsequent days), and allows the instructor to remotely reset the browser add-on state of an individual or the entire class.

Query cloud

The *query cloud* is a view of queries that have been issued in the class (Figure 3). This view is updated in real time; it both supports discussion about query formulation and provides "hints" for students who may be having difficulty with query formulation or refinement. Our first implementation of this component used *query-level* elements, in which each query as a whole was treated as an individual term, sorted alphabetically and sized according to frequency. However, this simple approach presents two scalability issues. First of all, alphabetical sorting can quickly lead to important (common) queries disappearing below the fold; although the view is scrollable, the teacher is frequently away from the display, and the need for scrolling reduces the "at-a-glance" value of the display for students. Second, until the class has issued a large number of queries, exact query matches among students may be infrequent, limiting the value of font scaling as a means to indicate trends within the class.

We made two modifications to the query cloud to address these issues. First, we sorted the cloud by frequency (most common queries first), rather than alphabetically, to support easier observation of class trends. Second, we created an *n-gram* mode, which shows all words that appeared in queries along with any higher-order n-grams (bigrams, trigrams, etc., possibly including whole queries) if they appeared in

Figure 3: Query Cloud (n-gram view). One of the two query cloud views, in which user queries are broken up into all possible words or phrases; these are displayed in the cloud, sized according to popularity. The instructor's mouse hovers over 'volcano,' revealing the list of queries using that term.

Figure 4: Site Cloud. Pages visited by students are aggregated by site and sized by popularity. The instructor can hover over a site to see all the pages visited within that site.

two or more queries. We included both n-gram and full-query modes in the deployed version of ClassSearch.

In n-gram mode, hovering the mouse over a term shows all the queries issued that contained a particular n-gram. For the query-level view, hovering the mouse over a term optionally shows the names of all students that issued that query. Clicking a term in the query cloud shows the associated search results in an embedded browser. This functionality is intended to support discussions like "Was this a good query?" or "Let's see the search results for the three most popular queries."

Site cloud

Instructors indicated that discussion about search skills often focused on site credibility; this motivated the *site cloud* view (Figure 4), which summarizes sites visited by the class. This view supports class-wide discussion about sites and provides suggested starting points to students who may be having difficulty with a search.

The site cloud shows the sites visited by the class, sorted and sized according to frequency (for the same reason we chose to sort the query cloud according to frequency). The reason sites are aggregated, rather than showing the URLs of specific pages, is twofold. First, we anticipated instances where multiple people are visiting the same individual pages to be relatively uncommon, so collapsing within sites is more likely to show trends within a class. Second, in our observations and interviews, credibility of pages was often assessed based on heuristics concerning the site and not individual pages.

When the instructor hovers over a site, all the pages students have visited in that site are displayed in a tooltip, sorted by frequency. Clicking opens a series of browser tabs containing each student-visited page within the site. This functionality is intended to support discussions like "Did this page contain the answer we're looking for?"

Search strategies view

While the primary goal of the ClassSearch system is to promote discussion about search behavior in aggregate, instructors also expressed a desire to quickly assess individual and group progress. The shared display, however, requires careful consideration of student privacy; frequent visibility of a student's complete search history would be potentially uncomfortable and might discourage exploration. The *search strategies* view (Figure 5) attempts to

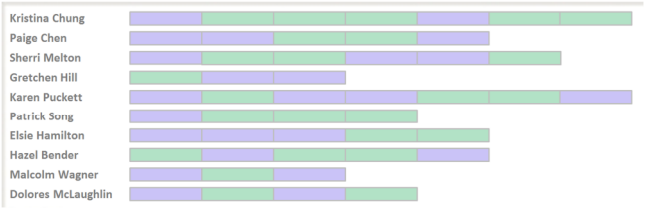

Figure 5: The *search strategies view* shows an overview of individual students' actions. Purple boxes represent queries, green boxes represent page visits. For privacy, we use pseudonyms for students throughout this paper.

balance these factors by abstracting student search patterns to allow instructors to quickly assess the progress of the class as a whole and identify students who might require assistance or students exhibiting behaviors worthy of highlighting to the rest of the class.

This view shows a graphical summary of each student's search history: purple boxes represent queries, and green boxes represent pages visited**.** Hovering over a purple box shows the corresponding query; hovering over a green box shows the corresponding URL. Clicking an item from the summary visualization opens the associated search results list or Web page in the browser.

Individual student view

An instructor can "drill down" on an individual student (by clicking their name in the search strategies view) to see a detailed view of that student's pages visited and queries executed during the search task (Figure 6). This view groups pages visited by query, presenting a short-term, search-centric browser history for that student, allowing quick summarization of a student's progress and search strategy. Furthermore, in cases where a student is identified as having a "good" strategy, this allows the instructor to share that student's strategy for discussion, specifically addressing our goal of transferring expertise from more-skilled to less-skilled students.

Student browser add-on

Each student's computer ran a Web browser add-on to collect search and page visit data and send that data to the teacher-controlled server computer. This component also presented a concise history of recently-executed queries on the student's monitor, attached to the bottom of the student's foreground browser window (Figure 7). This was inspired by feedback provided in preliminary interviews suggesting that over-the-shoulder assessment of students' search paths was a significant challenge for instructors as

Figure 6: The *individual student view* shows the queries executed and pages visited by a specific student, summarizing the student's progress and search strategy.

Figure 7a (top) and 7b (bottom): a) A student's browser with the add-in showing the query trail; a term has been clicked and instances of that term are highlighted in the browser. b) A close-up of the add-in showing each query, broken into terms.

they walked around the room visiting students. Queries can be concisely represented in text, and one query often summarizes an entire path of exploration, so we chose to include only queries – rather than a complete browser history – to maximize the amount of information provided to a visiting instructor while minimizing screen footprint.

Clicking any query in the history executes that query in the student's browser, and clicking any term in a query highlights occurrences of that term in the current page (Figure 7a). This functionality closely resembles a feature available in popular browser toolbars such as the Google Toolbar, and was motivated by teachers' lamenting students' inability to find relevant content within a page.

PILOT STUDIES

We conducted two pilot studies with our first iteration of ClassSearch: one with teachers and one with students in our target age range. In both cases, participants came to our lab and we created a "classroom" in a meeting room. We distributed pre- and post-surveys.

Pilot study with teachers

We brought six middle-school instructors into our lab for feedback on a small ClassSearch installation by participating in simulated classroom scenarios. Four of these instructors had also participated in our initial interviews. In the pilot study, we had each participant act as a "student," and the facilitator acted as the instructor, using search tasks adapted from [8].

Teachers remarked that being able to see which student had performed each query on the large display adds accountability. They believed this functionality would help mitigate off-task behavior, trigger interesting discussions, and group students who are doing complementary work. Instructors expressed a desire to turn off the display of student names for certain scenarios (a feature we subsequently added), but overall were not concerned with privacy in the classroom.

Instructors stressed the value of being able to delete individual terms from the display to mitigate awkward situations and to dissuade abuse. Instructors desired the ability to save and restore the state of the class's dataset because pedagogical sequences often span sessions. These features were also added after the pilot studies.

Instructors suggested that the history shown on the student's computer would discourage spending time on unrelated pages as the teacher walked around the room. Instructors also suggested that that the student history display would support discussions potentially useful for good conversations: "You can ask questions. Why'd you go here? Why'd you go there?"

Pilot study with students

Eight students entering the 6th grade came to our lab, where we set up a small ClassSearch installation. Two of the students knew each other beforehand. A researcher acted as the "instructor" and operated the shared display. The students completed four search tasks that varied in specificity, ranging from exploratory tasks to fact-finding tasks, based on tasks used by Druin et al. [8].

When asked their opinions about the query cloud, one student suggested it could help her think of a new query term, in particular by looking at the "big" (most popular) terms. In general, the students seemed focused more on the larger terms (rather than, e.g., the more unique terms or the more semantically relevant terms). One student said that her query "didn't work well" so she "looked at the query cloud and saw that 'vice president's birthday' was in a big font," so she decided to attempt that query. Another student remarked "it was easier to search when there were ideas up on the projector." When probed, the students generally acknowledged that a bigger (more popular) query does not necessarily mean a better query, and that popularity was not the only dimension upon which to evaluate queries or sites visited. Two students suggested that size of query or domain should indicate quality instead of popularity.

FIELD STUDY

ClassSearch represents a departure from existing practices of search instruction and classroom computer use. It ties individuals' usage together and creates novel feedback loops and visualizations of aggregated data that we hypothesized would impact search discourse and behavior. To investigate this hypothesis, we conducted a field study around a ClassSearch design that had been iterated to reflect the feedback from our pilot studies. Because the system presents new data in a novel format, our goals were to assess how it was appropriated and to highlight opportunities for further design iterations; a controlled, longitudinal assessment of ClassSearch's impact on learning outcomes is beyond the scope of the present work.

Participants

We conducted our field study in a public middle school (spanning grades 6-8 and ages 11-14) in an upper-middle-class suburb of Seattle. Students at the school were largely tech-savvy and regularly used computers at home. The school had a computer lab in the library that teachers used occasionally for class research projects, so computer use in class was unusual but not exceptional. Search was not given a great deal of emphasis in the school's state-defined curriculum, but was highlighted once or twice a year as part of a class library visit, typically coupled to a domain-specific research project.

For this field study, two instructors worked with us in their own classrooms. Both instructors were accustomed to the idea of using a computer with a projector at the front of the room to show students individual Web pages. Neither instructor had participated in our pilot studies.

One instructor was an eighth-grade American History teacher who used ClassSearch in four classes (approximately 50 minutes each) each day for two days in a row. Students in all four classes were conducting a research project; class time was not explicitly devoted to search instruction. The second was a sixth-grade instructor who taught search expertise in two of his classes on one day; each of these classes spanned two 50-minute sessions. Altogether, this gave us a total of six groups totaling 160 students that used the system in a total of twelve sessions.

Procedure

Prior to working with each instructor, researchers installed 30 laptop computers running the ClassSearch browser add-on in the appropriate classroom, along with an instructor (server) installation connected to a projector. We gave the instructors an opportunity to use the system the afternoon before the study began. During this session, we allowed them to test different features using a sample dataset preloaded into the shared display.

Students were not required to alter their normal searching practices to use either the shared display or browser add-on. They learned to use the system through real search tasks; no explicit demonstration was provided to students.

In the eighth-grade class, search expertise was not the goal of instruction, so students worked in small groups, where each group was focused on a particular research topic. For one such class, we installed a second server and projector in the classroom to explore the idea of having a shared display dedicated to a small, topic-specific group. The "main" projector was shared among the rest of the class, spanning several research groups.

FINDINGS

ClassSearch was actively used – students issued a total of 12,143 queries and visited 8,368 pages during the twelve evaluation sessions. Since our goal was to understand how teachers and students appropriated this social learning approach to search instruction, we present qualitative findings based on observations of these initial deployments. We discuss how the major features (query cloud, site cloud, and strategies view) were used, as well as unanticipated uses of the system that arose. Students used the re-finding features on their personal toolbars 411 times; however, as anticipated, the most interesting aspect of the system was the shared display, so we focus our findings there.

In general, students and teachers were enthusiastic about the technology (noting that ClassSearch was "seriously cool" and wondering "sweet, can I buy this?"). At the end of his class, one student pointed at the large display and noted "We need that *always*, so we can see what everybody is looking at and searching."

Query Cloud

Both instructors chose to keep the query cloud and site cloud open simultaneously (rather than maximizing one at the expense of the other). We had anticipated that instructors would click on items in the query cloud to display search results in order to compare and contrast the results returned by different variants of a query. In practice, though, they only clicked terms to display associated search results 11 times over the course of the observation period.

Rather than using the query cloud to explore the impact of query phrasing on results, the instructors instead valued the query cloud for the awareness it provided about synonyms, alternative phrasings, and syntax. One teacher commented that the cloud was helpful to the students because it helped them in "becoming aware of their own searching."

In particular, the query cloud enabled an awareness of aggregate behavior that was not previously possible. For instance, one teacher looked at the query cloud and noted that very few members of the class were using quotation marks when phrasing their queries. He pointed out to the class that one of the queries (*gay rights*) might be more effective if enclosed in quotations.

The query cloud also enabled the teacher to point out examples of interesting student behaviors. For example, when one student used an advanced operator in her query, the teacher called it out to the class as a point of discussion. The student explained that the operator had in fact dramatically improved her search results.

The students also paid attention to the query cloud, and it was often a topic of discussion. In particular, students were very interested in knowing who had done particular queries, and often asked about this out loud (since the list of names associated with a query was only visible if the teacher hovered his mouse over it).

Site Cloud

We anticipated that instructors would use the site cloud to select sites to visit in order to compare and contrast the quality of information found on different sites. Instructors did this occasionally, but not frequently (a total of 18 times). The teachers generally clicked on sites from *.edu* and *.org* domains in order to show the class examples of sites that they felt were credible, high-quality references.

Teachers found the site cloud valuable for spotting trends that were not previously apparent. For example, looking at which sites were large and small in the cloud, one teacher commented to his students about a trend: "people don't look at National Geographic, but it's a really good source."

The students also attended to the site cloud, as evidenced by discussions about the items it contained. In particular, the cloud prompted student discussions about the *quality* of particular sites ("the religious tolerance one is good, lots of info on it") and about the *process* of finding particular sites ("how did you search this one, Gretchen?").

Students requested an easier method to visit the sites others had found. For example, if an individual student wanted to visit a site that she noticed in the cloud, she had to type the URL into her browser – offering a shortcut to this behavior, such as the ability to "pull down" a site from the cloud, was desired. A teacher requested that students be able to indicate the quality of a site, and perhaps the site cloud font sizes could then reflect votes of quality.

Strategies View

The strategies view, which was not visible by default, was maximized 49 times, suggesting that this view was of special interest to instructors. It was generally used for short periods of time while the instructor specifically surveyed the progress of each student on the list; when the instructor walked away from the computer, he generally reset the view to display the query and site clouds. Instructors only drilled down to the individual student view 23 times.

The purple and green bars representing queries and sites in the strategies view became incorporated into the classes' vocabularies. For example, one teacher showed the strategies view and suggested that everyone in the class spend the last five minutes of the period in the "green" in order to focus on reading and taking notes on good quality sites, rather than continuing to search. Another teacher wanted to indicate that students should have completed their initial round of querying and begun reading content, and conveyed this by saying "you should probably have gotten to some green by now."

Both teachers and students desired metrics that would provide absolute indicators of search success and skill, and were eager to interpret the bars in the strategies view as such a metric. For example, referring to the strategies view, one student wondered how her search progress compared to another's, asking "am I doing well?" The bars also inspired a competitive feel among some students who assumed having longer bars was better – one even noted, "I'm going to do 1 million searches to win!" Two of the classes specifically discussed that the number of pages visited was not necessarily a meaningful metric of progress.

One teacher requested that the bars on the strategy view could indicate the duration of time a student spent on a particular page, perhaps to identify pages that might be of high quality or interesting to discuss with the entire class.

Unexpected Uses

Although ClassSearch was intended primarily as a tool to enhance search skill learning through social awareness, the teachers found additional value in the system as a means of keeping students on-task, both through their ability to monitor and address specific students' behavior and through the social pressures of behavior visibility which made students feel guilty about wandering off-task.

In the class where we introduced the second projector for the additional project groups, we were surprised to see that students not only looked at the public display for their own group (for all of the reasons mentioned earlier), but also found value in the other group's display as a glanceable summary of the other project. For example, noting the largest site "ciggyfree.com" and search term "cigarettes" gave students a quick overview of what the other group had decided to focus their inquiry on.

DISCUSSION

Our deployment of ClassSearch demonstrated that both teachers and students found pedagogical value in a public, interactive, real-time visualization of a classroom's query and browsing behavior. This simple intervention appeared to support a variety of pedagogical styles, including classes dedicated specifically to search skills instruction and classes focused on domain-specific research tasks.

Co-located search pedagogies

ClassSearch was designed for classroom scenarios in which students and instructor are co-located. However, our initial interviews, pilot studies, and field study revealed that a large spectrum of search tasks and pedagogies exists even within this space. In particular, we observed classroom tasks that varied along a continuum of individuation. We taxonomize the space of co-located search pedagogies, highlighting design recommendations to optimize ClassSearch for each scenario:

Individual topics

At the individualized end of the continuum, each student is researching a separate topic. This pedagogy was observed in eighth-grade classrooms during our field study, and is most appropriate for older students who have achieved some independence in both search and domain-specific skills. In these cases, few terms in the ClassSearch query cloud grew to large font sizes due to low overlap in usage. As a result, a student may notice query terms that may have nothing to do with his or her topic. While this presents some risk of distraction, it offers instructors a valuable opportunity for discussing overlap among related topics, and offers students lightweight awareness of other topics. Modifying ClassSearch to display structural rather than semantic commonalities among students might be a preferred alternative, for instance focusing the query cloud on trends in advanced operator use rather than keywords.

Topical sub-groups

In this type of classroom search task, the class is split into sub-groups searching for different material, but students

within each group share a topic. This led to a mix of individuated and group use of ClassSearch. One practice we observed was students working together without using the display but then taking a "break" to look at the ClassSearch display for inspiration in their own tasks or to reflect on the impact that their activity had on the large display. Groups could have some sense of other groups' topics and progress. Multiple ClassSearch installations within the same room, such as the two-projector arrangement explored during our field study, can benefit this pedagogical scenario.

Shared topic

Instructors often had a theme for a class or a portion of the class that they asked students to investigate, study, or write about. An example of this was "What is your favorite Greek god and why?" In these cases, the students are all searching for different specific information but it falls under a common umbrella, and may include common queries, particular early in the task. In this class structure, students often attempted to encourage peers to visit a page they liked by pointing at their screens or at the common ClassSearch display. This frequently-observed behavior may provide a valuable opportunity for discussion of search strategies among peers. Adding the ability to "pull down" a URL from the site cloud as a jumping-off point for individual exploration, a feature requested by students in our study, would add value to this scenario.

Shared task

In settings where search skills are a primary curricular objective, an instructor may ask all the students to complete the same task, e.g., "I want everybody to type in 'birth rate Norway' into the search box." These tasks allowed the instructor to guide student attention to some extent and might be followed by a free-form query to refine the search. This type of task is most appropriate for teaching search skills explicitly, and may be appropriate for younger students with only limited search skills (this approach was observed in our field study in sixth-grade classes). When ClassSearch is used in this environment, the initial (shared) query is very large (and not very useful), but the subsequent terms are understood easily as refinements of the common original. Modifying the query cloud to use a timeline or network-layout visualization, rather than one sorted by frequency, might benefit this pedagogical approach.

Metrics of search success

By providing visualizations of aggregate information, ClassSearch provided representations of data meant to prompt reflection on the search process. While ClassSearch appears to have been successful in this respect, it was also apparent that instructors and students desired objective metrics that would indicate search skill and quality. It is not clear that such metrics exist, as defining and modeling search expertise remains an open problem [2, 31, 32].

However, both students and instructors were eager to interpret various aspects of the display as evaluative. As mentioned earlier, the length of bars in the search strategies view was sometimes interpreted as a positive indicator of students' search skill, students and teachers conflated the number of results returned with query quality, and one teacher seemed to count the number of times queries using operators such as quotations appeared in the query cloud as a metric of the class's search sophistication. Additionally, in both the pilot and field studies participants requested indicators of query and site quality in addition to popularity.

The lessons in this regard are twofold: (1) designers of search education tools should be careful to avoid interface elements that can mistakenly be interpreted as metrics, since end-users seem eager to do so, and (2) research into identifying such metrics (if they exist) would potentially be highly valued, particularly by the education community.

In addition to short-term metrics of search success, long-term metrics reflecting students' learning may also be valuable, e.g., reflecting to the teacher and/or student whether and how a student's search behavior evolves over time. Due to the short-term nature of our deployment, which prevented measuring learning outcomes, we did not endeavor to develop such metrics. Adding an additional iterative stage to our design process, such as meeting again with teachers after a moderate-length deployment of Class Search to reflect on student behaviors, might be one way to develop such metrics.

CONCLUSION

Web search skills are an increasingly necessary literacy in the digital era, but existing tools do not effectively support search instruction, nor do they enable potentially effective social pedagogies. In this paper, we proposed that using a public display to show classroom-level information regarding querying and browsing behavior can support social learning of Web search skills.

We presented the results of an iterative design process involving interviews and pilot studies with stakeholders such as teachers, librarians, and middle-school students, that revealed areas for improving upon current search skills pedagogies, and directly informed the design of ClassSearch. We described ClassSearch, and reflected on how students and teachers appropriated the system during a deployment in twelve middle-school class sessions.

We found that ClassSearch enabled new types of discussion surrounding reflection on individual students' actions and on class-level trends. In addition to identifying opportunities for further improving the system's design, our observations also enabled us to articulate a design space for classroom-based search education interventions.

Future work includes improving ClassSearch based on the feedback from our initial deployment, arranging long-term deployments to assess learning outcomes, and considering the suitability of shared displays of aggregate activity for other educational goals.

REFERENCES

1. ACRL. Objectives for Information Literacy Instruction: A Model Statement for Academic Librarians. *ACRL Standards and Guidelines*, 2001.
2. Amershi, S. and Morris, M.R. CoSearch: A System for Co-located Collaborative Web Search. *Proc CHI 2008*.
3. Bhavnani, S.K. Domain-specific search strategies for the effective retrieval of healthcare and shopping information. *CHI 2002 Extended Abstracts*.
4. Bilal, D. Children's use of the Yahooligans! Web Search Engine: I. Cognitive, Physical, and Affective Behaviors on Fact-Based Search-Tasks. *J Amer Soc for Info Sci*, 51.7 (2000), 646-5.
5. Coiro, J. Rethinking comprehension strategies to better prepare students for critically evaluating content on the Internet. *New Eng Reading Assoc J*, 39 (2003), 29-34.
6. Collins, A., Brown, J.S., Newman, S.E. Cognitive apprenticeship: Teaching the craft of reading, writing and mathematics. *Knowing Learning and Instruction*, 8.1 (1989), 453-94.
7. Dewey, J. *Democracy and Education*. New York: The Macmillan Company, 1916.
8. Druin, A., Foss, E., Hutchinson, H., Golub, E., Hatley, L. Children's roles using keyword search interfaces at home. *Proc CHI 2010*.
9. Fernández-Luna, J.M., Huete, J.F., Macfarlane, A., Efthimiadis, E.N. Teaching and learning in information retrieval. *Information Retrieval*, 12.2 (2009), 201-26.
10. Flanagin, A.J., Metzger, M.J. The perceived credibility of personal Web page information as influenced by the sex of the source. *Computers in Human Behavior*, 19.6 (2003), 683-701.
11. Halttunen, K. Sormunen, E. Learning Information Retrieval through an Educational Game. Is Gaming Sufficient for Learning? *Education for Information*, 18.4 (2000), 289-311.
12. Henry, L.A. SEARCHing for an answer: The critical role of new literacies while reading on the Internet. *The Reading Teacher*, 59 (2006), 614-27.
13. Hightower, R. R., Ring, L. T., Helfman, J. I., Bederson, B. B., Hollan, J. D. Graphical multiscale Web histories: a study of padprints. *Proc Hypertext 1998*.
14. Hölscher, C. and Strube, G. Web search behavior of Internet experts and newbies. *Computer Networks*, 33.1 (2000), 337-46.
15. Jones, G. J. An inquiry-based learning approach to teaching information retrieval. *Information Retrieval*, 12.2 (2009), 148-61.
16. Jones, S. Graphical query specification and dynamic result previews for a digital library. *Proc UIST 1998*.
17. Kodagoda, N., Wong, B. L. Effects of low & high literacy on user performance in information search and retrieval. *Brit HCI Group Ann Conf on HCI*, 2008.
18. Kuiper, E., Volman, M., Terwel, J. The Web as an Information Resource in K-12 Education: Strategies for Supporting Students in Searching and Proceedings ofessing Information. *Rev of Educ Res*, 75.3 (2005).
19. Lave, J. & Wenger, E. *Situated Learning: Legitimate Peripheral Participation*. Cambridge Univ Press, 1991.
20. Lazonder, A. Principles for designing web searching instruction. *Education and Information Technologies*, 8.2 (2003), 179-93.
21. Leu, D.J., Zawilinski, L. The New Literacies of Online Reading Comprehension. *New Eng Reading Assoc J*, 43.1 (2007), 1-7.
22. Lin, X., Liebscher, P. Marchionni, G. Graphical representations of electronic search patterns, *J Amer Soc for Info Sci*, 42.7 (1991), 469-78.
23. Lin, C., Tsai, C. Navigation Flow Map Method of Representing Students' Searching Strategies on the Web. *World Conf on Educational Multimedia, Hypermedia and Telecommunications*, 2005.
24. Meltzoff, A.N., Kuhl, P.K., Movellan,J., Sejnowski, T.J. Foundations for a New Science of Learning. *Science*, 325.5938 (2009), 284-88.
25. Morris, M.R., Horvitz, E. SearchTogether: An Interface for Collaborative Web Search. *Proc UIST 2007*.
26. Shneiderman, B., Feldman, D., Rose, A., Grau, X. F. Visualizing digital library search results with categorical and hierarchical axes. *Proc ACM Conf on Digital Libraries*, 2000,.
27. Spoerri, A. InfoCrystal: a visual tool for information retrieval & management. *Proc Intl Conf on Information and Knowledge Management*, 1993.
28. Twidale, M. B. Over the Shoulder Learning: Supporting Brief Informal Learning. *JCSCW*, 14.6 (2005), 505-547.
29. Vine, R. Real people don't do Boolean: How to teach end users to find high-quality information in the Internet. *Information Outlook*, 5.3 (2001), 16-23.
30. Walton, M., Archer, A. The Web and information literacy: scaffolding the use of web sources in a project-based curriculum. *British J Educ Tech*, 35.2 (2004), 173-86.
31. White, R., Dumais, S., Teevan, J. Characterizing the Influence of Domain Expertise on Web Search Behavior. *Proc WSDM 2009.*
32. White, R. Morris, D. Investigating the Querying and Browsing Behavior of Advanced Search Engine Users. *Proc SIGIR 2007*, 255-262.
33. Zavesky, E., Chang, S., Yang, C. Visual islands: intuitive browsing of visual search results. *Proc 2008 Intl Conf on Content-Based Image & Video Retrieval*.