Learning Graphic Design Skills on the Web: Challenges in Locating, Understanding, and Employing External Help

Juho Kim[1], Benjamin Malley[2], Joel Brandt[3], Mira Dontcheva[3], Diana Joseph[4], Krzysztof Z. Gajos[5], Robert C. Miller[1]

[1]MIT CSAIL, Cambridge, MA 02139, {juhokim, rcm}@mit.edu
[2]Univ. of Michigan, Ann Arbor, MI 48109, malleyb@umich.edu
[3]Adv. Tech. Labs, [4]Learning Resources, Adobe Systems, San Francisco, CA 94103, {joel.brandt, mirad, dmjoseph}@adobe.com
[5]SEAS Harvard Univ., Cambridge, MA 02138, kgajos@eecs.harvard.edu

ABSTRACT

When people engaged in graphic design tasks have difficulty solving problems by themselves, they turn to external help resources such as the web, help documents, or nearby colleagues. The web has expanded the amount of diversity of help options available, both when they *learn on their own* and when they *learn with others*. We present findings from two laboratory studies designed to better understand learning challenges in locating, understanding, and employing external help in the context of graphic design tasks. The first study investigates learning on one's own by searching the web for information. We find that participants struggled to formulate accurate queries, failed to recognize appropriate webpages, and had difficulty in transferring knowledge from the web content to the task at hand. The second study focuses on learning with others by connecting with remote teachers for synchronous help. We observe that these teacher-learner pairs faced difficulties in building and maintaining shared context, and in managing the cost of synchronous social interaction. The findings contribute to a comprehensive understanding of learning with online resources for graphic design tasks, and suggest opportunities for better learning environments.

Author Keywords

graphic design, online learning, web search, synchronous help

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces—*graphic design*

General Terms

Experimentation, Design, Human Factors.

INTRODUCTION

An increasing number of professional and amateur designers are learning graphical design software to create websites, retouch photos, and merge images. While trial-and-error is a popular learning strategy for many [25], the complexity of modern design software can easily frustrate users [4, 20] and exploration can trap them into acute dead-ends or local optima [16]. This is why many users combine exploration with external help, learning on one's own using prepared resources and learning with others [29].

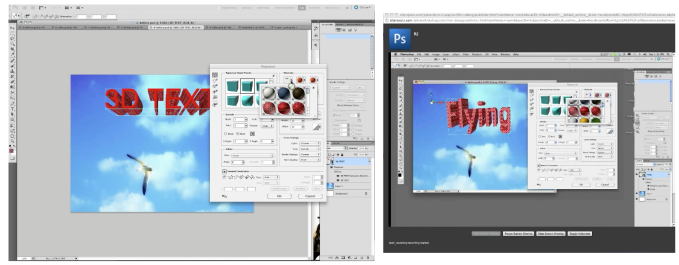

Figure 1. In the synchronous help condition (Study 2), both the learner and teacher had a dual-monitor setup, where the main monitor had Photoshop open (left) and the other had the prototype system we built for synchronous help (right). In this screenshot, the learner is following along the teacher by watching at the teacher's screen displayed on the right monitor.

The web has expanded the amount and diversity of help options available, affecting the way users find, access, and learn from external help resources. Where they used to rely on manuals and books for training on complex software applications, today's learners increasingly use the web to find tutorials and other learning resources [35]. The web is also used to seek direct help from others. Where previously learners had to ask a colleague or expert for help in person, today they can post questions on social Q&A sites and user forums [12]. This paper focuses on two learning approaches the web has made possible: *learning on one's own* by searching the web for information and *learning with others* by connecting with remote experts for synchronous help. Our goal is to better understand how to design learning environments that support these learning approaches on the web.

Searching the Web for Information

There is a wealth of free learning content available on the web. Users can search and browse through a repository of readily available learning materials. Although the web contains many images and videos, textual descriptions govern the indexing, retrieval, and display of images and videos. In order to search this space, users have to formulate a text query, even when their tasks are primarily visual. For example, how would a user approach finding a tutorial on how to select a good text color for a particular background image? Even when the user knows exactly how to approach a problem and just needs to know which tool to use, a lack of vocabulary can be paralyzing [8]. For example, to learn to resize or rotate an image in Photoshop, users have to recognize that the correct search term is *transform*. And to combine multiple images, users have to think in terms of *layers*. Even if users successfully navigate the text-based search interface, they face another challenge of translating the informational content to the task at hand.

Connecting with Remote Experts for Synchronous Help

A growing number of online user communities and Q&A sites make it easier to get help from others [11]. Despite the opportunity to learn from many experts, connecting with the *right* expert for help is challenging, especially if the goal is to provide synchronous help. Expert location attempts to identify the right expert for the task, and various approaches have been introduced to address the problem [33, 15, 23]. Even if an expert is found, motivating that expert to help remains a challenge. Many online communities such as StackOverflow embed leader boards or points systems to promote participation. With synchronous Q&A, there are issues associated with interruption costs and expert availability at question time [37]. Supporting interactions *during* synchronous help is also important, where building a common ground [5] and managing social interactions emerge as central problems. This paper explores learning challenges during synchronous communication in particular.

This research addresses the question of how to support learning graphic design software on the web. We specifically focus on learning on one's own by searching for learning materials and learning with others by connecting with remote experts. After discussing related work, we present two laboratory studies designed to better understand the learning processes, interactions, and difficulties users experience with visual tasks. The first investigates web search, while the second observes the experience of getting synchronous help from an expert. After reporting results from each study, we compare the two learning approaches for a deeper analysis. Our analysis focuses on identifying learning challenges in locating, understanding, and employing external help. The findings reveal the relative strengths and weaknesses of each approach, and more importantly, contribute to a comprehensive understanding of learning with online resources for graphic design tasks. We discuss the design implications of our studies towards future learning environments that support learning graphic design skills on the web, and briefly introduce a platform designed to realize the implications.

The key contributions of this paper are:

- An in-depth illustration of the learning process involving web search and synchronous help
- Identification of learning challenges in web search: missing vocabulary, failure to recognize appropriate content, context mismatch
- Identification of learning challenges in synchronous help: building and maintaing shared context, and managing the cost of synchronous social interaction
- Design implications for online learning platforms that support acquiring graphic design skills

RELATED WORK

This paper builds on three bodies of related work, which align with Rieman's categorization of software users' learning approaches: trial-and-error, consulting a manual, or asking for help [29]. One modification we make to the original classification is that in this paper we broaden the scope of consulting a manual to include web search. As the web has become a major source of learning materials, consulting a manual is complemented in large parts by searching for relevant web documents.

Trial-and-error: exploratory learning

Trial-and-error, or exploratory learning, is a popular learning strategy by many learners [25]. An observational study of interactive application users reports that task-oriented exploration is often the preferred method but users combine other learning approaches involving external help [29]. This finding relates to our motivation that exploration alone cannot easily accomplish the goal and looking for external help on the web is common in task-based learning for complex graphic design software. Because multiple learning approaches are often applied together, it is important to understand the comparative aspects of exploratory learning against external help. Comparative studies between exploration- and instruction-based techniques report that exploration yields higher performance [3] or more success in transferring procedures to a novel device [18]. These results suggest that learners accessing instruction-based tutorials on the web might find it difficult to transfer information on the web to the task at hand.

Searching for Learning Materials on the Web

Information seeking behavior is an active research topic in library and information science [24]. Pirolli and Card [26] propose an information foraging model derived from ecology and show that information seekers use similar strategies as food foragers. Just as predators find their prey by following clues about their location, so do information seekers follow hyperlinks, or *information scent*, in order to locate relevant information. Search engines have extensive algorithms for providing the best possible information scent for a given search result. In addition to the webpage title and url, they also display a snippet that is meant to describe why a webpage may be appropriate. Unfortunately snippets are typically text-focused and do not incorporate any visual information about a webpage. This presentation may have a larger effect on designers as compared to other populations as their tasks are visual in nature. Recent findings in search behavior show that when search becomes difficult, users start to formulate more diverse queries, use advanced operators, and spend longer time on the search results page [1]. Research in visual thumbnails [34] offers a possible approach to improving the information scent of search results for graphic designers.

Searching the web for learning has been studied in the context of programming tasks [2, 6] and the practice of craft [35]. Research shows that programmers pervasively use the web to find example code [2]. Example code helps reduce cognitive load during programming. Similarly, designers may benefit from considering example designs on the web without regard to their technical difficulty [14]. Torrey *et al.* find that people looking for help with craft projects regularly use the web but struggle in a number of ways. First, they have a hard time using keyword search because they don't always know the correct terminology. They are able to describe a tool or technique but don't know its technical term. Second, a tutorial may use tools or materials that differ from those the learner has at hand. Finally, the physical mechanism of certain techniques such as with knitting can be difficult to describe textually and even visually. Designers face some of the same challenges: they lack the necessary technical terminology to perform effective keyword queries, and the specifics of the visual tasks they want to perform may differ from those depicted in an online tutorial.

Learning with Others Synchronously and Asynchronously

Getting help from other individuals on the web has become a common learning approach. Education and communication researchers have studied the benefits and limitations of two modes of e-learning: synchronous and asynchronous [32, 17]. The contribution of this research on synchronous learning is that we address the affordances of the underlying communication systems and learning issues specific to visual tasks.

Although this paper focuses on synchronous learning, understanding how it differs from asynchronous learning can offer a comprehensive perspective. Asynchronous learning supports learners' relations with teachers even when both are not online simultaneously. Email, discussion boards, online courses, and user forums enable flexible e-learning [17], with relaxed time constraints. Research has shown that learners' information processing ability is enhanced with asynchronous learning [30]. The downside includes higher cognitive load and interruption [38] and the feeling of isolation [13]. Synchronous learning, on the other hand, requires that everyone participating in the communication be present at the same time. Personal tutoring, asking colleagues, screen sharing, video conferencing, or chatrooms all need someone else available at the time of learning. Synchronous learning is in-context and participatory [17], while enhancing learner's psychological arousal [22], motivation, and commitment [30]. However, expert availability is a scarce resource, often resorting to long waits or higher monetary cost to gain access to the right experts just-in-time [37].

STUDY 1: WEB SEARCH IN GRAPHICAL DESIGN TASKS

Our first study investigates learners' use of web search in graphical tasks to better understand the difficulties visual workers face when seeking information. The study explores the following research questions:

- How do learners approach looking for help on the web and do they find what they are looking for?
- How does the presentation of content affect the learning process?
- How do learners translate web learning resources into meaningful next steps?

Methodology

Participants

We invited eight participants (five male and three female) ages 19 to 58 to participate in 1 to 1.5 hour observation sessions. The goal was to recruit people with a wide range of expertise. We screened participants on profession to ensure we have both professionals and amateurs, and avoided relying on self-reported expertise. The self-reported measures, however, correlated with their profession. Four were professional graphic designers of varying levels of experience. Of the other four, three actively used graphic design software to communicate ideas with graphics as a non-primary function of their job. All participants were compensated $100 for their time. We refer to all the participants as *learners* regardless of their skill level.

Procedure

The study included an interview component and a task-based component. The learners performed the tasks on a laboratory computer and had the option of using a Windows or Macintosh machine. All tasks were performed in the application Adobe Photoshop. The learners were told to approach each task as they would at home or work and use resources they would use typically, including the web. The study moderator asked the learners to think aloud as much as possible and would occasionally ask a learner to explain certain actions. The study moderator would also prompt the learner when he or she was unsure of how to proceed (e.g., "how would you approach this at work?"). The study was recorded through screen capture software that records both the computer screen and any audio in the room. The tasks did not have a specified time limit and were concluded only when the learner had either completed the task or given up.

In the introductory interview we asked learners to describe their typical design tasks and their web use in the context of design. The learners were also asked to describe their processes for learning new tools or features in graphic design software. Finally, they were asked whether they had any favorite websites, whether they regularly visited or participated in forums and whether they preferred video or text tutorials.

Tasks

Each learner was given a warm-up and one or two main tasks to complete. The warm-up task was designed to be relatively easy and we expected that many learners would be able to successfully complete it. For this task, the learners were asked to recreate a target image by merging two input images.

The main study tasks were designed to be challenging, and we did not expect most learners to complete them without an external aid. For Task 1, learners were asked to scale and rotate a 3D object (in this case a soda can) to make it fit with a background image (see Figure 2). In order to ensure that learners would find this task challenging, they were pre-selected based on their lack of experience with Photoshop's 3D tools. The learners that were able to successfully complete Task 1 were given Task 2 that required adding a text label to the 3D soda can. For all tasks the study moderator described the task without using application-specific language that would aid the learners. For example, the learners were told that the soda can was a 3D object but the words "scale" or "transform" were not used.

Results

All of the learners used the web at some point during the session. Figure 3 visualizes the learning process over time. On average, each learner issued 6.5 queries and visited 7.9 web pages, 2.6 of which contained video. The learners spent an average of 36.5% of the study time on the web looking for help. This number was slightly higher for the amateurs (39.5%), than the professionals (33.9%). The professionals spent more time looking at text resources (46.9% of their time on the web) than video (26.5%). The amateurs spent more

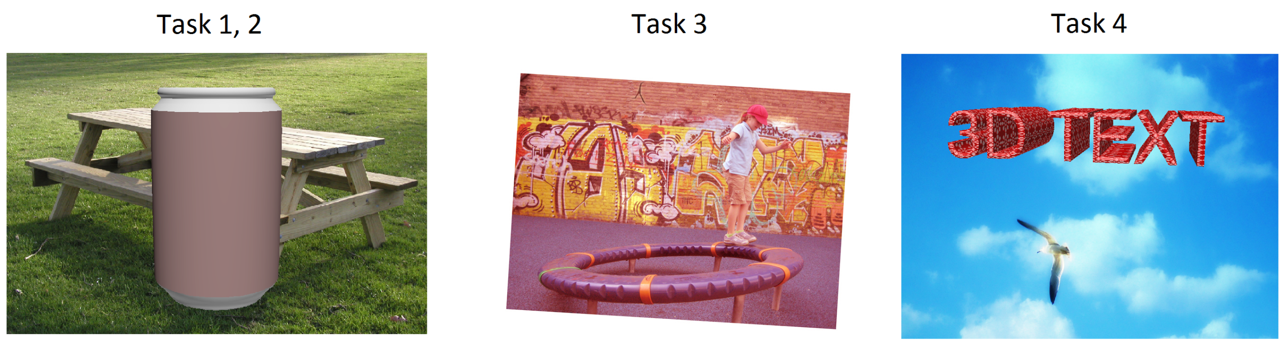

Figure 2. In Task 1, the learners were asked to make the 3D soda can look like it was sitting on the picnic table. Task 2 asked to place the label "soda" on the 3D can. Task 1 and 2 were used in both studies, while Task 3 and 4 were used only in Study 2. Task 3 was to improve the given photo with color correction techniques. In Task 4, the learners were asked to add a 3D-shaped text with customized texture and structure to the background sky image.

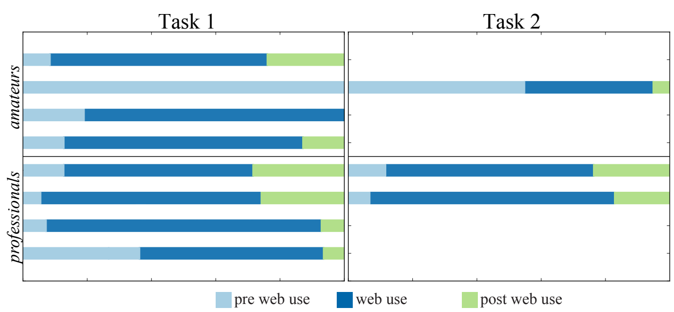

Figure 3. All learners used the web during Study 1. This graph shows the learners' relative time usage during each task. The white areas show that the learner did not work on the task. They began the task by trial-and-error, spent significant time searching for web learning content, and finished up applying the acquired knowledge to the task at hand. Note one amateur who successfully solved Task 1 alone and spent much more time exploring by himself on Task 2.

time with video (35%) than text (28.7%). However, the relative amount of time spent with either text or video was highly variable within each group and appeared more closely tied to personal preference than an indicator of expertise.

All the learners except for one amateur used the web during Task 1, and these learners spent on average 44.5% (min: 32% max: 58%) on the web looking for help. 34% of this time was spent searching and looking at the search results page, 37% of the time was spent looking through text web pages, and 29% of the time was spent looking through videos. The amateur who did not use the web discovered the 3D panel in Photoshop by accident while explaining to the moderator what was causing her difficulty.

Only three of the learners (two professionals and one amateur) were able to complete Task 1 and move on to Task 2. All three learners used the web for Task 2 and spent on average 42.6% on web resources. No learner was able to complete Task 2, despite finding appropriate web resources.

STUDY 2: SYNCHRONOUS LEARNING IN GRAPHICAL DESIGN TASKS

In Study 1, we observed learning approaches and challenges involving web search. With a different learning option available, how would the learning process change? We conducted a second laboratory study where participants were asked to work on graphical design tasks with a teacher available online to provide synchronous help, connected via screen sharing and voice chat. The study addresses the following research questions:

- When and why do learners turn to synchronous help?
- When does synchronous learning win or lose over working with the web alone?
- When learners get synchronous help, what does their learning process look like?
- What screen sharing modalities and interactions emerge from synchronous tasks?

Methodology

The setup of the experiment was identical to the previous study. We describe any nontrivial difference in this section.

Participants

We invited 8 participants (3 male and 5 female) ages 20 to 42 to participate in 1 hour observation sessions. Consistent with the first study, we had a mix of expertise and experience: four professional designers and four amateurs. The professional designers included graphic designer, urban designer, and design major student. The amateur designers included a documentary director, a software engineer, and an administrative assistant. The professionals reported much more frequent usage of Photoshop but none had prior experience with the 3D features, which the tasks require them to use. Each participant was paid $75 for their participation in the study. We refer to all the participants as *learners* regardless of their skill level to distinguish their role from *teachers* who give synchronous help in the study.

Additionally, we recruited 4 teachers (2 male and 2 female) to serve as just-in-time helpers for the learners. The teachers were intermediate to advanced Photoshop users, and two of them were professional designers. Each teacher was trained before actual sessions on the 3D Tools in Photoshop to make sure they knew how to approach the tasks. Teachers were matched with learners of same gender, and each joined 1 to 3 learner sessions. Each learner only worked with a single teacher throughout the entire session.

Apparatus

We built a prototype system as shown in Figure 1 to provide a synchronous learning experience to the learners. We chose to build one rather than use an existing system because we wanted to tweak the interface and features while iterating on the study design. With the prototype, users can use start and stop buttons to begin and end sharing their screens. The main portion of the screen asset is assigned for viewing the partner's screen when available. When one side decides to share screen, the partner automatically starts seeing the other side's screen on which Photoshop is open. The system supports both text and voice chat, but contains no shared pointer and users are not able to take control of their partner's screen. Thus all references to objects on a shared screen occur through voice and text channels.

Procedure

In addition to using any help resources and web search, the learners had an option to get teacher's help on the tasks. Upon pressing the call button on the screen, the learner and teacher were connected with voice chat. Screen sharing was not enabled by default, and each side had an option to share their screen as desired. This configuration allows four different screen sharing modalities (nothing shared, learner shared, teacher shared, and both shared). The study moderator did not require the learners to use a particular problem solving method, although in a warm-up task the learner was required to test voice chat and screen sharing.

Tasks

After completing a simple warm-up task, each learner was given two to four tasks to complete. The warm-up task used in Study 2 was different from the one used in Study 1: it was a UI location problem that is simple once familiarized but few people know how to do. This led the learner to naturally engage with the teacher, experiencing synchronous help before the actual tasks. The first two main tasks were the 3D tasks identical to Task 1 and 2 in the first study: scaling and rotating a 3D object, and adding a text label to the object. Task 3 and 4 were only presented when enough task time remained. We added them because we expected learners to complete the tasks faster than Study 1 with teacher's help available. For Task 3, we wanted to see how people work when confronted with a familiar, open-ended problem. Because all of our learners reported at least some familiarity with color correction, we asked them to use this technique to improve a photo (Figure 2). Task 4 was a multi-step, more complex 3D task that required more than locating the right UI element (Figure 2).

Results

Overall, each learner successfully completed 3.6 tasks on average, and no task was incomplete or stopped by the moderator. Six of the learners completed all four tasks, one completed three of the tasks, and one completed two.

	Task1	Task2	Task4	Task1+2+4
Exploration	28.71%	34.74%	39.99%	34.61%
Sync Help	65.09%	61.51%	54.38%	60.38%
Wrapping up	6.20%	3.75%	5.63%	5.01%

Figure 4. All learners worked with synchronous help during Study 2. This table shows the learners' relative time usage during each task. They started each task by exploring the interface, connected with the teacher for help, and finished up applying the acquired knowledge to the task at hand. Note the increasing time spent on synchronous help as the task progresses (more complex, more comfortable with the teachers).

All of the learners used the teacher's help at some point during their tasks. 23 of 29 tasks (79.3%) included synchronous communication between a learner and teacher. All learners used synchronous help for Task 1, 2, and 4. For Task 3, the color correction task with which everyone had prior experience, only one learner asked for teacher's help. Every task except for one had exactly one call between the learner and teacher. Learner 4 called the teacher again after getting stuck after the first call. No teacher or learner used text chat at any point during the sessions. This is likely because every synchronous help instance started with voice chat by default, eliminating the need for text chat.

Search on the web or in-application help menu was used in only 6 of 29 tasks (20.7%). Five learners searched for learning materials at least once. Results also seem to suggest that search never led the learners to the final answer: in all six tasks search was followed by synchronous help, and no search took place *after* a synchronous conversation.

In all 23 tasks with synchronous help, the learners chose to share their screen, but the teacher's screen was shared only 10 of 23 times (43.5%). The teachers seem to be in charge of determining whether to share their screen: two teachers shared their screen in every task while the other two never shared their screen. No learner explicitly asked the teachers to share their screen.

Compared to the average of 36.5% of the task time spent on the web in Study 1, the learners in Study 2 spent 60.38% of their time talking to the teachers. We remove data from Task 3 because there was only one participant who used synchronous help. Task 1, 2, and 4 are all 3D-related tasks that progressively get more complicated. One trend to note is that as the task gets more complex, the learners' time spent on exploration before talking to the teacher increases. But the time spent with the teachers decreases (See Figure 4).

COMPARING LEARNING PROCESS & APPROACHES

This section provides an in-depth analysis of the learning process from the studies. We discuss interesting learning patterns the learners exhibited and attempt to explain why such patterns emerged in each condition. For the rest of the discussion we refer to teacher-learner pairs as *pairs*.

Deciding to Get External Help

In both studies, most learners started by exploring the application by themselves to find a solution. For the 3D tasks, most failed to find an answer solely by trial-and-error. In many ways, the learners' decisions to turn to external resources indicated that they had "given up" on the interface. At what point do they decide to turn to external help, and why?

How and why questions indicate a need for external help

The learners decided to get help when they were feeling *stuck*. We observed two different cases. In the first case, the learner could not locate any relevant tool to apply, having no clue *how* to approach the problem. The learner in the other case tried a solution he originally thought would work, but it failed. The learner does not understand *why* the approach failed. The two different causes of frustration seem to affect the way the learners describe their problem to the teachers. In the first case, the learner often gave a problem-oriented description to the teacher, as in "I need to move and resize this can so that it can fit on the table". In the second case, the description was more approach-oriented, explaining what the learner tried and how that failed to produce a desired outcome. One learner asked, "I tried free transform on the can, but it gives me this weird error message. I don't know why transform fails here. The can's already a separate layer."

The learners tended not to ask for help when they could achieve any result. In the color correction task in Study 2, locating the color adjustment panel was easy for everyone, and trial-and-error returned satisfactory results to many learners. The learners avoided teacher's help regardless of the quality of their results. The only learner who contacted the teacher in this task also successfully had located and tried a few color adjustment tools. Nonetheless he decided to call the teacher saying "I want to get his [the teacher's] subjective feedback. Also, I want to see how an expert does color correction."

Describing the Problem

Once the learners decided to use either web search or teacher's help, the first step was to describe the problem in a form suitable to the help medium. In search this was formulating a search query to retrieve results containing useful information, and in synchronous learning this was verbally describing the problem and goal to the teacher. The learners had difficulty translating their visual task into written or spoken language, running into what is known as the *vocabulary problem* [8]. They had to revise queries multiple times or backtrack the logged history of all operations until the helper system or person could return useful suggestions.

Web search: formulating a query

The learners in both studies often used multiple search engines and, in particular, transitioned from using a general search engine (e.g., google.com) to a domain specific one (e.g., adobe.com). Of the thirteen who used search, eight began at google.com, four started with an in-application help client, which displays web resources, and one began by searching for video on YouTube.

The average query length of the queries we observed was 4.9 terms (min: 1, max: 9) and on average learners formulated 4.5 queries (min: 1, max: 9) per task. Although our sample is small, this behavior is different from the reported general web search behavior of an average 1.94 queries per session with query lengths of 2.93 terms [19].

Learner 4 - professional	
Task 1	
photoshop cs5 transforming a 3d object	google
photoshop cs5 scaling a 3d object	google
Task 2	
3d object, type	adobe
adding type to a 3d object	adobe
adding text to a 3d object	adobe
Learner 6 - amateur	
Task 2	
how to shrink 3d image photoshop cs5	google
how to shrink 3d image photoshop cs5 youtube	google
shrink 3d image photoshop cs5 youtube	google
shrinking 3d image photoshop cs 5youtube	google
reducing 3d image size photoshop cs5 youtube	google
shrink 3d image size photoshop cs5 youtube	google
shrink 3d image size photoshop cs5	youtube
3d image photoshop cs5	youtube

Table 1. These are the queries some of the learners made during Study 1. The amateurs had a harder time finding useful vocabulary for the task at hand and evaluating the available resources.

The observed high query length was due to several factors. First, learners included many function words such as "in" and "a" in their queries. Second, they frequently appended the application name and version to the query (i.e., "photoshop cs5"). Finally, because the learners faced the vocabulary problem, in addition to changing keywords (e.g., changing *shrinking* to *reducing* as seen with Learner 6 in Table 1), they often added new keywords to their searches, creating longer queries.

Overall, the average number of terms per query was much higher for general search engines (approx. 5 terms) than for in-domain searches (2 terms). This is primarily due to learners dropping the application context from the query, as can be seen with Learner 4 in Table 1. In many cases, dropping the application context from their searches led to learners arriving at content written for a different application (e.g., a Flash 3D tutorial). Further, many learners failed to realize this discrepancy and ended up spending significant time trying to use resources designed to help users of a different application.

With web search, we noted differences between the amateurs and the professionals. In contrast to previous findings on the effect of domain expertise on query length [36], the professional designers formulated shorter queries (mean: 4.1) than the amateurs (mean: 5.5). Amateur users tended to include more function words in their queries and to formulate their queries as questions, which accounts for this difference.

Synchronous: building shared context

We observed patterns the pairs exhibited to reach the shared understanding of problem status. Building and maintaining shared context helped the teachers to better understand, analyze, and resolve the problem.

When the learners called the teachers for help, the teachers' first question in most conversations was "What have you done so far?". Sharing the learner's screen took place in this problem description phase to aid the communication. Having a shared screen made the communication for the pairs easier, resulting in less verbal description and the faster establishment of a shared context. One learner said, "[with screensharing,] you don't have to waste time explaining everything". On average, the learners started sharing their screen after 11 seconds (min: 0, max: 61) the synchronous communication started. The first-time screen sharing was often accompanied by introduction or asking for permission (e.g., "is it okay for me to share my screen with you?").

The pairs reached the agreement to share the learner's screen faster in later tasks than in earlier tasks. On average, it took 23, 8, and 1 seconds for the learner's screen to be shared in Task 1, 2, and 4, respectively. Task 3 was left out because only one learner used synchronous help. The learners shared their screen even before they started talking in 14 of 23 (60.9%) synchronous tasks. Because the study design required each learner to work with a single teacher for all the tasks, the pairs could skip the social interactions of introducing and asking for permission to share screen in later tasks.

A problem arose when the learners omitted crucial information in their problem description. One learner, while working on resizing a 3D object, converted the object into 2D before talking to the teacher. There was a dialog that asked her if she wanted to rasterize the 3D layer, and she clicked "Yes" not knowing the effect of such action. When she asked the teacher to help out, the teacher noticed that her solutions did not work for some reason. She asked, "did you do anything else to the file?", but the learner replied, "No".

The above example shows that the vocabulary problem persists in the synchronous case, but the pairs applied various solutions to this problem, from navigating the operation history together until the teacher was confident that the file was at a clean state to closing the working file completely and reopening it to start over.

The learners often asked follow-up questions to the teachers. The questions were sometimes about the specific tool applied at the moment (e.g., "Can I wrap an image on the soda can, not just text?"), but many general or tangential Photoshop questions were asked on-the-fly (e.g., "Is there a shortcut for opening the history panel?"). Once a shared context is built, the cost of asking additional questions is very low. One learner noted, "I was able to ask questions and get immediate response/answers." This was not observed during web search, probably because the learner had to formulate a separate query from scratch.

Locating a Solution

Once the learners described their problem, they either browsed through search results or worked with the teachers to find a solution. This phase featured evaluating multiple options until they found a reasonable solution to apply.

There is a trade-off between the diversity of answers the learner explored and the depth in understanding why one approach works or not. In Study 2 the learner talks to a single teacher, and therefore learns only one approach, even if it is not the *best* approach. The teacher's expertise in Photoshop and teaching therefore has much influence on the learning experience. It is also true that more in-depth discussion about a solution takes place in synchronous conversations (why one approach works and the other does fails). With web search, we observed multiple instances of information assimilation [7], which involves piecing together different answers. The learners quickly skimmed through multiple answers and obtained a higher-level view.

Web search: foraging for text and video content

As in other information seeking behaviors, web search for visual tasks also included the foraging phase [26]. More than one learner indicated that they used the presentation and aesthetic attributes of a resource as a measure of its value. In regard to one tutorial site, a learner said "I've learned that usually anybody who has a web page header that's that bad is not going to be helpful." Video as well was often regarded in terms of the quality of presentation. One learner commented on on the proliferation of videos where "you have 14-year-old kids showing you how do something," noting that it was frequently not "constructive."

Several learners exhibited a strong preference for either text or video content in both the interview and task portions of the session. One learner indicated he liked the idea of video, but had found web videos to be varying in quality. This general sentiment was echoed by a number of the learners.

The observation sessions also demonstrated the poor information scent of video resources. This manifested itself in two ways. First, learners were often unable to differentiate text and video links on search results pages. One learner was surprised when a link she opened in a background tab began playing video and quickly closed it. Another learner who preferred text resources commented on the annoyance of "surprise video," e.g., accidentally opening a video link during a meeting, as a reason he preferred text-based tutorials. Second, learners often decided a particular video probably wasn't going to help them, but were reluctant to close the video or skip ahead for fear of missing something valuable.

Learners also arrived at online content for different applications including a video on importing 3D objects into Adobe After Effects, a tutorial on 3D in Flash and a page on layouts in Adobe InDesign. Further, the learners were generally not able to recognize the inappropriateness of one of these resources for the task. When asked what he discovered by watching a video on 3D in After Effects, one learner said "there has to be a simple way to do this."

Synchronous: search-like behavior

Although they were not using a search engine, the interaction between the learners and teachers resembled typical search behavior. Analogous to the foraging and query reformulation process in web search, the synchronous learners continuously updated their working image during their conversation with the teachers. The pattern appeared clearly when the teacher did not have a clear answer to the problem. The teacher suggested multiple possible solutions, and the learner applied them to the task image. This behavior is consistent with the information seeking literature, corresponding to updating search representations as new options were found [31].

In the six tasks where the learners did use web search, they accessed only text-based resources. This is starkly different from the learners in Study 1, six of whom used video content. One interpretation of this difference is that a synchronous connection replaced the need for video content.

Applying a Solution

The phases of locating and applying a solution are tightly coupled, often forming a loop until the final solution is reached. This phenomenon is consistent with social search literature [7], which noted the foraging and sensemaking loops in informational search efforts. In this phase, the learners attempted to apply a located solution to their working image, similar in fashion to sensemaking [27].

A notable difference between the two studies in the applying phase is the speed of loop iteration. The synchronous learners showed much faster locating-applying cycles than the web searchers. This was partly because the teacher could provide realtime feedback on the learner's work. Sometimes the teacher stopped the learner from applying a wrong solution.

Web search: transferring knowledge from the web

The learners were unsuccessful in translating the web resources to the task at hand. This was as true of content that was relevant to the task as it was of content that was not relevant or appropriate. While six learners located an Adobe webpage on Photoshop's 3D object and camera tools, only one was able to use this information to complete the task. The remaining five learners could not make use of the content because their application context was not identical [9] and they couldn't figure out why certain menu options were not available. One learner found a video that showed her how to accomplish the task but with a different 3D object, a cone instead of a soda can. Instead of trying to apply the technique to her 3D object, she stopped working on the task and started following the tutorial.

Synchronous: watching or listening

The process of transferring the teacher's knowledge varied significantly based on screen sharing modality. In all tasks, the learners ended up sharing their screen with the teacher. But the reverse was not true: the teachers only shared their screen in 10 of 23 tasks (43.5%). We observed two configurations: both screen shared and only learner's screen shared.

When both screens were shared, the learner frequently moved back and forth between watching the teacher's screen and working on their solution. Access to the teacher's working process served as a visual learning aid: one learner said, "you could see in real time the results of each step the expert took." Some learners lost focus on the demonstration and struggled with following along. Other learners chose to watch the teacher demonstrate until certain point before they worked on their image. The synchronous task length was slightly shorter when both screens were shared (mean=345 sec, σ=204) than when only the learner's screen was shared (mean=366 sec, σ=192).

When only the learner's screen was shared, the solution finding and applying phases were happening almost at the same time. The learners were listening to the teacher's suggestions and applied them instantly. One repeated difficulty the pairs faced was locating a UI element on the learner's screen. Because the teacher had no way to demonstrate from their own screen or control the learner's screen, the teacher had to verbally reference visual elements. In helping his learner locate the 3D tool button, one teacher said, "to the right, right, right, now too much. To the left, right there".

The teachers in the study had a clear preference for a specific modality. Two of the teachers shared their screen in all tasks, and the other two never shared their screen, regardless of who the learners were. We attribute this divide to the preferred teaching style that the teachers have. Some teachers were confident in demonstrating while others were more comfortable with talking.

DISCUSSION

Limitations of the Studies

Despite the various learning interactions observed during the sessions, our two studies face limitations due to their simplified laboratory settings.

In the synchronous study, the teachers were pre-trained with relevant tools and aware of the tasks the learners were given. They also had more than one learner sessions, with the possibility of learning effect. Our informal observation suggests that learning effect did not occur, as the there was no notable difference in the learners' time usage in first-time students and others. This resulted in much faster establishment of shared context, faster and higher quality solutions suggested, and probably more enjoyable communication and learning experiences for the learners. Despite the confound, however, most pairs still struggled with nontrivial learning challenges.

Our setting ignored any overhead in finding the right teacher and motivating the teacher to help out. Teacher training was a simulation of manual expert location. Although this was by design in order to limit the scope of our observations to actual interactions between learners and teachers, the full learning experience would include more social and time cost for both sides. Again, despite the simplification, the learners faced difficulty dealing with social interactions. We plan to address these issues as we iterate on the design of our learning platform, which we discuss below.

We also note that the studies only look at two specific learning approaches. The design space for learning on the web is much larger: Search can be synchronous, algorithms can suggest queries to aid text-oriented search, etc. The two learning methods we picked, however, suggest interesting insights that could be applied to other methods.

Design Implications

The findings from the studies highlight the learning challenges users of graphical design software face. We outline a few design implications for the next generation learning platforms to support better integration with external help.

Support the visual language of design

Because a majority of learning tasks in the graphic design domain are visual in nature, the learning environment should allow users to visually communicate. The vocabulary problem could seriously cripple the communication between users. One suggestion is to let images be the primary unit of communication, complemented with textual descriptions. Supporting visual comparisons between different images or information search using graphic design primitives, such as fonts, color gradients, or images, can significantly enhance learning.

There are interaction design issues around supporting visual communication during synchronous sessions. Because the pairs in Study 2 did not have a direct way to refer to their partners screen, locating UI components required tedious communication. The learning interface can have a shared pointer where each party can visually point to the other side's screen or visualized keyboard usage to capture shortcuts.

Encourage learning-by-doing

Make it easy to move back and forth between the application and help resources. Recent research on interactive tutorials presents new opportunities to bridge the gap between the two [10, 21, 28, 39]. Graphic design applications themselves can allow designers to easily apply solutions presented in external materials, in a context-aware manner to deal with version or tool mismatches.

In synchronous sessions, there are many screen sharing solutions that support taking over the partner's screen. This might not be useful if the goal is to encourage learning-by-doing. One learner in the study said, "It was even better than giving over my screen to another person, because I was in charge of tools and I had to learn by doing."

Support various learning interactions

Our findings suggest that each learning method has its own strengths and weaknesses. Learners will benefit from platforms that support flexible transitions between different approaches. One example is to embed chat inside an application or tutorials to foster synchronous discussion while exploring or reading materials.

Flexibility in user roles and screen sharing modalities for synchronous sessions is also desired. Because it is very hard to accurately model expertise in complex graphical applications, defining a global expert is also difficult. A more flexible role definition than mere expert-novice would motivate learners to help each other, sharing their learning experiences with others than being passive receivers of information. In Study 2, we observed cases where even professional designers learned from amateur teachers. Also, even the small number of pairs used two different screen sharing modalities in Study 2. More user control on how to communicate with each other might better support learning.

Lower the overhead of synchronous sessions

We noticed that the social and physical cost of a synchronous session is high. A good synchronous learning platform should make it simple to start and stop a session. The cost on the learner side comes from looking for the right teacher and socially interacting with the teacher. The learner might be discouraged by the feeling that he is bothering the teacher with trivial questions. The teacher faces different challenges in allocating their scarce resource to the right learners. One idea might be to impose a time limit in each session, so that the teachers know how much time they will spend in the session even before joining one. Brief help from teachers might be useful, based on the observation that many questions simply require locating the right tool or pointing to the right next step. Teachers can translate users' problems into domain-specific vocabulary and link the most useful learning resources in synchronous sessions.

A Community of Learners Exchanging Synchronous Help

Based on the results of the studies and design implications presented, we have begun to build a learning platform that provides flexible and accessible synchronous learning to anyone. It aims to support a community of learners exchanging just-in-time help for graphic design tasks. There are few communities in service supporting synchronous learning for visual design tasks.

To support the visual language, the system uses images as primary unit of communication. The system encourages browsing through user-posted images in a grid layout, with text descriptions supplementing each image. Pairs start a synchronous session upon joining a virtual room and agreeing to work together. To lower the barrier to joining in synchronous sessions, we implement two features. First, a lobby chatroom allows pairs to clarify their objective before starting a session. This prevents the vocabulary problem from misguiding the teacher and helps learners revise their description. Second, each session is limited to 5 minutes to minimize the teacher's time commitment and make exiting sessions socially appropriate. The system also supports other learning methods: 1) learners can choose to record a session so that they or others can review the recording later and discuss, 2) the system can recommend relevant web tutorials and forum entries for each session to further the learning, 3) teachers can add video replies to learner questions asynchronously by self-recording their solution, and 4) spectators can join live sessions and watch other pairs work on graphical design tasks.

We plan to address community-building, teacher motivation, and matchmaking more carefully, as our studies have not explored these in depth.

CONCLUSIONS AND FUTURE WORK

This paper presented two studies designed to better understand challenges in locating, understanding, and employing external help on the web in learning design skills. The first study looked at web search behavior, finding that learners struggle with query formulation, context mismatch, and knowledge transfer. The second study focused on synchronous learning, finding that learners struggle to maintain shared context, manage the burden of synchrony and social interactions. Based on the observations we outlined a set of design implications, and described an idea to incorporate some of the features recommended. We plan to extend the work by building the proposed system and launching it publicly.

Our studies showed learning challenges specific to the visual design domain, which is one of the main contributions, but many of the presented findings might be applicable to more general GUI learning, as users also suffer from the vocabulary problem, context mismatch, or the difficulty of finding and learning from an expert.

REFERENCES

1. Aula, A., Khan, R. M., and Guan, Z. How does search behavior change as search becomes more difficult? In *CHI '10*, ACM (New York, NY, USA, 2010), 35–44.

2. Brandt, J., Guo, P. J., Lewenstein, J., Dontcheva, M., and Klemmer, S. R. Two studies of opportunistic programming: interleaving web foraging, learning, and writing code. In *CHI '09*, ACM (New York, NY, USA, 2009), 1589–1598.

3. Carroll, J. An overview of minimalist instruction. In *System Sciences, 1990., Proceedings of the 23rd Annual Hawaii Intl' Conference on*, vol. iv (jan 1990), 210 –219 vol.4.

4. Ceaparu, I., Lazar, J., Bessiere, K., Robinson, J., and Shneiderman, B. Determining causes and severity of end-user frustration. *International Journal of Human-Computer Interaction* (2004).

5. Clark, H. H., and Brennan, S. E. Grounding in communication. In *Perspectives on socially shared cognition*, . S. D. T. L. Resnick, J. M. Levine, Ed. Washington, DC, 1991, 127–149.

6. Dorn, B., and Guzdial, M. Learning on the job: characterizing the programming knowledge and learning strategies of web designers. In *CHI '10*, ACM (New York, NY, 2010), 703–712.

7. Evans, B. M., and Chi, E. H. Towards a model of understanding social search. In *CSCW '08*, ACM (New York, NY, USA, 2008), 485–494.

8. Furnas, G. W., Landauer, T. K., Gomez, L. M., and Dumais, S. T. The vocabulary problem in human-system communication. *Comm. of the ACM 30*, 11 (1987), 964–971.

9. Gentner, D., and Markman, A. B. Structure mapping in analogy and similarity. *American Psychologist 52* (1997), 45–56.

10. Grabler, F., Agrawala, M., Li, W., Dontcheva, M., and Igarashi, T. Generating photo manipulation tutorials by demonstration. In *SIGGRAPH '09*, ACM (New York, NY, USA, 2009), 1–9.

11. Harper, F. M., Moy, D., and Konstan, J. A. Facts or friends?: distinguishing informational and conversational questions in social q&a sites. In *CHI '09*, ACM (New York, NY, USA, 2009), 759–768.

12. Harper, F. M., Raban, D., Rafaeli, S., and Konstan, J. A. Predictors of answer quality in online q&a sites. In *CHI '08*, ACM (New York, NY, USA, 2008), 865–874.

13. Haythornthwaite, C., and Kazmer, M. M. *Bringing the Internet Home: Adult Distance Learners and Their Internet, Home, and Work Worlds*. Blackwell Publishers Ltd, 2008, 429–463.

14. Herring, S. R., Chang, C.-C., Krantzler, J., and Bailey, B. P. Getting inspired!: understanding how and why examples are used in creative design practice. In *CHI '09*, ACM (New York, NY, USA, 2009), 87–96.

15. Horowitz, D., and Kamvar, S. D. The anatomy of a large-scale social search engine. In *WWW '10*, ACM (New York, NY, USA, 2010), 431–440.

16. Howes, A., and Payne, S. J. Semantic analysis during exploratory learning. In *CHI '90*, ACM (New York, NY, USA, 1990), 399–406.

17. Hrastinski, S. Asynchronous and Synchronous E-Learning. *Educause Quarterly*, 4 (2008), 51–55.

18. Kamouri, A. L., Kamouri, J., and Smith, K. H. Training by exploration: facilitating the transfer of procedural knowledge through analogical reasoning. *Int. J. Man-Mach. Stud. 24* (February 1986), 171–192.

19. Kamvar, M., Kellar, M., Patel, R., and Xu, Y. Computers and iphones and mobile phones, oh my!: a logs-based comparison of search users on different devices. In *WWW '09*, ACM (New York, NY, USA, 2009), 801–810.

20. Kaufman, L., and Weed, B. Too much of a good thing?: Identifying and resolving bloat in the user interface. In *CHI '98*, ACM (New York, NY, USA, 1998), 207–208.

21. Kelleher, C., and Pausch, R. Stencils-based tutorials: design and evaluation. In *CHI '05*, ACM (New York, NY, USA, 2005), 541–550.

22. Kock, N. Media richness or media naturalness? the evolution of our biological communication apparatus and its influence on our behavior toward e-communication tools. *IEEE Transactions on Professional Communication 48*, 2 (june 2005), 117–130.

23. Lappas, T., Liu, K., and Terzi, E. A survey of algorithms and systems for expert location in social networks. In *Social Network Data Analytics*, C. C. Aggarwal, Ed. Springer US, 2011, 215–241.

24. Marchionini, G. *Information seeking in electronic environments*. Cambridge University Press, 1997.

25. Novick, D. G., Elizalde, E., and Bean, N. Toward a more accurate view of when and how people seek help with computer applications. In *SIGDOC '07*, ACM (New York, NY, USA, 2007), 95–102.

26. Pirolli, P., and Card, S. Information foraging in information access environments. In *CHI '95*, ACM (New York, NY, USA, 1995), 51–58.

27. Pirolli, P., and Card, S. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis (2005).

28. Pongnumkul, S., Dontcheva, M., Li, W., Wang, J., Bourdev, L., Avidan, S., and Cohen, M. Pause-and-play: Automatically linking screencast video tutorials with applications. In *UIST 2011*, ACM (New York, NY, USA, 2011).

29. Rieman, J. A field study of exploratory learning strategies. *ACM ToCHI 3* (September 1996), 189–218.

30. Robert, L., and Dennis, A. Paradox of richness: a cognitive model of media choice. *Professional Communication, IEEE Transactions on 48*, 1 (march 2005), 10–21.

31. Russell, D. M., Stefik, M. J., Pirolli, P., and Card, S. K. The cost structure of sensemaking. In *INTERACT '93 and CHI '93*, ACM (New York, NY, USA, 1993), 269–276.

32. Russell, T. L. *The No Significant Difference Phenomenon*. International Distance Education Certification Center, Montgomery, AL, 2001.

33. Singley, K., Lai, J., Kuang, L., and Tang, J.-M. Bluereach: harnessing synchronous chat to support expertise sharing in a large organization. In *CHI EA '08*, ACM (New York, NY, USA, 2008), 2001–2008.

34. Teevan, J., Cutrell, E., Fisher, D., Drucker, S. M., Ramos, G., André, P., and Hu, C. Visual snippets: summarizing web pages for search and revisitation. In *CHI '09*, ACM (New York, NY, USA, 2009), 2023–2032.

35. Torrey, C., Churchill, E. F., and McDonald, D. W. Learning how: the search for craft knowledge on the internet. In *CHI '09*, ACM (New York, NY, USA, 2009), 1371–1380.

36. White, R. W., Dumais, S. T., and Teevan, J. Characterizing the influence of domain expertise on web search behavior. In *WSDM '09*, ACM (New York, NY, USA, 2009), 132–141.

37. White, R. W., Richardson, M., and Liu, Y. Effects of community size and contact rate in synchronous social q&a. In *CHI '11*, ACM (New York, NY, USA, 2011), 2837–2846.

38. Wright, P., and Lickorish, A. Menus and memory load: navigation strategies in interactive search tasks. *International Journal of Human-Computer Studies 40*, 6 (1994), 965 – 1008.

39. Yeh, T., Chang, T.-H., Xie, B., Walsh, G., Watkins, I., Wongsuphasawat, K., Huang, M., Davis, L. S., and Bederson, B. B. Creating contextual help for guis using screenshots. In *UIST 2011*, ACM (New York, NY, USA, 2011).