Supporting Graphic Designers' Online Inspiration Seeking Practices via Software Tools

By
Volodymyr Dziubak

A Thesis submitted to
the Faculty of Graduate Studies of the University of Manitoba
in partial fulfilment of the requirements of the degree of
Doctor of Philosophy

Department of Computer Science
University of Manitoba
Winnipeg

Abstract

Inspiration seeking is an integral part of a graphic designer's work: they often seek inspiration for new design ideas by exploring examples of others' designs on the Internet. Designers also seek inspiration for new ways to use complex design software, such as Photoshop and Illustrator. Driven by their intrinsic motivation to master these complex tools, graphic designers often explore online software learning materials, seeking ideas for new software techniques that would extend their arsenal. However, notwithstanding the importance of inspiration seeking, finding and managing suitable inspirational material on the Internet is a hard task that presents designers with numerous challenges.

This thesis investigates how to design and evaluate software tools that support graphic designers' online inspiration seeking practices. Chapters 3 through 5 describe the design and evaluation of three systems, Switter, Instory, and Prism, respectively, each instantiating a different approach of helping designers with inspiration-seeking related challenges. Summarizing our findings, Chapter 6 presents a set of guides towards designing and evaluating inspiration-seeking support systems for graphic design.

Acknowledgements

First and foremost, I would like to thank my adviser and mentor, Dr. Andrea Bunt. I could not begin to describe everything that I have learnt from Andrea over these years, but perhaps the most important lessons that she has taught me are to not make rushed decisions and to not forget about the big picture when solving hard problems. I am eternally grateful to Dr. Bunt for her positive attitude towards my successes and failures, for her patience, her support, as well as her encouragement and guidance.

I would also like to thank the members of my doctoral advisory committee: Dr. Pourang Irani, and Dr. Richard Hechter. Their feedback and encouragement during our annual meetings gave me assurance that I was on the right path and helped me to see this work through to the end. Additional gratitude is owed to my external thesis examiner, Dr. Wolfgang Stuerzlinger, because his feedback and constructive criticism contributed greatly to this thesis.

I am also thankful to all my HCI Lab friends who helped me to stay sane during hard times and made good times even better. Thank you, Stela, Dan, Patrick, Adnan, both Aniks, Ellie, Joel, Masa, Hayato, Michelle, Troy, Brian, Diljot and Raquel, Suneet, Shahed, Lorena, Ananta, Hayley, Brittany, Kenzie, and the Goat.

Last but not least, I want to thank my family. Than you, my beloved Denise, for reading and proof-reading countless iterations of this thesis, my papers, and other documents. Thank you for your support, for tolerating my eternal grumpiness, and for making my life better ☺. Thank you, Maguie, for reminding me to take breaks from work with your gentle Castellano “miau”. Thank you to my parents, my parents in-law, and to my brother for being there for me.

Publications

Sections of this thesis have been published in conference proceedings and journal publications, either previously or forthcoming at the time of publication. Permissions for these works to appear in this dissertation have been granted by their respective publishers. Following is a list of prior publications in which portions of this work appeared, organized by chapter.

Chapter 3

Volodymyr Dziubak, Patrick Dubois, Andrea Bunt, Michael Terry. (2016) Switter: Supporting Exploration of Software Learning Materials on Social Media. ACM SIGCHI Conference on Designing Interactive Systems. pp. 1209-1220.

Portions of Chapter 4 and Chapter 5

Volodymyr Dziubak, Andrea Bunt. (2018) Prism: Enhancing Graphic Designers' Visual Research with Interactive Search Trails. ACM SIGCHI Conference on Designing Interactive Systems. pp. 1283-1294.

Table of Contents

List of Figures

Chapter 1

Introduction

Inspiration-seeking can be characterized as a process of searching for something that would stimulate your creative thinking, but without upfront knowledge of what those stimuli might be [5,34,45]. Given this lack of defined goals, inspiration-seeking represents an exploratory search process [56], through which inspiration-seekers gain deeper understanding of their task and its domain, as well as generate, investigate, and refine their creative ideas [5,34,45].

Graphic design is one of the domains where inspiration-seeking plays a prominent role; designers can seek inspiration for design ideas and ways to implement these ideas into final products. A common example of an inspiration-seeking activity that helps graphic designers inform ideas for their designs is *online visual research* – the process of exploring design examples on the Internet. Graphic designers often engage in online visual research when starting a new project to help them understand their task and generate new design

ideas [5,33,62,75]. Many graphic designers also research images on the Internet in the later stages of the project, when they need inspiration for specific elements of their work, such as the shape of a logo, or a background pattern [62].

To implement their ideas into designs, graphic designers often use specialized feature-rich software programs, such as Photoshop. Informal analysis of online forums and social media suggests that professionals who work with complex feature-rich software often seek opportunities to improve and to expand their software skills, for instance, by exploring software tutorials which might be relevant to their practice. Such exploration is a type of inspiration-seeking, as through this exploratory search, professionals refine their learning goals and eventually find software tutorials that satisfy those refined goals.

Despite the prominence of inspiration-seeking in a graphic designer's workflow, finding suitable inspirational material is not trivial (e.g., [33,62,75]). Numerous search engines (e.g. Google, Bing) and specialized websites (e.g., Pinterest, YouTube) provide designers with practically limitless pools of inspirational material. However, exploring such an abundance of information without knowing exactly what you are looking for can be difficult. Also, designers' search strategies are often influenced by their prior experiences and biases, such as personal tastes and common trends in the industry, which can impact a designer's creative process [62,67]. Lastly, the high number of potential directions for exploration can make it difficult to keep track of already explored directions and of one's motivations and goals when deciding on the new paths to follow [33].

1.1 Thesis Goals and Approach

The goals of this thesis are to explore the challenges that graphic designers face when seeking inspiration on the Internet and to investigate the potential for software systems to support graphic designers in this practice. In particular, satisfying these goals involves answering the following research questions:

1) What types of challenges do graphic designers encounter when exploring inspirational material on the Internet?
2) In what ways can software systems help graphic designers with inspiration seeking? Specifically, this thesis investigates how software systems can aid designers to:
    a. explore high volumes of inspirational materials on the Internet
    b. increase awareness of the potential effects of bias on exploration strategies
    c. keep track of the exploration process
3) How would graphic designers incorporate these tools as a part of their work practices, and how can such tools impact graphic designers' inspiration-seeking goals and strategies?

Graphic designers seek inspiration to come up with design ideas and to find the best way to implement their ideas into an end product. To consider both of these types of inspiration-seeking practices, this thesis answers the aforementioned research questions from the perspective of two HCI domains: software learnability and ideation. Our investigation within the software learnability domain brings new insight about the ways in which graphic designers seek inspiration for new software techniques and how software tools can support these practices. Similarly, our research within the ideation domain contributes new knowledge about the challenges that designers face when seeking inspiration for design

Figure 1. Overview of thesis chapters and the different types of inspiration seeking investigated in this thesis.

ideas and how augmenting their practices with software tools can support designers in their inspiration hunt (see Figure 1 for the overview of thesis chapters and the different types of inspiration seeking investigated in this thesis).

To answer the first two research questions from the perspective of software learning, Chapter 3 presents results of a formative interview study with experienced graphic designers. Through this study we characterize *ad libitum exploration* – a practice of habitually exploring software tutorials shared by a design community on social media. Our findings reveal that during ad libitum exploration, graphic designers often struggle with selecting tutorials to follow, as information provided by social media is often not sufficient to determine a tutorial's inspirational value. Later in Chapter 3, we use findings from our formative study to inform the design of Switter – a system aimed to help graphic designers appraise and select inspirational tutorials on Twitter.

To answer the first two research questions from the perspective of ideation, we aggregate knowledge about graphic designers' challenges from existing literature. An extensive body

of literature about ideation in graphic design (e.g., [5,33,34,37,62,69,75]) suggests that when graphic designers conduct online visual research, they might face the challenges of design fixation and information loss (Chapters 2, 4, and 5). We leverage findings from existing research to inform the designs of Instory (Chapter 4) and of Prism (Chapter 5) – two systems which aim to help graphic designers mitigate the harmful effects of design fixation and to retain information about visual research, respectively.

To answer the third research question, we preset results from three system evaluation studies (Chapters 3, 4, and 5). Chapters 3 and 5 describe two open-ended field evaluations, of Switter and Prism respectively, and Chapter 4 describes a laboratory evaluation of Instory. Our results describe how our participants incorporated our software systems into their work practices and how this augmentation impacted their creative processes. Additionally, Chapter 6 presents a high-level analysis of our collective results and outlines a set of recommendations for how to design and evaluate inspiration-seeking support software systems.

Overall, our systems demonstrate three unique, yet complementary ways of supporting inspiration-seeking practices among graphic designers. Our results from evaluating these systems illustrate the benefits of augmenting graphic designer's work practices with software tools and bring novel insights about designers' inspiration seeking goals and strategies within the respective domains.

1.2 Summary of Contributions

Artifact Contributions. This Ph.D. work contributes three prototype systems – Switter, Instory, and Prism – each instantiating a novel approach for supporting exploration of inspirational material on the Internet.

Empirical Contributions. This thesis contributes findings from three system evaluation studies, which validate their respective approaches, highlight the strengths and the weaknesses of each system, and demonstrate a variety of scenarios for how each system can support graphic designers' inspiration seeking practices.

In addition to the aforementioned primary contributions, this Ph.D. work also provides new insights into the types of inspiration-seeking goals and strategies that emerge from augmenting graphic designers' processes with software tools. Finally, this thesis contributes a set of design guidelines aimed to help researchers to design and evaluate inspiration-seeking support systems.

1.3 Thesis Outline

The next chapter presents background and related work in the areas of software learnability and ideation. Chapters 3 to 5 describe the motivation, design, and evaluation of the three systems contributed by this thesis: Switter, Instory, and Prism, respectively. In Chapter 6 we reflect on our experience designing and evaluating the three systems and generalize our findings into a set of design guidelines for building and evaluating systems to augment graphic designers' inspiration seeking practices. Finally, Chapter 7 summarizes the thesis

work, highlights its contributions in greater detail, and discusses potential directions for future research.

Chapter 2

Background and Related Work

Research in this thesis relates to the two broad Human-Computer Interaction (HCI) research areas: software learnability and ideation. This chapter provides background knowledge and characterizes inspiration seeking activities in these two HCI domains. We also review common challenges faced by those who seek inspiration and summarize the ways in which prior work has supported inspiration seeking practices via software tools.

2.1 Overview of Software Learnability Domain

Modern design workflows often involve using specialized design software, such as Adobe Photoshop, Adobe Illustrator, Bohemian Sketch, etc. These software, similar to most modern complex software, often provide their users with hundreds of commands and features, which makes them suitable for a wide range of tasks of various complexity levels [35]. However, such versatility makes learning these software challenging [31]. For

example, with such a high number of features, it is practically impossible to know each feature's functionality. Software users can easily get overwhelmed with the visually crowded ("bloated") interfaces of these software and see no clear path for completing their tasks [38].

To learn how to accomplish their goals, users often refer to online learning resources, such as online tutorials [44,52], and support forums [1,55]. Even though these strategies allow software users to find ways to use software to achieve desired results, software users face numerous challenges, such as finding appropriate learning materials [23,26], applying techniques illustrated in tutorials [6,24,52], locating commands mentioned in tutorials in the interface [39], etc.

2.2 Task-Free Software Learning as a Form of Inspiration Seeking

Prior work in software learnability mostly considered task-driven software learning, when a need to learn feature-rich software is motivated by a necessity to complete a specific task. For example, a user might search the Internet for instructions on how to change the background of a photo in Photoshop because they want to alter their Facebook avatar picture, not because they want to learn Photoshop. With task-driven software learning, the learning itself occurs as a side-effect of completing the task (e.g., [23,57,60]).

Prior work contains evidence of a different form of software learning, which is motivated by the curiosity and general desire to improve one's software skills, rather than a specific task [17,44,70]. Rieman in his field study of software learning strategies [70] reported that some users were browsing the interface and documentation because of curiosity about the features of the software. Dorn and Guzdal in their investigation of learning practices among

web designers and web developers [17] also observed some users learning software due to curiosity and a desire to stay up-to-date with the latest industry standards. Lafreniere et al. in their work on roles and uses of web tutorials [44] found that some software learners follow software tutorials to increase their general software expertise rather than to find answers to specific questions.

From the perspective of learning theory, this task-free software learning is a form of self-regulated learning, as according to Zimmerman's definition, self-regulated learners are "metacognitively, motivationally, and behaviorally active participants in their own learning" [88]. From the perspective of software learnability, this task-free learning is a form of extended learning, as discussed by Grossman et al.'s survey on software learnability [31]. Their survey differentiates between initial learning, where novice users gain initial proficiency, and extended learning, where users' performance changes over time. Importantly, this form of learning is open-ended: users are not seeking to learn a specific thing, but rather to generally improve their abilities, potentially in ways they have never considered.

We consider this task-free curiosity-driven type of software learning to be a form of inspiration seeking, as software users who engage in task-free software learning do so without knowing beforehand what they want to learn. Instead of searching for specific answers, these people explore learning materials hoping to find inspiration for new creative ways to use software, such as learning how to use unfamiliar software features, learning new use cases for familiar software features, and learning new software techniques.

Despite the evidence of designers engaging into task-free learning, we still do not know much about how professional software users engage into such a form of inspiration seeking and what types of challenges they encounter. Chapter 3 presents an in-depth investigation of one of such practices that is common among professional graphic designers, namely a practice of habitual monitoring social media, such as Twitter, Reddit, etc., hoping to find inspiration for new creative ways to use design software (we refer to such a practice as *ad libitum exploration*).

2.3 Supporting Software Learning

While there has not been much research on supporting task-free software learning, the approach illustrated in Chapter 3 draws inspiration from prior work on supporting task-driven software learning. In what follows, we overview research in the area of supporting task-driven software learning that has been most inspirational for our work.

2.3.1 Exploration and Retrieval of Inspirational Software Learning Materials

Since accessing information on the Internet relies on search and exploration, both of these practices play a crucial role in seeking inspiration for software techniques. Prior research in software learnability has explored ways to help software learners find relevant software learning materials on the Internet. For example, the Sikuli system [87] allowed users to search an online reference of software commands using screenshots of the application's interface instead of a keyword search query. Ekstrand et al. [23] suggested supplementing keyword search queries with software commands recently used by a software user. Their system then returns search results that are most related to the user's recent software usage. The Delta system [42] improved tutorial exploration experience by visualizing software

workflows demonstrated in tutorials – sequences of commands used to achieve a specific result. However, Delta did not account for a fast-paced nature of the modern Internet.

Our research described in Chapter 3 adds to this body of work by introducing a novel interface, tailored specifically for the task-free exploration of software learning materials disseminated via fast-paced social media.

2.3.2 Appraising and Accessing Learning Materials

Considering the vast variety of learning materials on the Internet, appraising and accessing suitable learning materials can be challenging. To help software learners access relevant software learning materials, such as online tutorials and documentation snippets, prior research investigated ways to leverage knowledge about users' command activation history. For example, prior research has used a user's command activation history to retrieve and recommend relevant learning materials [26,28,57,82] and to recommend potentially relevant new software commands [40,51,60].

Researchers also have investigated how to simplify access to software learning materials by embedding them directly into the respective software application. For example, the ToolClips system [30] used tooltips for interface elements to display documentation snippets in the form of short video clips that demonstrated the usage of the respective tool. Similarly, the Intertwine system [26] showed tooltips with relevant help snippets from the recent browser history. The IP-QAT [58] and LemonAid [14] systems used interface elements for accessing contextual community troubleshooting specific to the respective tool.

The third chapter of this thesis demonstrates a novel way of using software application context for appraising and retrieving software learning materials. Specifically, our approach uses a replica of the software interface as an interactive index into a body of software learning materials.

2.3.3 Motivating Software Learning

A common challenge for task-driven software learning is motivating users to expand their existing software skills. A phenomenon known as *paradox of the active user* [11] tells us that software users do not wish to spend time learning new software techniques that would replace their existing knowledge, even if these new techniques would lead to more efficient work. Prior research has investigated various ways to motivate software learners, for example by enhancing learning with gamified dynamics [16,47,48], recommending relevant software commands [40,51,60], and by promoting reflection on learning and one's skills [42,54,59,72]

The approach we use in Chapter 3 also facilitates reflection, however for a different purpose. The task-free software learning, which is the focus of the chapter, does not require extra motivation, as it is driven by intrinsic curiosity and the desire to improve one's software skills. Therefore, instead of facilitating reflection to motivate learning, we facilitate reflection to help designers identify those areas of their software knowledge that they might want to improve.

2.3.4 Leveraging Community Contributions

Another way to support software learners is to leverage the content provided by the community. For example, prior work investigated ways to improve tutorials using valuable information found in user comments [6,19,44].

The approach that is most relevant to our work is demonstrated in the TwitApp system [49]. TwitApp uses tweets to enable software users to share their progress with other people, to follow each other's projects, and to monitor others' work in search for inspiration for new ways to use software. Similar to TwitApp, the approach described in Chapter 3 uses Twitter as a source for inspirational material. However, our approach focuses on leveraging existing rich material on Twitter, rather than on generating specialized content, as it was the case in TwitApp.

2.4 Inspiration Seeking within the Ideation Domain

Prior research in cognitive psychology characterizes design problems as ill-defined and open-ended. Initially, designers have only an incomplete and imprecise mental representation of design goals. As they iterate through phases of problem-framing and problem-solving, they gradually shape what used to be an ill-defined problem into a fully-formed problem that has become specific to them [5,22]. This explains why different designers might present different solutions to the same problem.

Ideation (a process of generating design ideas) is a core component of the design process. Herring et al. [34] identified 19 different techniques used by creative professionals to generate new ideas. Among these techniques, the authors highlight active inspiration seeking, which represents the targeted searching for images that reinforce and refine

specific design ideas, and passive inspiration seeking, which represents an exploratory browsing of images and design repositories without any concrete goal. Both these types of visual research are usually conducted via online image repositories, such as Google Images, Dribble, Pinterest, Behance, etc. From the standpoint of cognitive psychology, one of the ways to generate new ideas is by making analogies. For example, Bonnardel suggests that to come up with new ideas, designers construct a "constrained cognitive environment" by looking up knowledge from external sources (e.g., examples of other designs) and combining it with elements of their internal knowledge (e.g., personal experiences and memories) [5]. Bonnardel's ideas align with a theory of idea generation called Search for Ideas in Associative Memory (SIAM). According to SIAM, pictorial references help our brain activate knowledge records from long-term memory that would not have been accessible otherwise, which facilitates generation of new ideas [65,66,78].

Designers often collect inspirational images, for example by saving them to folders on their computers. These collections then act as visual frameworks for evaluating their design ideas and communicating them to other stakeholders, such as their colleagues and clients [33]. Designers also use these collections of inspirational design examples to reflect on their design process in the future. This reflection allows them to analyze the flow of their design workflow and to use it as a template for future projects [75].

The prior research described in this section provides background for Chapters 4 and 5 of this thesis, which focus on two specific challenges associated with this type of inspiration seeking: design fixation and information loss.

2.5 Supporting Seeking Inspiration for Design Examples

The ill-defined nature of design problems makes finding inspiration challenging. In what follows, we review prior research about a few of the most prominent challenges that graphic designers face during visual research.

2.5.1 Exploration and Retrieval of Inspirational Design Examples

As discussed in the Introduction, numerous online image repositories, such as Pinterest, Dribble, and Behance, contain endless amounts of inspirational design examples. However, such a vast pool of inspirational materials makes filtering and finding the right image difficult. For example, a traditional way of accessing the information on the Internet is via keyword searching. However, if a designer does not have a specific goal in mind, it might be difficult for them to come up with the appropriate search keywords [33,62,75]. Struggling to verbalize their ideas, graphic designers might make their search too narrow and end up not exploring a variety of topics [33]. Also, not being able to find the right set of keywords for visual research, designers often give up and refer to inspirational material that they are familiar with (that are either well-known or that they used in the past), which has been shown to hinder the creativity and novelty of their design ideas [67].

A plethora of HCI research has investigated ways to help designers to better explore and search for inspirational materials on the Internet. For example, Hashimoto and Igarashi developed a method of retrieving design examples using user-generated sketches as queries [32]. The CueFlik [25] system augments traditional image search capabilities by allowing its users to rank search results based on image visual characteristics instead of relevancy to a query. Yee et al. [86] presented a system that allowed designers to browse images and design examples by the image metadata (e.g., media, themes, location, date, shapes). Lee

et al. [46] introduced style-based exploration, allowing designers to navigate through a corpus of web page designs using stylistic features, such as colors, fonts, number of columns, and visual density. Ritchie et al. [71] moved this approach a step further by introducing techniques for style-based searching, recommendation, and filtering.

Our approaches described in Chapters 4 and 5 investigate how to further enrich a graphic designer's search and exploration experience by allowing them to reflect on their exploration process.

2.5.2 Biased Exploration and Design Fixation

Despite the benefits of online visual research, various forms of designer bias can hinder the novelty and creativity of the designer's final output. For example, prior work suggests that drawing inspiration from images which are semantically homogeneous [66], illustrate common ideas [67], are too close to the target problem [50], or which are familiar to a designer [69] might limit the designer's creativity. These forms of bias can cause design fixation – a well-researched phenomenon which describes a premature commitment to a limited set of ideas [9,10,78,81,12,15,37,43,50,66,67,69]. Researchers commonly consider design fixation to be a negative phenomenon, as it limits the novelty and creativity of final designs [67].

Prior work has also found that when experienced designers seek inspiration on the Internet, some of their search strategies increase their chances of collecting "fixating" images [62]. Prior research has also proposed ways to reduce these fixation effects, for example, by augmenting the image-collecting process with a system that displays fixating characteristics of saved images [62]. Examples of these fixating characteristics are the

familiarity of the images to the designer and similarity between saved images. The idea of such a system is built on the assumption that designers reflect on collected images during and immediately after collecting them – a process we refer to as *immediate reflection*. There is currently limited insight in the research literature, however, on whether immediate reflection takes place and, if so, designers' motivations and goals for such reflection. Even if immediate reflection takes place, it is also not clear whether designers would see benefits of a system that displays fixating characteristics of their collections.

Chapter 4 aims to fill-in the above gap by investigating whether and why graphic designers engage in immediate reflection on images collected during visual research and the potential for software tools to support such immediate reflection.

2.5.3 Information Loss in Visual Research

Prior work studying graphic designers' practices found that designers often lose information when conducting visual research [33,75]. For example, they forget to save or bookmark their search and hence have to re-do the search again later [33]. They can also lose their train of thought as they get sidetracked by alternative ideas sparked by the images that they see [33]. Saving images also does not capture the designer's rationale and motivation for saving these images – information that is as important to the graphic designers as the images themselves [75].

Chapter 5 builds on these findings, demonstrating an approach for helping graphic designers to retain some of the information lost during visual research.

2.5.4 Reflecting on the Design Process

It is common for designers to reflect on their past practices (reflection-on-action, as defined by Schön [73]). Such a reflection helps them to gauge the progression of their style and skill [54,75].

Prior research in creative practice has demonstrated a number of knowledge management systems that help designers to reflect on their design process by capturing and displaying certain aspects of their process. For example, the Freed system [61] allows designers to manually organize their image collections into spatial views and to indicate relations between them. ReflectionSpace [74] automatically organizes images found on the designer's computer using metadata from the images, such as creation date and file name.

Chapters 4 and 5 add to this research by investigating the importance of reflection on the variety of collected images, and by studying reflection on past design thought processes, respectively.

2.6 Using Search History to Support Information Seeking Tasks

Prior research has investigated browsing history to understand how people use web search to satisfy their information needs. For example, White and Drucker investigated Internet users' search trails which they defined as sequences of search queries and visits to web pages from the respective search results [84]. White and Drucker analyzed these trails to study the variability of searching behavior among Internet users. Chapters 4 and 5 extend the idea of search trails to the context of visual research.

Our approaches in Chapters 3 to 5 are also motivated by prior work that has reified Web search patterns, for example, to improve user's general searching capabilities or their searching efficiency. For instance, the Search Dashboard system [4] allows users to reflect on and improve their search strategies by allowing them to compare their search strategies to those of expert users. The SearchBar system [63] uses a hierarchical history of search topics to help users resume their tasks after an interruption and to help them re-find information in the future. Chapters 4 and 5 bring the idea of capturing search information to the domain of creativity and visual research, which, to our knowledge, has not yet been investigated.

2.7 Summary

In summary, this thesis seeks to extend existing knowledge on inspiration seeking practices in the areas of software learnability and ideation theory. In the domain of software learnability, Chapter 3 seeks to provide new insight about how and why graphic designers engage into a task-free software learning, which challenges they encounter during this form of learning, and about how we can design software systems to help graphic designers to achieve their learning goals.

In the domain of ideation, our work seeks to extend existing research on supporting graphical designers' online visual research. Specifically, Chapter 4 explores the importance of reflecting on the variety of collected inspirational material and whether such a reflection could help graphic designers to mitigate harmful effects of design fixation. Chapter 5 investigates an approach of using search trails to help graphical designers to retain their

goals and motivations for choosing the specific exploration paths when exploring inspirational images.

Chapters 3 to 5 also contribute the design of three systems. The designs of these systems add to existing research on using software interface cues for accessing and appraising learning material and on using search history information to support information seeking and knowledge management.

Chapter 3

Supporting Selection and Appraisal of Inspirational Software Learning Materials

Working with complex software, such as Adobe Photoshop, Adobe Illustrator, Bohemian Sketch, etc., graphic designers know that there are always new things to learn, such as new software commands, new software techniques, or even a completely new piece of software. Driven by the desire to improve their skills, designers actively seek out software learning resources without any specific task at hand, but rather hoping to find inspiration for new, creative, and more efficient ways to use their software [17,70]. Despite the evidence of this type of inspiration seeking taking place among graphic designers, not much is known about designers' motivations and strategies for such an exploration, and about the challenges that this exploration presents to them.

This chapter first investigates the designers' motivation and strategies for this form of inspiration seeking via an exploratory study with 11 graphic designers. Guided by our

findings, we present Switter – a tool that helps graphic designers to learn new ways to use Photoshop by aggregating and organizing the latest Photoshop tutorials posted on Twitter (Figure 2).

To evaluate Switter, we conducted a week-long field study with 11 experienced graphic designers. During this week, we asked our participants to use Switter every day to try to learn new things about Photoshop and to record their experiences in a journal. To get feedback about Switter's features and its utility for supporting inspiration seeking, we held two interviews: one in the middle of the study (day 3) and one at the end of the study (day 7). During these interviews, we asked participants to tell us about how they used Switter and whether its features helped them with seeking inspiration for new Photoshop techniques.

Research presented in this chapter has been published and presented at the ACM SIGCHI Conference on Designing Interactive Systems (DIS) in 2016 [21].

3.1 Exploratory Interview Study

Informally, we observed that some graphic designers (including the thesis author) seem to regularly check online resources (e.g. Reddit, Twitter) seeking inspiration among new, potentially relevant software learning materials, such as tutorials. To determine whether this activity was more widespread, we conducted semi-structured interviews with 11 graphic designers (6 female). We recruited participants via snowball sampling, using the thesis authors' personal contacts, and through notices posted on a university campus, Reddit, and Twitter (see Appendix B for study materials). Our participants were between 18 and 48 years of age and had at least one year of experience with image-manipulation or

graphics software (e.g., Sketch, Adobe Photoshop, Adobe Illustrator). Nine participants used the software professionally, while two used the software extensively as part of their current training. Interviews were conducted either in-person or via Skype and lasted between 30 and 45 minutes. Participants were remunerated with a $15 gift card.

In our interviews, we asked participants what, if anything, motivates them to learn about the software they regularly use. We also asked participants to describe the specific learning strategies they use and how well their current strategies support their learning objectives.

Interviews were transcribed in full. Data from the transcripts were analyzed by creating affinity diagrams using a bottom-up inductive approach [79]. From these affinity diagrams, we held joint data interpretation sessions where we extracted common themes.

3.2 Findings: Motivation for Inspiration Seeking

All our participants emphasized their desire to continually seek out inspiration for using the design software that they use in their work, as the following two quotes illustrate:

> *[I've been using Photoshop] since I left school, which was when I was 15. So, 10 years [of experience] ... I am learning all the time. You learn stuff you didn't realize you could do. (P3)*

> *I usually look at posts that... describe something I haven't done before, so maybe [a] new style or new tools I haven't accessed, or I might not have been really comfortable with them. So I am always looking for stuff that's pretty much new to me. And I also... look at the stuff that I've done before, [that's in] similar styles to mine. (P2)*

These quotes not only emphasize that the desire to improve is still strong after 10 years of experience (P3), but also highlight that participants seek out topics completely new to them (P2), or which enable them to compare others' methods of solving a problem to their existing practices (P2).

Across all interviews, participants expressed a variety of objectives including: staying up-to-date with the latest industry standards, improving the efficiency of their workflows, improving their end products, reinforcing existing skill sets, and uncovering new tools or capabilities. These findings are consistent with Lafreniere et al.'s analysis of comments users post to online tutorials once they have completed them. Their findings revealed a number of tutorial users that go beyond task-specific learning, such as seeking to expand one's skills set, or to shadow the techniques of other expert users [44].

3.3 Findings: Characterizing Ad Libitum Exploration

Our findings indicate that graphic designers regularly seek out resources that provide pointers to new instructional materials, which they hope will inspire them for new software techniques – a process we refer to as ad libitum exploration. However, participants also recognized that they may not find anything new, and thus stressed the need to filter content. We expand on these themes below.

3.3.1 Habitual Monitoring

For some participants, the inspiration seeking process is an integral part of their daily routine. Participants periodically check their favorite "trusted" websites, such as official Adobe forums or Reddit, with the goal of staying aware of what is happening in their field:

[Design] Reddit is nearly daily. For work, I'd say it'd be every other day... I definitely like to... pop [in] and to see what's going on. (P5)

Another popular trusted source is Twitter, where participants reported following key individuals for new tips and tricks:

On Twitter, for example, ...I follow people that are in [the] design community, that are also in my field. They regularly share links to interesting websites, or articles, or tutorials... and if it fits my interest, I will click on it and investigate further. (P1)

These behaviors suggest a clear desire to stay current and to discover new techniques that may be useful to the participant.

3.3.2 Looking for the Unexpected

Our participants reported that curiosity and accidental discovery also play important roles in inspiration seeking process. Due to the large number of features in programs such as Photoshop, even our experienced participants were not sure they knew how all the available tools work. Some participants described specializing in certain aspects of the program they use, yet still being keen to discover new tools that might be relevant to their work:

Sometimes I will be stuck on a painting and I will start staring off into the space and I will look at the buttons and will be like "hey, wait, I do not know what that does". So, I will open up a new document and I will start playing with it. (P10)

As another example of an individual seeking the unexpected, one participant reported watching tutorials on YouTube, in hopes of seeing something new through recommended videos:

And I start with a YouTube video of something I do know and then just seeing what else pops up. So that you learn stuff that you don't think to ask. (P11)

3.3.3 Low Expectations for Exploration

In some cases, participants expressed low expectations for finding inspirational content, but still engaged in watching instructional videos, in case there was something they could use in their own work:

You see what people do in interesting ways and they take these tools and they use them in new ways and they post their weird creations... and I will look at that, and even though the method is not going to be useful to me, you never know. You might learn something from that method. (P5)

3.3.4 The Need to Filter Content

Our interviews showed that with more experience, finding content for learning that is personally interesting and new becomes more and more difficult. The main challenge is that the more knowledge people have, the less likely it is they will learn something from a tutorial. Consequently, participants spoke of the high cost of viewing a tutorial and the desire to ensure that they would benefit from it:

Tutorials are usually time-consuming, so you have to be sure that it's worth your time... and to determine that... I would really be in need of the skill [covered in the tutorial]. (P1)

In summarizing our exploratory study, we found that experienced graphic designers continually seek inspiration for new ways to use design software, but engage in ad libitum

exploration with the knowledge that they may not always find something useful to them. The cost associated with finding new and relevant information raises the importance of assessing a resource's learning potential quickly and accurately.

3.4 Findings: Limitations of Current Strategies for Ad Libitum Exploration

In our study, we found a number of ways ad libitum practices could be improved. For example, social media, such as Twitter or Reddit, provide users with streams of links to new content, but often lack adequate information to help users determine whether it is worth viewing the full source material. Conversely, dedicated learning portals, such as the official Adobe forums, often provide high quality learning materials, but the pace of information flow on such portals is not as rapid as on social media, such as Reddit and Twitter.

We also note that search engines help people find instructional materials for specific problems but do not necessarily help users to discover or browse materials that push one outside one's typical work practices.

3.5 Design Desiderata for Supporting Ad Libitum Exploration

Drawing on the findings from our exploratory study, we define the following set of design desiderata for systems that support ad libitum exploration:

- **Continuous flow of recent content.** To support the habitual nature of ad libitum exploration, the system should present new content that users can access on demand

(e.g. when they are taking a break from their work). Ideally, this content should reflect recent trends to help people stay up-to-date with the field.

- **Serendipitous discovery.** The system should support the discovery of new content that diverges from a user's typical practices.
- **Filtering and browsing.** The system should allow users to browse and filter learning materials. For example, the presentation of an individual content item should be descriptive enough to help people decide whether watching or reading the content is going to be beneficial for them. The system should also include filtering capabilities that help the user target topics of interest.

3.6 Switter

We designed and developed Switter – an alternative Twitter client for software-centric learning that allows users to browse tutorials for a specific software product (in our implementation, Photoshop). Switter (Figure 2) projects the content referenced by a tweet

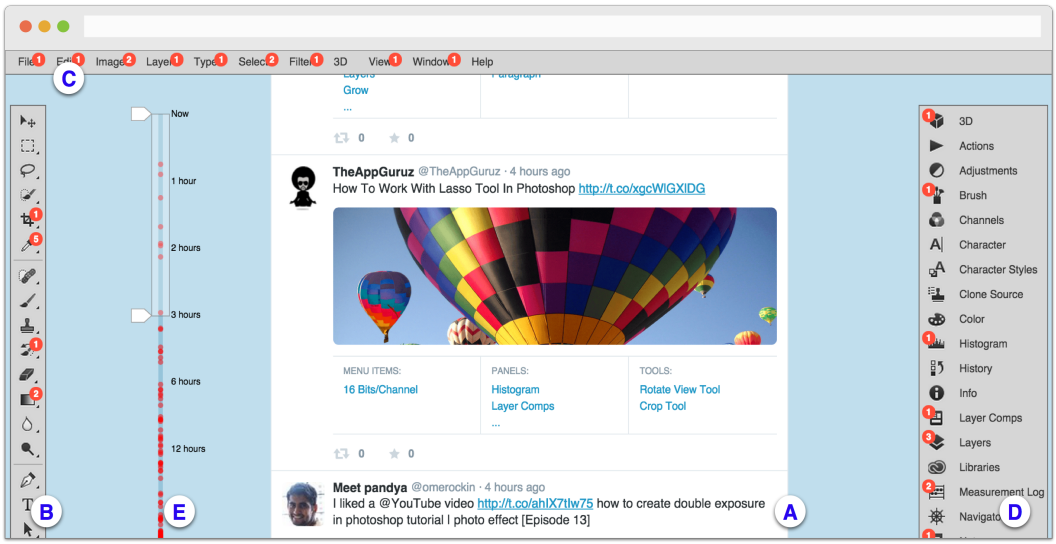

Figure 2. Interface of Switter

Integrated Twitter timeline (A); indicators of user interface elements mentioned in tutorials (B, C, D); list of commands mentioned in a tutorial (A); historical summary of tweet activity (E).

into a user interface replica and provides a Twitter timeline in place of where the document would be found in the replicated application. This enables users to browse and explore tutorials broadcasted via Twitter using the target application's interface as a navigation aid. Our current implementation of Switter is a prototype that makes use of a Wizard-of-Oz backend for extracting commands from tutorials and inserting tweets about Photoshop into a read-only Twitter timeline.

Switter's interface consists of three main components: the Twitter timeline (Figure 2, A), the interface replica (Figure 2, B-D), and the historical summary (Figure 2, E). We discuss the role of each interface component in the following subsections.

3.6.1 Twitter as a Data Source

To provide users with a continuous flow of recent content (our first design goal), we use Twitter as a source for discovering new tutorial-related information. This decision was grounded in the results of our exploratory study, where our participants positioned Twitter as one of the trusted places they used to discover new relevant information.

To get a sense of the potential volume of tutorial-related information on Twitter, we collected and examined a stream of sequential tweets that contained the word "Photoshop" over a one-hour time period. During this single hour, about 30 tweets (of a 300 total tweets) linked to Photoshop learning resources. We believe this rate of information flow to be sufficient for providing users with fresh learning material in a continuous manner.

3.6.2 Projecting Tutorials onto the Software Interface

To address the need to filter and browse learning materials, Switter highlights the areas of the interface mentioned in the tutorial. This technique 1) provides awareness of the volume

and breadth of commands used in the tutorials, and 2) enables users to browse tweets from an interface-centric perspective. For completeness, our replica contains the main menu (Figure 2, C), the toolbar (Figure 2, B), and the list of accessible modal panels (Figure 2, D).

Switter highlights a specific command or tool by placing a red dot next to it, with a number indicating the number of linked tutorials. To test the initial proof-of-concept prototype, we manually extracted commands from online content, as we focus on interaction and design. In the future, this could be automated using recent advances in command extraction (e.g. [2,27,57,68]).

To support navigating and browsing tweets, the user can click on a tool or command in the interface replica. Switter responds by reducing the list of tweets to only show those that mention the item selected.

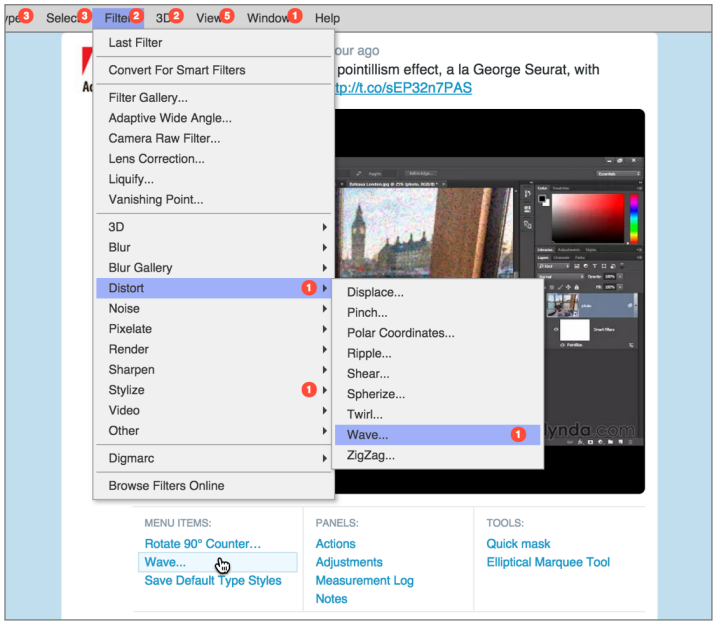

Figure 3. Switter revealing the location of a software command referenced in a tutorial.

Switter also augments each tweet with the list of commands that it references. This supplemental information seeks to help people make decisions about the utility of the individual tutorial. When the user clicks on a menu item or a tool in the list below the tweet, Switter automatically reveals the referenced user interface element in the replica (see Figure 3).

3.6.3 Temporal Awareness via Historical Summary

Switter includes a historical summary of tweet activities (Figure 2, E) to provide both awareness and filtering capabilities. For awareness, each tweet's posted date is mapped to the respective point on the timeline. This allows users to get a sense of the volume of tutorials as well as their distribution over time. For example, a user can see many tweets were posted in the last three hours, providing cues as to whether or not the volume of new information is worth browsing. Users can also use the historical view to filter the tweets projected onto the interface replica by specifying a range of time.

3.7 Field Evaluation of Switter

To test the utility of Switter, we conducted a week-long field study. The main goal of the study was to see how people would adopt the tool, and to gain initial insight into Switter's ability to support ad libitum exploration. The consent form, the recruitment script, and other relevant materials are presented in Appendix C.

3.7.1 Participants

We recruited nine designers and photographers (3 female) who use Photoshop as their main working tool. All of our participants were between 21 and 45 years old and have been working with Photoshop for at least one year. We recruited participants through online

postings on Reddit and via snowball sampling. Participants received a $75 gift card for their participation.

3.7.2 Procedure and Data Collection

Participants were asked to use the system at least once per day over a period of seven days. We did not give participants any specific task to perform, but rather asked them to browse through the content, to look for something that would catch their interest, or that could potentially teach them something new. We did not specify how much time they should spend in the system each day, nor how many tutorials to view.

We asked participants to fill out a short online journal entry at the end of each day to record their experiences with Switter that day. We also conducted two semi-structured interviews with each participant: one in the middle of the study (day 3) and one at the end of the study (day 7). Finally, we logged key interactions with the Switter interface.

We analyzed the qualitative data using the same analysis techniques as in our exploratory interview study.

3.7.3 Switter Content

We deployed Switter as a standalone web application. Throughout the study, two of the researchers working on this project actively monitored Twitter for new tweets that linked to Photoshop tutorials. For each such tweet, we manually labelled the tweet with all commands referenced in the tutorial, and added the labelled tweet to a repository. A randomized script then gradually delivered Tweets from the repository to the participants. This procedure made Switter's data flow similar in nature to the continuous data flow of

Twitter. Over the 11 days, Switter displayed 311 tweets with a median of 30 tweets per day (IQR=9).

Because our participants were joining at different times, the Switter's initial state for different participants was not the same. For example, participants starting the study on day 1 started with no content, while participants starting on day 4 could see tweets published during the last 3 days. These different initial states helped us understand Switter's utility in a variety of contexts.

3.8 Findings: Usage Logs and Daily Journals

To gain some insight into how much people used Switter, we calculated the number and duration of usage sessions. We define a usage session as an interval of activity within Switter longer than one minute, separated from other sessions by at least one hour of inactivity. We used the above heuristics to dismiss sessions that likely did not include true interaction. For example, if the user leaves the system open in the browser and occasionally hovers over the page while switching tabs, we did not count this as a session. For these reasons, we believe our summary statistics represents only a conservative report of engagement with Switter.

We observed a total of 56 sessions, with a median session duration of 24.2 minutes (IQR=27.5). Overall, the duration of user sessions was skewed towards longer times, with Q3=45min. In most cases, participants had one session per day. However, in some cases the system was used more extensively. For example, on most days, P6 had three interaction sessions and on one day, P3 had four sessions of interaction.

As part of the journal entries, we asked participants to indicate how many new things they learned that day and how many tutorials caught their interest. Despite daily email reminders, most participants forgot to make an entry at least once. Additionally, P8 encountered technical difficulties with their company's firewall. For these reasons, the total number of entries we received was 49 (as opposed to 63).

In 46 out of 49 journal entries, participants indicated learning at least one or two "new things", while only three entries indicated no learning. Interestingly, about half of the time people reported learning one or two new things, the number of tutorials they found interesting was effectively double that number (from three to five). On several occasions, participants reported seeing 10 interesting tutorials, while indicating they learned only one or two new things. These results suggest the difficulty in finding useful content, even against the backdrop of interesting content.

In terms of what participants reported learning, they described discovering underused tools, learning ways to combine several tools for neat effects, learning unknown techniques, and brushing up on their existing skillset. The following excerpts illustrate these findings:

> *The smudge or sponge tool are tools I rarely use so I learned how and why are other people using them. (P3)*

> *I was a bit weak at [the] pen tool, and today I was easily able to find a tutorial for [the] pen tool and had a good practice... I am more confident with [the] pen tool. Apart from [the] pen tool, I explored some style tutorials. I was familiar with the techniques, but discovered that there are some pretty good alternatives to get same style for the text with different techniques. (P7)*

I learned about deeper use of the spot healing brush. I gained a deeper understanding of mixing usage of the tool with the clone stamp. (P9)

We note that the above findings describe participants' self-assessments of what they learned each day. Because we provided a diverse set of tutorial content, and our participants varied in their expertise and knowledge, we did not attempt to quantify their learning beyond self-reports.

3.9 Findings: Interview Data

In the interviews, we asked participants to describe their impressions of the system and their experiences using it. We also referenced the logs and journal entries, and asked participants to elaborate on certain behaviors and learning outcomes.

We structure our findings by considering how Switter supports learning via its ability to browse and filter content; how it can be used as a reference tool; and opportunities for improving these types of systems.

3.9.1 Browsing and Filtering Content

Overall, participants were positive about the system, indicating that it was helpful in finding useful content:

It's much easier than Google. Like I said, if you just [added] a search bar, it's competition to Google now for tutorials. (P7)

Participants indicated that Switter's features allowed for a number of new browsing and searching behaviors that might be hard to achieve with currently available tools. In what follows, we highlight participants' feedback on these capabilities.

Projecting Commands on the Interface Replica. Participants appreciated Switter's ability to project the commands and tools referenced in a tweet onto the Photoshop interface replica. This domain-specific rendering of the tweets helped them browse and filter the content referenced by the tweets non-linearly. The following quotes demonstrate this appreciation:

> *I haven't seen such [a] thing at all. I haven't imagined that there can be a system where you could use commands to filter tweets and learn from that. That part is really awesome. (P7)*

> *I really, really liked that, actually – that it has the overlay [highlighted commands]. I was actually showing a couple of friends of mine... and they thought it was so cool, the layout. They liked overlay. (P1)*

Participants adopted a range of new browsing techniques that leveraged command-specific filtering and the projection indicators in the replica, which we describe next.

Engaging in Popularity-Driven Exploration. Participants reported that the interface projections guided browsing behavior by attracting their attention to commonly used tools and commands. For example, P3 indicated that the projections allowed her to identify commands frequently referenced in the tweets, which made her curious about what people use them for:

> *I just think it's really cool that you can click on everything and you see people using [commands] and I am, like "I wonder why, or how", you know? ...And I feel like this is more... for the curious mind, that is "ah, what is this? I want to learn more about this thing". And then you can... from there you click and it will show you this many possibilities. I do like that a lot. (P3)*

As hinted above, P3 would filter tweets by clicking on the specific tools that she found most interesting. In turn, command-level filtering satisfied her curiosity by exposing her to the variety of use cases of the selected tool.

Exploring to Address Weaknesses. While P3 used the popularity of highlighted elements to guide her exploration, a few other participants started by first identifying weaknesses in their skill set they wished to address:

> *So, I can just focus on my weaknesses here, like whenever I jump in a system, I just go to the pen tool and start practicing with it. I do not have to go through other stuff that I do not want... I can just focus on my skills. (P7)*

> *I am really bad at masking, so I always kept certain edges when I did masking. And I saw the tool listed as masking and I particularly got interested and just went there and checked the video when it's listed there. (P5)*

Comparing Alternatives. Our interviews revealed that Switter's command filtering capabilities could be used to draw comparisons between several commands, a use case we had not anticipated. For example, P9 described the ability of "jumping" between two commands as one of his favorite aspects of the system:

> *For instance, one of the tools I was learning about yesterday was the patch tool. That was the tool I haven't used before. And the thing that I kind of found when I was using it was it would be a good alternative to the clone stamp tool. And the clone stamp was the thing I used in the past most frequently for that type of work... One of the things that I found the most useful about the program was the ability to quickly jump between tutorials that use the patch tool and back to tools with the clone stamp tool, to see if there's overlap between the two, to kind of see which tool would be most effective in which situation. And being able to do that quickly was what allowed me to make that comparison. Because obviously, if I did not have something that allowed me to jump quickly... it's kind of difficult to see where that overlap is. Just because you kind of lose track of where you are at. (P9)*

What is notable about this use case is that the participant learned not by consulting a single resource, but by explicitly juxtaposing the content of multiple learning resources to compare and contrast alternative methods. Switter aided this process by helping them first locate this similar content, and then swiftly move between the resources.

Discovering Synergies. In addition to comparing workflows, we found that participants would also use Switter's filtering capabilities to learn which commands could be used in conjunction with others:

> *I was looking into ones that have burn and dodge, but then it was interesting to see what they use in conjunction with [burn and dodge] ... So that's kind of when I would look into, like, the menu items... It's cool to see combinations. Like, it's not always about the one. It's about how they fit together. (P1)*

At present, Switter lets users filter tweets based on single selections. However, several participants expressed interest in filtering tweets by multiple commands at the same time. We believe that adding such functionality might help users compare command capabilities and search for tutorials that illustrate specific tool combinations.

Filtering Tutorials Using the Command Summaries. Participants also made use of the summary of commands Switter provides below each tweet. One participant preferred these summaries over interacting with the interface replica directly, which she sometimes found overwhelming:

> *It's a little bit overwhelming, you know? ...If I click on layers, there's a huge amount here... Whereas if I am looking at this one video and I click on the layer there, it takes me to that specific tool it's going to be using, in that context. (P6)*

As hinted at by this quote, clicking on commands in the summary list carries with it the advantage of teaching one where the command can be found in the interface.

Seeing the list of commands also enabled experienced users to quickly assess what is being covered in a tutorial. For example, several participants described how this information allowed them to essentially recreate the tutorial in their heads without having to look inside:

> *I see that they use "Desaturate," I see that they use "Gaussian Blur," and it makes sense. It's like, I do not need to go into the video and spend 15-30 minutes there, because [by] just looking, and "Oh, desaturation is what takes away the gloss from the photo." Like I know how to use the texture, or I know how to blend. I know WHY would they use "Gaussian Blur," I know why they would use desaturation. (P4)*

Uncovering Unexpected Usages. Conversely, when participants could not imagine how the list of commands could achieve a given result, they pursued the tutorial to learn a new workflow. In these cases, participants compared their existing workflow to the one in the tutorial to find out which one is more efficient:

> *So, it's pretty much just looking for a way to do things quicker than you are already doing. So, sometimes seeing the commands it's like "ah, that makes way more sense to be doing it this way instead of the way I've been doing that." (P1)*

Similarly, unexpected uses of a specific command or tool incited participants' curiosity. For example, P4 noted that looking for such "out-of-the-box" knowledge was something he was most interested in:

> *If someone, for example, comes and tells me that "I used my hammer to eat my noodles," I'd be very curious, like "how did you do that?" Because a hammer is for banging a nail in its head... It's something that's out-of-the-box knowledge that you get from people. (P4)*

3.9.2 Switter as a Reference Tool

While we originally requested that participants use Switter at least once a day, a number of participants integrated Switter into their existing workflow. For example, P6 indicated that he used Switter as a convenient place to look up how others use the tool he had just struggled with:

> *I was trying to use the clone stamp tool. So, I wasn't quite getting it the way I wanted it to, so I ended up in [Switter] and clicked on a clone stamp tool and simply filtered*

down to all of those videos using clone stamp tool, and that's a much easier way for me to find tutorials in context with certain things I am struggling with, you know? (P6)

3.10 Findings: Opportunities for Improvement

Participant comments also provide insight into ways that Switter could be improved to better support ad libitum exploration.

3.10.1 Curating Incoming Tutorials

Many participants appreciated being constantly exposed to fresh learning materials. However, a few reported that the quality of the delivered tutorials often did not meet their expectations:

What deterred me the last few times, was that some of the content did not appeal to me at all [...] another one is like "how to make water drops" and the picture of water drops [...] Look at those water drops. No offense, but those are really bad water drops. (P1)

This quote suggests that low-quality content might discourage some users from exploring new learning materials. In future design iterations, one could explore mechanisms to curate twitter-retrieved tutorials, so that only higher-quality tutorials are delivered to the users.

3.10.2 Improving Tutorial Summaries

Participants appreciated the command summaries below each tweet, but many wanted more detailed information about each tutorial. For example, some requested that command summaries reflect the order in which operations are performed in the tutorial. Some participants also wanted the tutorials labelled according to the higher-level skills covered,

similar to the approach explored by Kim et al. [41]. Moving forward, the challenge will be both obtaining accurate labels and finding ways to provide this additional information without visual overload.

3.10.3 Re-finding Tutorials

Some participants reported finding a useful tutorial early in the study, but had difficulties re-finding it later. Participants also commented that they did not always have enough time to watch and follow a tutorial that caught their attention. These two needs suggest that it may be worthwhile to include search-based capabilities, or bookmarking functionality, to enable re-finding interesting content at a later date.

3.11 Discussion and Future Work

Social media like Twitter provide a platform for users to easily share and discuss instructional materials for software. Our field study provides encouraging insight that Switter's approach of projecting these tweets onto a replica of the application's interface helps support ad libitum exploration, by helping users browse and locate learning materials of interest. Our results also suggest that Switter's organization of these resources helps preserve the curiosity-driven component central to ad libitum exploration: users in our field study described a number of instances where they used the tool to uncover unanticipated and sometimes unorthodox bits of new knowledge.

In light of these promising initial results, we discuss a number of directions for future work.

3.11.1 Measuring Effects of Ad Libitum Exploration

While our participants reported seeking inspiration each day, our study did not include a control condition to provide a baseline for comparison, nor did we explicitly test the outcome. Thus, one fruitful path for future work is to compare Switter to existing, general-purpose social media clients, with respect to their ability to support ad libitum exploration. From our study results, it is clear that Switter provides some clear benefits compared to existing clients (e.g., the ability to quickly toggle between sets of learning resources by clicking on commands), but it is worthwhile to quantify the impact of these features.

3.11.2 Scalability of the Approach

During the study, Switter published around 30 new tutorials per day. This information flow was sufficient for the purpose of our study but did not allow us to test how well our design scales to a larger volume of tweets. For example, our preliminary analysis suggests that there might be closer to 30 Photoshop tutorial tweets per hour. Our current design attempts to address this issue by providing a historical overview and time-based filtering of tweets. Nonetheless, how well this approach scales needs further research.

3.11.3 Integration with Software

We implemented Switter as an independent web tool, rather than integrating it into Photoshop itself primarily for ease of prototyping. However, we also believe that Switter's independent format has its own benefits, for example, enabling users to browse the resources on devices that do not have Photoshop installed. Investigating the tradeoffs between in-application instrumentation and independent application is an area for future work.

3.11.4 Tutorial Annotation

Our field study included a Wizard-of-Oz component by virtue of us hand-annotating each tweet with the commands used in the referenced tutorial. Prior work has examined ways to automatically extract commands from both text [27,40] and videos [2,57,68]. Incorporating automated extraction techniques into a system like Switter will introduce noise into the system, which will alter the user experience. However, there are alternative approaches to determining the commands mentioned in a tutorial. For example, if workflows are automatically captured within the application, as Chronicle does [8], precise command metadata is readily available. Command annotations could also be crowdsourced, for example, within an application like Switter.

Beyond command annotations, some participants desired more descriptive labels for the instructional content. For example, they wanted to know whether it described a particular technique, or the intended general audience for the material (e.g., an interface designer vs. graphic artist). These labels are further candidates for crowdsourcing approaches, given that automatically categorizing tutorials across dimensions (such as audience) is currently a challenging problem. Alternatively, one could provide explicit support to organize tweets by hash tags, as described above.

3.12 Summary

This chapter investigated the practice of seeking inspiration for new software techniques, which is common among graphic designers. Through our exploratory study, we found that expert designers often participate in such a form of inspiration seeking by habitually monitoring social media – a practice that we refer to as ad libitum exploration. We also

found that when engaged in ad libitum exploration, designers often struggle selecting and appraising suitable inspirational materials among the abundance of information on social media, such as Twitter.

To support ad libitum exploration, we developed Switter – an alternative Twitter client that contextualizes tweets that mention Photoshop tutorials into a replica of a Photoshop's user interface. Switter augments tweets that mention Photoshop tutorials with a list of commands mentioned in the respective tutorials and provides designers with novel ways to browse and discover Photoshop tutorials via the interface replica. Our weeklong field study revealed a number of compelling use cases for Switter, suggesting the utility of its approach to support the practice of seeking inspiration for new software techniques.

Chapter 4

Supporting Immediate Reflection on Collections of Inspirational Images

As mentioned in Chapter 2, prior research has shown that some of the inspiration seeking strategies employed by experienced graphic designers increase their chances of getting affected by design fixation [62]. Prior research has also proposed the idea of mitigating these fixation effects by designing systems that inform designers about fixating characteristics of their image collections, such as familiarity of the images to the designer and similarity between saved images [62]. However, the success of such systems would depend on whether designers reflect on familiarity and similarity of saved images during and immediately after visual research – a process we refer to as *immediate reflection*. To the best of our knowledge, there is limited insight in the literature about immediate reflection in the context of visual research in graphic design: whether it takes place and (if so) what designers' motivations and goals are for such reflection. Moreover, it is also not

clear whether displaying fixating characteristics of image collections would impact designers' inspiration seeking strategies.

This chapter describes results of an exploratory study with 14 experienced graphic designers that investigates i) whether and why graphic designers engage in immediate reflection on images collected during visual research, ii) whether designers might recognize the benefits of displaying fixating characteristics of their collections, and iii) the potential for systems to support immediate reflection. During the study, we asked our participants to use a technology probe [36] – a system that we named *Instory* (Figure 4) – to gather inspiration for three distinct ideas for a hypothetical design task. Instory aims to facilitate immediate reflection on the variety of collected images by augmenting collections of saved images with elements of a designer's search history.

Parts of the research described in this chapter has been published and presented at the ACM SIGCHI Conference on Designing Interactive Systems (DIS) in 2018 [20].

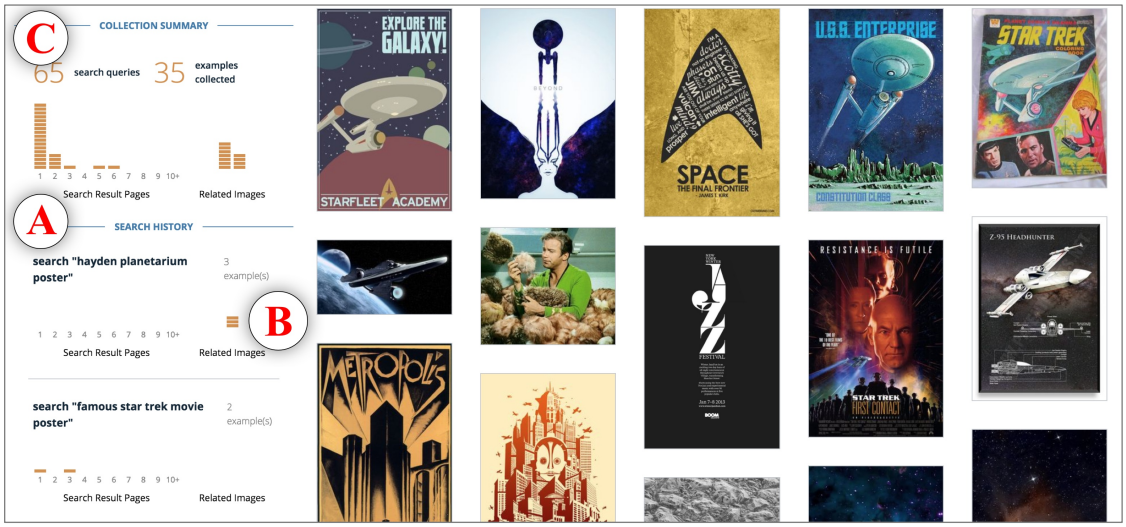

Figure 4. Interface of Instory.
Full list of search queries used in visual research ordered by the number of saved images per query (A); relevance of an image to the respective search query (B); summary of search trails for the collection of images (C).

4.1 Instory

In light of our research questions, we were particularly interested in whether graphic designers would see value in reflecting on the variety of collected images, as the lack of variety has been shown to increase the chance of design fixation [66]. To better convey the variety of collected images to a designer, we decided to show them how they found each image. Specifically, we decided to display the list of keywords that a designer used for retrieving saved images and the relevancy of each saved image to the respective keyword. By showing this information we aimed to give a designer an additional dimension for assessing the similarity between collected images. This was based on an exploratory idea that images highly relevant to the same search query might share greater similarity than either i) images with low relevancy to the same search query or ii) images collected from different queries.

We implemented our approach in Instory (Figure 4) – a system we used as a technology probe in our study. Instory captures and displays the full list of the images saved by a designer during visual research and the search queries used to find these images (ordered by the number of images saved for the query; Figure 4 B). For each saved image, Instory also aims to show the relevance of the image to the respective search query by showing which page of search results the image came from (or whether it was found via searching related images; Figure 4 B). The decision to use the location of an image in search result pages was motivated by the fact that search engines return results in decreasing order of relevance to the respective search query. Therefore, images that are located higher in the search results are more likely to illustrate similar ideas than images located lower in the search results. Instory also provides a compact overall summary of relevant search activity,

depicting the number of queries and overall relevance of saved images to queries (Figure 4 C). We included the summary information to see if designers might find this as an useful way to reflect on the overall similarity of saved images.

We implemented Instory as an add-on for Pinterest – a popular platform for saving and sharing images on the Internet. Pinterest allows its users to browse images and to save them to personal or shared "boards" – user-defined containers that they can use to store and share the images that they like. Instory provides an alternative view for a Pinterest board: it shows a designer all the images that they saved to their Pinterest board with the relevant search information on the side (Figure 4). The search information also serves as an interactive index into the collection of images, allowing designers to filter images by query or by a specific page of search results in a query.

We chose Pinterest as a target platform because of its core functionality of maintaining curated collections of images natively supports the type of visual research studied in this thesis. In addition, informal conversations with a few professional designers indicated that they often use Pinterest for visual research.

Our integration with Pinterest relies on a Google Chrome extension that monitors AJAX (Asynchronous JavaScript and XML) calls done by the Pinterest webpage, such as searching by keywords, loading search results, and saving an image to a Pinterest board.

4.2 Exploratory study

To investigate whether and (if so) why graphic designers engage in immediate reflection on images collected during visual research we conducted an exploratory study, where

Instory was used as a technology probe. We chose the technology probe method because of its ability to solicit answers to open-ended research questions [36] and we chose an in-laboratory evaluation method to create realistic reflection tasks which would allow us to get insights into Instory's utility in a single inspiration seeking session per participant. The consent form, the recruitment script, and other relevant materials are presented in Appendix D.

4.2.1 Participants

We recruited 14 expert graphical designers (4 female) by posting advertisements on local message boards and on social media (Reddit). All participants had at least one year of professional design experience. Participation was rewarded with a $40 gift card.

4.2.2 Procedure and Study Task

Participants recruited through local message boards participated in the study in person, while participants recruited via the Internet participated through Skype, Google Hangouts, and similar services.

At the beginning of the study, we introduced our participants to Instory. After a brief demonstration of its features, we gave participants 5 minutes to familiarize themselves with the prototype. The study task asked participants to come up with three different ideas for a hypothetical design task and to present them in the form of mood boards – collections of images that convey a general feel and style for a chosen design direction [53]. We used the help of a local graphic designer to craft a plausible design scenario, in which a hypothetical client asked our participants to deliver three mood boards, each representing a different design idea for a poster advertising a Star Trek-themed event:

The New York Planetarium asked you to help them with designing a poster for their new Star Trek-themed science show. Your client asks you to come up with three different ideas for the design and present them in form of mood boards – collections of images that would convey style, feel, colors, textures, etc. of your idea.

After the participants created the three mood boards, we conducted a semi-structured interview. During the interview, we asked participants to describe the instances of immediate reflection which we observed during the study. We also solicited feedback on Instory's approach of augmenting collections of saved images with elements of search history.

On average, each study session lasted approximately one hour (30 minutes for the study task and 30 minutes for the interview).

4.2.3 Data collection and analysis

During the study, we observed our participants' inspiration seeking strategies while they were working on the study task (for remote participants we asked them to share their screens). We also tracked saved images and search information, and recorded user interviews which were later transcribed. Data from the transcripts were analyzed by creating affinity diagrams using a bottom-up inductive approach [79]. The researchers involved in the project then held joint data interpretation sessions where we extracted common themes from the affinity diagrams.

4.3 Findings: Usage Data

On average, participants saved 25.9 (SE=3.2) images from 26.9 (SE=5.3) search queries. The largest number of saved images was 46 over 70 search queries (P7), and the smallest was 8 saved images over 13 search queries (P9). This degree of variability in the volume of saved images and searches is at least in part owing to different strategies that the participants used to complete the study task. The variety of observed strategies for seeking inspirational images aligns with the prior findings of Miller et al. [62], who found that there is no one universal strategy for seeking inspirational design examples.

4.4 Findings: Immediate Reflection on Collections of Inspirational Images

The majority of our participants approached their visual research in an iterative manner, alternating periods of exploration (during which they collected design examples) with periods of immediate reflection (during which they examined what they had collected). From our participants' interview responses, we synthesized three main goals for such immediate reflection: refining exploration directions, assessing collection saturation, and preventing confirmation bias.

4.4.1 Refining Directions for Further Exploration

Several of our participants took a bottom-up approach for their exploration. Specifically, starting with some initial direction, they would collect every single image which appealed to them, after which these participants would go back to what they had saved and search for emerging themes. These participants would then use these themes to inform further

exploration and to decide when to stop it. The following quotes by P10 illustrate this strategy:

> *It's not searching for very specific things that you want, it's searching for things in that general area and then figuring out [...] what the actual thing you are looking for is. (P10)*

> *There's no criteria [for stopping exploration]... it's just me going back and looking at what I have and deciding if it's enough for me to... talk through what I am trying to convey and if it illustrates the concept well enough. (P10)*

4.4.2 Assessing Collection Saturation

A few of our participants approached the study task in a top-down manner, by first coming up with a direction for a mood board and then seeking images to represent various facets of their idea. For example, one of the mood board ideas generated by P13 was "space exploration". To represent this idea, P13 wanted to show images of a space mission, views from inside and outside the space station, etc.:

> *I definitely come in with an idea in my head, with, like a visualization of it. So, once enough of those pictures that I've gathered fit that visualization, that's when I generally deem it successful. Like, going back to the "explorers" example, I wanted inside and outside the space station, someone exploring a new world and some kind of ship in itself. So, that's where I got... like the six images on that board. Like, there's someone in the ISS taking a photo, there's a picture from inside the [...] space station itself, the [...] astronauts from the moon, and someone exploring the cave – I felt that [those pieces] solved the puzzle of "explorers" topic. (P13)*

These participants used immediate reflection as a method for determining whether their collection covered enough facets of their idea or whether they needed to extend it, which is illustrated by the following quote:

> *The Apollo mission pretty much perfectly nails down space explorers. So, that part of the topic is satisfied. I want to now show additional parts of being a space explorer, that is unrelated to pure astronauts. (P13)*

4.4.3 Preventing Confirmation Bias

Half of our participants mentioned reflecting on the variety and familiarity of saved inspirational images as a way to prevent confirmation bias from affecting their work. In psychology, confirmation bias is the term that "connotes the seeking or interpreting of evidence in ways that are partial to existing beliefs, expectations, or a hypothesis in hand" [64]. Based on our interviews, we found that in the context of online visual research, confirmation bias presents itself as an inclination towards a certain design direction or imagery due to personal taste and preference, which could lead to design fixation.

In their interviews, several participants told us that their personal preferences towards certain design styles might make them want to pursue these styles in their inspiration seeking:

> *A lot of the times, [...] I will spend all day thinking of some topic and getting images put together or sketching some things, or just figuring things out on my own... and [later,] I will go back and look through the source material that I am working with and I see that all the [images] are basically the same. This means that in my head, this is the way I've been thinking this was going to go, and if that's the case, why would I*

even search for inspiration if I am just keep finding the same things that I want to find? [...] if my taste is about Star Wars posters is predominantly purple, my results for things are going to be purple for the most part [...]. (P14)

... the first thing you see, and you like it, your mind gets stuck there. But you always have to [keep searching] for more options. (P2)

However, at the same time the participants spoke of realizing that their work is created not to please their aesthetic needs or their client's needs, but rather to appeal to their target audience. Therefore, they stressed the importance of reminding themselves to get out of their comfort zone and consider stylistic choices that might not be their favorite, but might be suitable for the project:

There's always confirmation bias when you do this kind of search. That's why you have to be careful. And also, you'd have to take your pre-conceived notion of how to tackle the problem, and state it for yourself, that you know that it's there. And then ignore it. But you have to be conscious of that, like "oh, I am thinking of doing it this way", you know? Like I've said, when I [started the task], I was really inclined to do modern, and in my first search I did the retro and I kind of [felt discouraged], because I don't like retro. But then I think that this is a way to do it, there's an audience, so I got to take my bias from out of there, and then I remembered that NASA did retro really well, tastefully, right? And that's why I made this [mood board], right? If I was going with the one approach, like the client wanted one direction, I would not even consider this, right? Unless the client specifically pushed on it. (P4)

Unfortunately, our participants did not have a definitive strategy for dealing with confirmation bias. The few participants who were aware of these biases told us that they have to rely on their experience and on constant conscious effort to prevent them from getting affected by design fixation:

> *[You have to] keep telling yourself that you have to focus on the task, like, there's no system or there's no way of doing it. (P2)*

> *I would write down the ultimate idea that I [am pursuing] [...] [When I get new ideas as I look through images,] I would constantly go back to that script to see if these new ideas match that original idea. Because the original idea [...] was conveying a certain message that I was excited about [and] that client also agreed [to]. (P13)*

One participant also mentioned that their prior experiences often result in them using the same keywords and inspirational images to inform the designs for multiple projects. This participant mentioned that pro-active reflecting on the familiarity of collected images helps them to stay aware of this bias:

> *If I am searching for something, and then I notice that all the images returned become too familiar [...]. Maybe I searched for them, like a thousand times [in the past] and I should probably figure out some other way to get different keywords. Sometimes I've been driven to use a dictionary to find synonyms for words in order to get Google to get some different results, to kind of see newer things. (P14)*

This strategy aligns with prior work suggesting that designers use synonyms and antonyms to generate more search keywords for visual research [62].

During the study, we observed how one of our participants (P12) was affected by confirmation bias. During the study, P12 thought that they were collecting a variety of images. But when they looked back at what they had collected, they saw that their collection had the same type of image displayed from different perspectives.

> *...when I did my initial search of galaxies, I liked the bunch of images, but I found that they were all the same color, and they were all of kind of similar variety. And [that happened] because I've been pulling them together image by image [and I thought they were different], but [in reality] I was looking at the same picture. (P12)*

Immediate reflection on the variety of saved inspirational images helped P12 to detect that their exploration was impacted by their personal bias. It also motivated them to backtrack their process and fix their collection by pulling more diverse images from already explored searches:

> *So, I like when I [can] go back to that search [query] again and say, "hey, let's find some more with more orange hue, or red hue", something I was not looking for initially. (P12)*

Interestingly, P12 was one of the few participants who talked about being aware of the potential effects of confirmation bias, and pro-actively trying to avoid it through conscious effort. When we asked P12 why they thought they were affected by this bias despite their effort, the participant suggested that it could be attributed to the sequential nature of image collection, where saving images one by one does not provide awareness of how adding a new image may impact the rest of the collection:

I was on Pinterest and I was pinning the images, I was seeing them individually. And then when I went to the system, I saw them all together. And then I noticed that I was pinning all these similar images. (P12)

As we indicated earlier in this section, preventing design fixation via conscious effort was a common theme across our participants. Given this prevalence, the above quote from P12 suggests that this strategy might not always be effective, which opens up interesting opportunities for system design. We return to this point in our Discussion.

4.5 Findings: Utility of Augmenting Image Collections with Elements of Search History

Instory aims to convey the varied nature of saved inspirational images by displaying the list of search keywords used to find those images. However, the majority of our participants told us that they did not feel the need to refer to the search-history information provided by Instory, as they could remember what they searched for and how they found each image:

Because I still had all those [images] open [in browser tabs]. (P3)

If it [happened] a day or two [ago], [then] I can remember [what I have searched for]. But if it is past a week, then [list of search queries] would be very beneficial. (P9)

When asked whether they see any potential uses for displaying their search queries, the majority of our participants suggested that this information could be useful as a "breadcrumb trail" which would allow them to return to a specific point of their process in later stages of their work

> *That's where that really comes in, where I can keep track of what I have already searched, what are the roads that I have already gone down and I do not need to go down those roads anymore, because I feel like I have already gotten all the ideas, I picked the flowers from it. (P13)*

Our participants said that such a trail could be useful for when they need to change to an alternative design idea, or when they require additional images for a specific design element. The next section discusses this finding in more detail.

Instory also attempts to convey the proximity of collected images to their respective search queries by showing designers the location of saved inspirational images in the search results pages. We hoped that such information might help designers to recognize when they are getting too specific in their visual research. However, our participants told us that it was not obvious how to interpret this information:

> *What would I do with the infographics of the queries? Like, search for Star Trek type word and I can see three examples and it's on page 1, 2, and 7. So [...] what do I do with [this] information? (P2)*

In addition, after we explained the purpose of this information to the participants at the end of the study, they did not believe such information to be valuable to their task:

> *This kind of distribution you had in the user interface, to me it was just... too small to get anything useful from it. I did not really feel any... need to click any of the... any of the distribution, the bars to [focus on images from] those searches. (P12)*

In this quote, P12 refers to distributions of collected images across the pages of search results, which were displayed by Instory. In their case, these distributions were very sparse, as P12 collected only a few of images from each page of search results. This quote suggests that that such information could be too low-level for easy interpretation, however, more research is needed to determine whether higher-level summaries of this data could be useful for designers.

4.6 Findings: Potential Utility of Search Information as a Long-Term Memory Aid

As mentioned in the previous section, the majority of our participants told us that displaying a list of their relevant searches alongside collected images could be useful as a long-term memory aid. These participants indicated that most of their visual research is done in multiple sittings, mediated by pauses that might span a few days. With long breaks between the sessions, participants would not be able to rely on their memory and their browser's history to recall which ideas they have explored, and which search queries have resulted in them finding inspiration. These participants indicated that seeing search information snippets relevant to their visual research could help them to recreate their thought process:

> *...I did not just type, like, "space" to get it. I typed like "deep space exploration"... And it lets me know [that] I was going for something that portrays [space] as a big, large, never ending idea. (P12)*

> *...if I am working on something, I have a lot of things going on. So, I will leave it for a while and then I will come back to it. And [when] I came back to it, […] I was*

confused about what I've done so far. Having something like Instory [would help me] to go back and see what my thought process was. [It] would just bring me up to speed and help me start on it again. (P9)

A few participants also told us that the search history could help them to craft follow-up searches, for example, after a client gives them feedback on their initial ideas. These participants felt that seeing their prior successful and unsuccessful searches could help them to better choose paths for more detailed investigation and to avoid paths that did not lead them to anything inspirational:

...if a client says "we really want astronauts", I would go back and figure out how I got that astronaut to get more, you know, pictures of astronauts. (P13)

...for example, I searched for "star trek t-shirt" there [and did not find anything interesting]. [If I came back to this task later,] I would searched that again. I would definitely fall into doing that again, I would definitely type that in, I just know I would. Whereas [seeing that I have already searched for it and did not find anything interesting], I already know [that I should] not to do that [again], because there was not anything for me to see. So, that way I could save some time. (P7)

...this is cool, it shows me even searches [that] I did not save [any images from] [...] [I find it cool] because these are kind of things that I just discarded. I am not going to be searching for anything like this [in the future] if I have no examples from it. (P9)

Collectively, these findings suggest that a summary of search history could provide designers with a memory aid for a long-term visual research. The next chapter of this thesis provides a more in-depth investigation of this direction.

4.7 Findings: Opportunities for Improvement

Participants also spoke to a number of ways in which we could improve the way we display search history to better support visual research.

4.7.1 Documenting Visual Research Instead of Augmenting Saved Images

Many participants suggested to expand Instory beyond just saved images, as saved images alone do not capture all the aspects of their visual research process. This feedback aligns with prior research which found that collections of images that designers save often do not contain "stories" behind those images, such as the reasons why the designer saved each image [75].

In our study, many participants said that they often want to return to an image that they saw at some point in the past but did not save because it was not immediately relevant to what they were doing. They suggested to augment the search summary with all the images that catch their attention during visual research, in case they want to return to them in future:

> *So I've shown interest in this image, so I open it and then closed it. So it should be in the system. Like, […] it shows me every image that I opened. (P5)*

Chapter 5 provides further investigation into the utility of capturing visual research beyond explicitly saved images for retaining "stories" behind saved images.

4.7.2 Chronological History of the Search Summary

One of the comments received from the participants concerned the ordering of search queries shown on the side. By design, Instory shows search queries in decreasing order by the number of images saved. When working on the task, many participants refined their search keywords in an incremental manner, e.g., by adding or changing adjectives to a target keyword. Since these incremental searches represent one idea, participants expected Instory to group or somehow relate them. Participants told us that seeing the progression of their searches might help them better understand their thought process at that time and would help them to plan their further investigations.

4.7.3 Relevance of Collected Images to Search Queries

Instory attempted to convey the proximity of collected images to the respective search queries by showing our participants the location of saved inspirational images in the search results pages. We hoped that such information might help our participants to recognize when they are getting too specific in their visual research. However, our participants did not know how to interpret such information. Specifically, our participants told us that it was not obvious for them that the location of an image on a list of search results could indicate the relevance of the image to the search query. After we explained the purpose of such information to the participants at the end of the study, they did not think that such information would be valuable for their task. We observed that the participants' visual research focused on finding specific images, hence the participants saw more value in seeing whether or not a search query had helped them to find a suitable image than in seeing the relevancy of an image is to the search query.

4.8 Discussion and Future Work

In this section, we reflect on our findings in light of our study goals and discuss promising venues for future investigation.

4.8.1 Utility of Search History Information for Immediate Reflection on Image Collections

As mentioned in the previous section, our participants found only limited use for the search information presented by Instory within the scope of our study task. We suspect that this could be partially attributed to the high familiarity of the domain to our participants. Almost all our participants were familiar enough with Star Trek to be able to generate the ideas for three mood boards before conducting any visual research. As a consequence, our participants could generate design ideas for solving the study task without engaging in preliminary research, and had no difficulty coming up with keywords to describe and research their ideas. While such a strategy is valid, it skips an initial domain research stage, which, according to prior work, is meant to give designers a better understanding of the problem domain, possible design ideas within this domain, and domain-specific vocabulary to describe these ideas [33,83].

One direction for future investigation is to explore the potential utility of search history information to facilitate visual research for a design task in an unfamiliar domain. We assume that a less familiar domain could motivate designers to research the problem domain before generating mood board ideas. Search history information provided by Instory might be useful during such research, as reflecting on the variety of searchers might give designers a better sense of how much they researched the domain, which in turn could help them avoid preliminary commitment to certain design ideas.

4.8.2 Immediate Reflection on "Fixating" Characteristics of Image Collections

Our findings suggest that designers see value in reflecting on the similarity and familiarity of collected images, which are both considered to be "fixating" characteristics of image collections. Specifically, we found that such reflection helps graphic designers stay aware of the effects of confirmation bias. Future work might investigate whether graphic designers see the value of reflecting on other "fixating" characteristics of image collections.

One idea for future work is to investigate how to facilitate reflection on the commonality of ideas represented in image collections. To facilitate such reflection, one can imagine a system which retrieves not only images saved by a designer, but also images that are related to them. Displaying these related images alongside collected images might help the designer to assess whether ideas represented in their image collection are common in other designers' approaches. However, it might not be obvious to a designer why a system displays related images. Moreover, this additional information might introduce noise into collections, making it more difficult for a designer to work with the collected images. Further research is required to investigate the potential advantages and disadvantages of utilizing of such an approach.

Future research may also explore ways to facilitate reflection on the proximity of saved images to the design task. Based on our observations, we suspect that displaying search queries used to find images might provide value for such reflection when accompanied with some degree of interpretation. As an example of such interpretation, a system could compare the search queries to the task description and inform a designer of the degree of deviation between their keywords and the task description. Alternatively, a system might utilize the keywords used by a designer to deduce the domain of the design task and suggest

additional topics to research in this domain. However, one of the challenges of designing a system like this is finding the right amount of interpretation for the system to provide to a designer. For example, too little interpretation might make it unclear for a designer why a system is showing them additional information, while too much interpretation might be perceived as taking over a designer's creative judgement.

4.8.3 Potential of Tools to Reduce and Facilitate Recovery from Confirmation Bias

Our findings suggest that one of the main reasons for immediate reflection on collected images is to assess the potential impact of confirmation bias. This result indicates the value of providing the ability to reflect on the diversity of the images that have been collected. Similar to what was suggested in [62], a system could augment the collection process by informing a designer about how each of the images that the designer sees during their research relates to other images the designer has already saved. For example, a system could use image processing algorithms to compute joint visual characteristics of the collection of saved images, such as color information, composition, and types of objects depicted on each image. Before a designer adds an image to their collection, the system could notify the designer how adding this image would impact the varied nature of their collection, e.g., how it might change the overall color profile of the collection, or whether it may impact the variety of objects depicted in their collection. However, prompting a designer with this additional information might be perceived as distracting, and may cause the designer to lose their train-of-thought.

When providing feedback about Instory's features, the participant who experienced confirmation bias during the study (P12) despite their efforts to prevent it, told us that their search information could help them to recover from confirmation bias. Specifically, they

suggested that after detecting that all their images were alike, the tracked search terms might allow them to go back to each query and find a more varied set of images. Future work could investigate the potential of a system specifically designed to support recovery from confirmation bias. For example, one could imagine a system which allows a designer to easily substitute an image in their collection with an alternative image from the same search query. This could be achieved by saving a diverse set of related images alongside each collected image. The system could then allow a designer to substitute individual images or even an entire collection with related images, with the intent of allowing the designer to explore the same design idea but from different perspectives.

4.8.4 Search History as Long-Term Memory Aid

Our findings suggest that the search history approach has the potential to support long-term visual research. Specifically, our participants told us that displaying their relevant search history alongside their saved images has the potential to capture and visualize their design process. Capturing their process could help designers to remember "stories" behind images, such as the reasons why they saved specific images, as well as capturing design directions that they have considered but never explored. Participants also suggested a number of ways to augment Instory's functionality to better support their long-term needs, such as capturing images that they considered but did not save and enabling advanced filtering by image metadata. The next chapter of this thesis investigates this direction in more depth.

4.9 Summary

This chapter investigates graphic designers' motivations for reflecting on collected inspirational images during or immediately after collecting them – the process which we

refer to as immediate reflection. We described Instory – a prototype system that aims to facilitate immediate reflection on the variety of images collected during visual research by displaying relevant search information alongside the images.

Our findings from an exploratory laboratory study with 14 experienced designers contribute new insight about graphic designers' motivations for immediate reflection on collected images. For example, we found that one of the reasons for such reflection is to assess the familiarity and similarity of collected images to prevent design fixation arising from confirmation bias. Our results contribute insight into the nature of confirmation bias in the context of this type of visual research, as well as how designers approach overcoming it. Our findings also indicate the potential for tools to help increase designers' awareness of confirmation bias and to provide streamlined paths to recovery. Finally, we use our findings to motivate and outline a number of promising directions for future research.

Chapter 5

Retaining “Stories” about Graphic Designers’ Visual Research

Findings from the previous chapter suggest that search information recorded during visual research could be used as a memory aid for long-term inspiration seeking. These findings align with the findings of prior research on this subject, which tell us that saved images do not completely satisfy designers’ informational needs [33,62,75]. For example, we know that collections of saved inspirational images do not retain the “stories” behind these images – a combination of information on how the designer found each image and why he/she decided to save it [74,75]. Other examples of information that is not retained by saved images include the design directions that the designer intended to explore, but did not explore because he/she got distracted, and the images that the designer intended to save, but simply forgot to do so [33]. This chapter investigates the idea of retaining “stories”

about visual research by augmenting collections of saved images with rich search history snippets, which we refer to as *search trails*.

This chapter describes the design and evaluation of Prism (Figure 5) – a system that automatically collects all images that a graphic designer interacts with on Google Images and annotates them with the designer's search trails. Prism then acts as an extended image library, where the designer is able to see not only the images they liked and would normally save, but also all the other images that they have inspected, as well as the alternative design directions that they have considered.

To validate Prism's approach, we conducted a two-week-long field study with 11 professional graphic designers. During these two weeks, we asked the participants to use Prism whenever they had to seek inspiration or references on the Internet. To collect feedback about Prism's utility, each participant went through two semi-structured

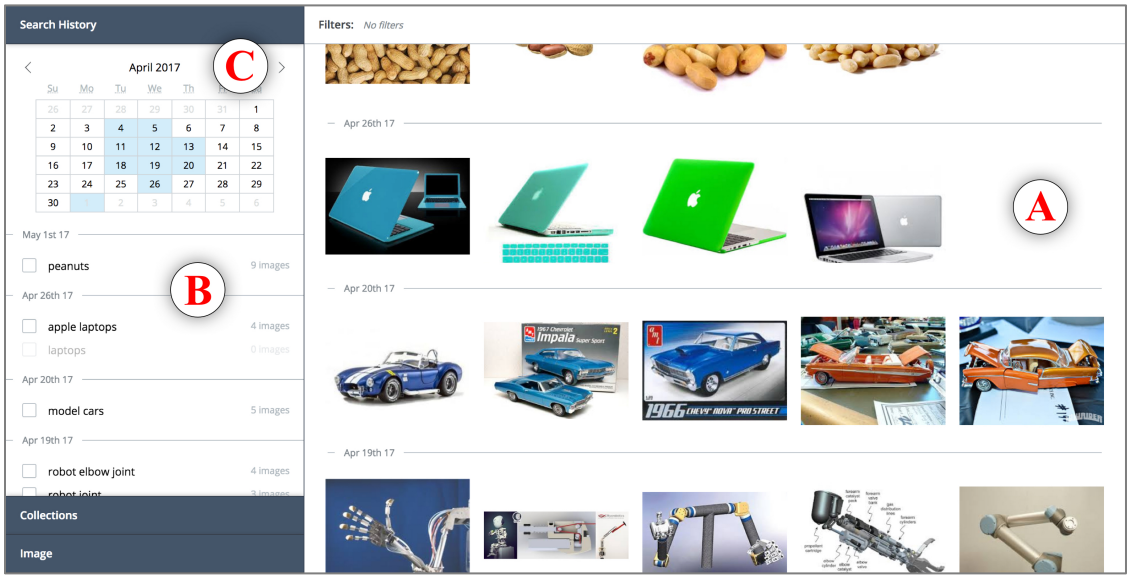

Figure 5. Interface of Prism.
Extended image library (A); index into the library via search queries (B); calendar highlighting days when Prism captured search trails (C).

interviews: one after the first week of the study and the other one at the end of the study. During these interviews we asked the participants to give us feedback about Prism’s features, and about the utility of search trails for their inspiration seeking workflows.

Research described in this chapter has been published and presented at the ACM SIGCHI Conference on Designing Interactive Systems (DIS) in 2018 [20].

5.1 Enhancing Visual Research with Interactive Search Trails

Inspired by existing research in the domain and by our findings from the evaluation of Instory, we investigate the idea of supporting graphic designers’ visual research by recording and displaying search trails associated with each design direction explored during visual research. In this context, a search trail would contain all of the search queries used and all of the images saved for a particular design direction. For example, if a designer decides to include elements of a stormy sky in their design, then they might use the following search trail in their visual research: “cloudy sky”, “sky with clouds”, “stormy sky”, saving one or more images per query. This specific search trail would contain the three search queries and all saved images.

The search trails approach continues our idea of supporting inspiration seeking practices via non-intrusive objective information about the process that does not make any assumptions about designers’ goals.

5.2 Prism

Prism is a system that supports visual research by automatically capturing all inspected images and annotating them with the graphic designer's search trails (Figure 5). The system's name originates from an analogy between a collection of images and the light spectrum: like the visible part of light spectrum, a collection of inspirational images does not reveal much to a designer about their history. Similar to a prism that reveals "hidden" spectral colors of light, Prism reveals "hidden" stories behind images.

5.2.1 Automating Image Capture via Click-Based Inspection

Motivated by our findings from Chapter 4, Prism captures visual research process beyond explicitly saved images. Prism tracks a designer's interaction with images that he/she explores and automatically captures all those images that catch the designer's attention. Specifically, Prism captures all the images that a designer clicks during his/her exploration, so that he/she can see it up-close. The mechanics of bringing up a larger view of an image on mouse click is common for online image repositories (e.g., Google Images and Pinterest) and when evaluating Instory we observed that our participants often clicked images that caught their interest before making a decision on whether to save it or not. We elaborate on potential benefits and drawbacks of this method in comparison to more liberal (e.g., image hover) or conservative (e.g., an explicit save action) ones in the section 5.9.

5.2.2 Annotating Images with Interactive Search Trails

Prism's other main feature is that it saves search trails for each design direction explored. Both prior work and results of our evaluation of Instory in Chapter 4 suggest that designers often lose their train of thought when seeking inspiration or references for their design ideas [33]. The literature and our study results also suggest that over the long term,

designers often do not remember why they saved a certain image [75]. Unlike Instory, which captured only those search queries that were related to the saved images, Prism automatically tracks all search queries that designers use in their visual research, so that the designers can backtrack their process and remember why they pursued a specific design direction.

5.2.3 Collection Overview and Filtering

Prism displays a chronological list of all searches split by days on the left-hand side of the image library (Figure 5 B). Unlike Instory, the chronological presentation of this information was tailored to facilitate its retrieval from a designers’ episodic memory [80]. Specifically, by showing the evolution of a designer’s thoughts over time we aimed to facilitate the recall of the exact look and feel that the designer was pursuing at that moment of time.

Prism also captures rich information about each image, including the date when it was collected, the search query, the link to the search results page where they found the image, the image’s resolution, and a list of most dominant colors. This information could help a designer to recall why they decided to save a specific image.

Given that the automated image collection algorithm implemented in Prism naturally leads to capturing more images than designers would normally save, the system provides a range of navigation and filtering options (Figure 6). As one example, Prism uses the list of search queries as an interactive index into the image library. A designer can click one or more search queries, which filters the image library to only show those images that were collected via the selected queries (Figure 6 A).

Prism also allows designers to use their memory of when they worked on certain design directions to filter their image library. An interactive calendar above the list of all search queries highlights the days in which Prism recorded search trails. Clicking on a day automatically filters the list of search queries and the image library to show only those from the selected day (see Figure 6 A for an example).

Finally, following the participants feedback from our previous study, we included the features of filtering by image color and resolution (see Figure 6 B for an example).

5.2.4 Manual Grouping

Prism also allows designers to manually group images into ‘collections’ (by accessing the “collections” tab in Figure 5). This feature was added in case designers were not satisfied with Prism’s automated organizations.

5.2.5 Implementation

The current implementation of Prism works with Google Images, but ideally such a system could work with various search engines. The decision to switch to Google Images from

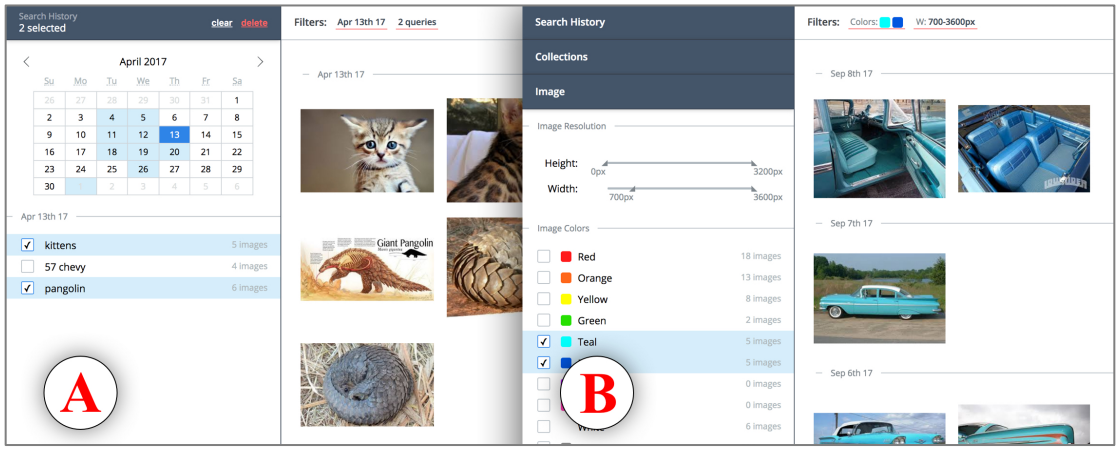

Figure 6. Examples of filtering in Prism.
Filtering by calendar and search queries (A); filtering by image resolution and colors (B).

Pinterest was motivated by our participants’ feedback from the Instory study described in Chapter 4. While our participants in that study found Pinterest sufficient for completing their study task, many of the participants told us that Pinterest does not provide as much variety of images and topics as Google Images. For example, those participants told us that Pinterest is a good source of inspiration for images that relate to arts and crafts, but not as good when it comes to factual knowledge and reference material. We decided to use Google Images for Prism because all our participants considered it to be the most general and versatile source for inspirational images.

Our implementation uses a Google Chrome extension to record search queries and image clicks. Prism allows designers to temporarily disable tracking when, for example, they are doing personal image searches. Clicking the Chrome extension icon in their browser will toggle the tracking on and off. The extension conveys the state of tracking by a blinking red dot (Figure 7).

Prism is built as a web application written in JavaScript using the React library and NodeJS server. The current implementation does not store images, but rather links to images and their thumbnails on Google Images. We chose this approach to minimize storage requirements, which would become particularly important with long-term use.

Figure 7. Prism tracking indicators.
Left: tracking is on (red dot blinks); right: tracking is off.

5.3 Field Evaluation of Prism

We evaluated the utility of Prism in a two-week long field study with 11 professional designers and digital artists. The goal of the study was to see how these professionals might adopt such a tool as part of their workflow and to gain insight on the potential strengths and weaknesses of its main features. The consent form, the recruitment script, and other relevant materials are presented in Appendix E.

5.3.1 Participants

We recruited 12 professional graphic designers and artists (5 female) by advertising the study on Reddit. One participant withdrew part way through the study (for reasons unknown) leaving us with data from 11 participants. Participants were 19-39 years old and had at least 3 years of professional experience (3 participants had over 10 years of professional experience). Collectively participants had experience working in a variety of design-related fields including logo design, print design, web design, branding, and illustration. Participants received a $75 gift card. Of the 11 participants, one had also participated in our evaluation of Instory (described in Chapter 4)

5.3.2 Procedure

Our study consisted of an initial meeting, a two-week use period and two semi-structured interviews. During the initial meeting, we guided the participant through prototype installation and gave a brief demonstration of its features. We then asked participants to use the system for the next two weeks whenever they were seeking inspiration or reference images for their ongoing projects.

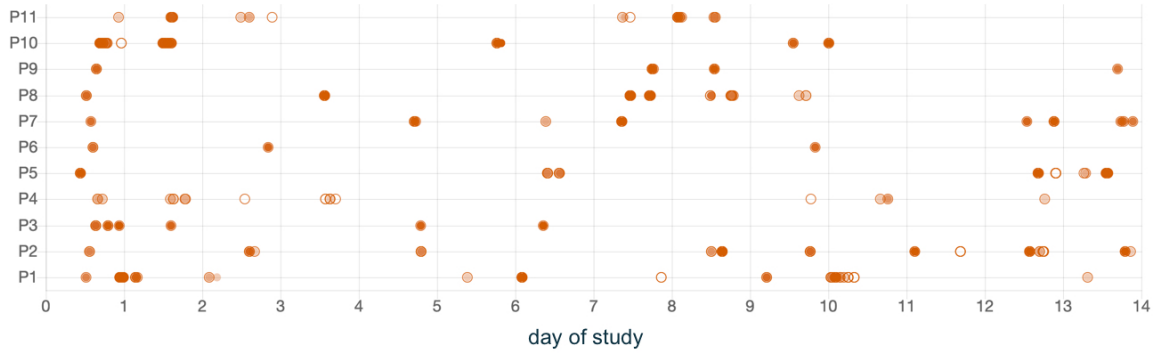

Figure 8. Prism log data, showing participants’ activity.
Filled and empty circles represent captured images and search queries, respectively. Darkness of circles represents the frequency of use, e.g., dark filled circles indicate a high number of captured images.

Our semi-structured interviews with each participant took place in the middle of the study (day 7) and then again at the end of the study (day 14). During the semi-structured interviews, we asked participants to walk us through their experience with Prism thus far. We asked participants to describe the projects they had worked on, the types of images they had been looking for, why they had selected specific images, and to comment on any ways that Prism had helped or hindered their design tasks.

5.3.3 Data Collection and Analysis

We audio recorded the interviews and logged all interactions with Prism. Interviews were transcribed in full.

5.4 Findings: Usage Data

Figure 8 shows that all participants used Prism multiple times throughout the study. On average, each participant interacted with the system for 5.4 days (SE=0.5) and collected 125 images (SE=31.6) across 42 search queries (SE=9.1).

As expected, not all participants used the system to the same extent. The most active participant (P1) interacted with Prism for 9 out of 14 days, during which they collected 397 images over 106 search queries. The least active participant (P4) collected only 13 images across 20 search queries, despite interacting with the system for 7 out of 14 days. More than half of the participants (6 of the 11) used Prism to collect over 100 images and six participants continued to use Prism after the end of the study period.

A couple of participants (P1 and P3) used other image repositories in parallel to using Prism during the study period, as they felt that Google Images did not provide them with the functionality they needed. For example, P3 used specialized websites (e.g., freepik) for retrieving vector images. Two participants also discussed using existing image collections in parallel to using Prism: P2 used images that they received from their client as a part of project requirements and P9 referred to printed collections of inspirational images. Importantly, both expressed desire to upload their existing images to Prism, which indicates the desire to adopt the new workflow that Prism supports.

5.5 Findings: Validating Prism's Approach

As we expected, with Prism, participants indicated that they collected far more images than they would normally through manual curation and annotation mechanisms. In some cases, participants indicated that the difference in their collection practices was extreme, going from rarely explicitly saving images during visual research to the detailed collection that Prism stores. For example, one participant indicated that they would normally not save even a single image during visual research, typically inspecting each image for only a short time. Using Prism was a complete change of workflow, but they were happy about it:

Typically, I just kind of go for something without even moodboarding it. It’s just how I am. But this was a complete change of my workflow, really... but I liked it a lot. (P9)

Participants’ comments revealed that the combination of automated image collection and the search-trail annotation enabled use-cases that would not be possible with either of those two features by themselves. We elaborate on these use cases below, along with participants’ comparisons to existing practices.

5.5.1 Reifying Design Thinking

All participants indicated that the ability to see all the images that caught their attention annotated with the respective search queries allowed them to easily re-create the way they had approached their design tasks. For example, when asked to describe their design approach for a particular design task, P8 used the combination of search trails and the extended library to recall the kinds of images they had been searching for, the rationale behind their searches, properties of unsuccessful directions, and motivations for further searches:

...when I search for just “prince”, I just got like the musician, and I was like “this is not helping”. I am looking for a “fairy-tale prince”, but I don’t like the fairy-tale ones I am looking for. I was like “oh, I don’t know if that is going to read into a Japanese style”, so [I searched for] “prince in Japan” and I got nothing but his performances when he went to Japan. So, I had to be specific, like “prince fairy-tale Japan”. (P8)

This quote illustrates a common theme in our data, which supports our idea of retaining the designer’s train of thought by saving their search trails and all the associated images.

Participants found that capturing their evolution of thought via search trail helped them both to explore the design space by retaining the context surrounding prior ideas and to re-engage with their design tasks after time away. We elaborate on each of these use cases below.

Supporting design space exploration through retained context. Almost all participants said that having all of their images and searches saved and organized allowed them to track and return to alternative design ideas that they had considered during inspiration seeking. For example, in the following quote P11 explains how their habit of rapidly switching between search queries often leads them to losing design ideas that they liked. P11 told us that Prism helped them keep track of all the nuances of their searches that they would typically loose with their existing workflow, and that these nuances gave them the ability to go back to alternative ideas:

> *[When] you develop designs [...] there's like a point A and a point B and there's a lot of points in between...And with the computer, it's really fast to move between those points without really realizing the thought process... And there's something in between that you thought that it was genius. But sometimes it's really hard to remember those things. Like, it's little nuances in design that make design really, really cool. But it's also those little nuances that you forget, because you are not writing them down. You are just quickly searching google. So, if you had an idea, you clicked on an image... and then it's there [saved in Prism]. I think [the image in Prism] will just serve as just a breadcrumb for your mind to go back to that place... (P11)*

Aiding recall and task resumption. Participants also felt that Prism helped them pick up their work after a break. Almost all participants told us that it is often difficult to remember their train of thought after they take a break from working on a project, such as switching to a different project, or even simply going for lunch. Prism provided participants with enough context to remember ideas they had explored and those that had given them inspiration, which helped the participants to get back on track:

> *...if I had multiple searches going at once, it's nice to see exactly where I was in each search, because sometimes you are just searching for too many things and I just get confused with all of it. (P4)*

> *Sometimes I am drawing a character or something, and I stop for lunch, right? So, I close everything and go lunch and go back [to work]. And when I go back, I have to remember what was the search query again. [With Prism] I don't have to remember the search query. [...] It helps me to pick up my work from where I left easier. (P7)*

5.5.2 Serendipitous and Intentional Inspiration for Future Projects

All of our participants mentioned that when seeking inspiration or reference images, they see many images that they think might be useful for them in the future. Validating our approach, participants told us that Prism's automatic image tracking and support for re-finding enabled them to not only collect references for their current project, but also to intentionally capture inspirational images for potential future projects:

> *I started to look for things related to what I want to do in future, not necessarily right now. (P8)*

Participants found Prism's annotation and filtering capabilities to be sufficient for re-finding images from the two-week study period. They also thought that the existing features would be useful for re-finding images after even longer time periods, such as several months. However, a longer-term evaluation would be needed to investigate Prism's ability to support re-finding these types of 'off-topic' inspirational examples long term.

5.5.3 Automated, Objective Organization as an Alternative to Manual Collections

Participant feedback also provides validation for Prism's automatic approach of organizing collected images. Participants told us that manual organization is a tedious process that takes their mind from the task at hand. As they don't want to put much effort into organizing references, their collections often get too disorganized to be useful:

> *...it is like chaos. I don't even want to look at it [my reference folder]! (P1)*

All participants found Prism's organization of images by search trails intuitive and sufficient for their task. In fact, only four participants created any manual collections for their images, whereas most of participants just did not feel need to:

> *It kind of [organizes] everything for me. I guess I could use collections, but I kind of did not see need to. (P3)*

One participant (P8) specifically commented on how they appreciated that the computer didn't try to be too "smart". They emphasized that Prism did just right amount of automatic organization to augment their process, but did not attempt to 'think for them':

> *...what's nice about this tool is that it does not think for me, you know? I still have to put in my creative process to it, because I know specifically for this project, it's*

intended for children, [...] so I need to keep that in mind, and that’s not something your tool can provide. That’s only something [you can get] through training or just considering. I like that it doesn’t supplement my knowledge and [does not] think for me on what that looks like, or what that inspires me to do. (P8)

5.6 Findings: Unexpected Use Cases

In addition to validating our initial insights regarding Prism’s potential to support designers with visual research, our interview data also revealed a number of use cases that we had not anticipated. These include using Prism to reflect on design alternatives, establish a shared vocabulary with clients, and identify themes in vaguely defined design directions.

5.6.1 Reflecting on and Presenting Design Alternatives via Custom Views

Participants felt that Prism’s filtering capabilities provided them with customized views of their collections, and that these customized views enabled a range of productive reflection activities.

Gaining a holistic perspective. A few participants told us that they liked the ability to view the entire collection to gain a holistic perspective on their alternatives and ideas in the early stages of the design process. For example, P8 mentioned how this overview allows them to “step back and look at things”:

I really feel I get a fuller picture, because I am literally stepping back and looking at things. (P8)

Creating on-demand mood boards. When focusing on a specific idea, or when presenting potential ideas to a client, participants filtered their images by the search queries that they

felt best represented the specific idea or theme. This allowed the participants to create custom ad-hoc mood boards with minimal effort:

> *I can click [queries] and then I can only show those [images]. I don't want to show [the client] different things that I don't necessarily want to show them, like the [images of] bags. That's a personal thing. So, I can do that... it's like automatic Pinterest board for my client. (P1)*

> *I used it as a little library, to just turn things on and off. So, I know I want "Japanese foxes" and "red", I also want just like regular looking foxes [is checking respective queries in filters] So, just having like my own little pin-board [...] is really nice, just keep going back and forth while I am working. (P8)*

Soliciting feedback on alternative directions. Two participants, P7 and P10, told us that Prism's features helped them to capture the full breadth of their visual research and that they were able to get more feedback from their colleagues than they would typically obtain. For example, P7 told us that typically they present and discuss a set of images that represents their favorite direction. Prism allowed P7 to show his/her colleagues all the ideas that they had considered in visual research, which gave their colleagues the "bigger picture" and elicited more constructive feedback:

> *I showed [my colleagues] the whole search. So, that was really useful, because I could save what I want locally and show to everyone what I actually searched for. And they could give their opinions on that [...] if I had showed them only what I have saved, like what I always do, because, you know, you don't save everything you see, they would*

have seen only my tastes for what I wanted. When they saw every image that I had searched for, it gave them the bigger picture. (P7)

Another participant suggested that Prism could also be used in the same way to get feedback from a client.

These sentiments align with prior work suggesting that presenting multiple design alternatives elicits better feedback and can lead to better final results [18]. Design ideas are also often best presented by considering the space of alternatives in combination [7,29]. The fact that Prism supports these practices in a light-weight manner is encouraging.

5.6.2 Establishing a Shared Vocabulary

Several participants saw Prism as a tool that could help them better communicate with their clients about the nature of the design task. Participants told us that their clients often don’t know how to verbalize the look or the style that they are going after because they don’t know the proper stylistic vocabulary. In the quote below, P11 describes how it is difficult to understand what their clients mean. For example, if their clients says “brick wall”, they might be imagining a picture that is different from a picture that P11 would associate with a “brick wall”:

So, if a client says “hey, can you put a brick wall here?” [...] they [might] have a different idea than what I do of what that means. [That is] a very subjective request, where they may not know how to put that what they’re asking for into words. (P11)

P11 suggested that giving Prism to their clients could help them with this problem. Specifically, P11 indicated that having a client’s visual research in Prism would allow a

graphic designer to see not only images and styles that caught the client's attention, but also the words the client had used to find those images.

5.6.3 Refining Design Directions via Thematic Analysis of Liked Images

For two participants, P2 and P9, Prism's automatic image and trails collection helped them to refine their design directions in cases when they did not have a clear vision of what they wanted to do. Specifically, the participants dived into inspiration seeking with a vague and ill-formed idea of the design direction. Then, they reviewed the extensive collection of all the images that they liked for commonalities and subtle differences in style, composition, colors, etc. The participants then used this information to refine their initially vague design direction:

> *...you can see there's a lot of gradients, there's a lot of this bluish-green, it's pretty common in all of them. You can just visually see that all over the place. Those are the things I did not notice on google images when I was looking at stuff earlier... the divine effect on the edges that a lot of these have... So, just being able to spot common themes within a type of imagery you are searching... (P2)*

> *[Looking at all images I clicked on] was a really quick way to see what I have liked and what I kind of want to do in terms of direction. [...] So, [Prism] helps me a lot to identify what I wanted to go for and just and do it a lot quicker. (P9)*

These comments suggest that the high volume of collected images and search trails was the primary factor that enabled them to refine their initially vague design direction. Achieving a similar effect via manual collection would be tedious.

5.7 Findings: Weaknesses and Open Issues

While participants were generally enthusiastic about Prism’s approach, with a number of participants asking to continue using it after the conclusion of the study, they did note some important considerations for future iterations.

5.7.1 Extending Prism to a Variety of Image Repositories

One of the main limitations that participants cited was that the tool only supported visual research using Google Images. Participants indicated that they typically use a variety of platforms, taking advantage of their strong sides, for example, using Freepik to find vector images, Pexels to find free-to-use stock photos, or Pinterest for its recommendation engine. A full-featured tool would therefore have to include support for image search via a variety of image repositories and image search engines to fully integrate with designers’ current visual research practices.

5.7.2 Extending Search Trails with More Information

One participant (P8) wanted to add even more information to their search trails. They said that they often make notes in their reference books about why they like a certain image. They wanted the system to allow them to make similar notes for collected images, so that they could make a memo for themselves about why they saved that image:

> *I would write small details on why I wanted to save this, or why I clicked this, or what something I want to remember about this. [...] [For example] for some of the stickers, I would talk about the use of texture here [...] that’s why I save those images, not that I want to draw a daruma in this way. (P8)*

P8 was worried that without these additional notes, it would be difficult to remember their design thinking if they came back to the project after an extended period of time:

> *... a few years down the line, I would not remember why I saved these polka dots. [...] When I am going to look at my [notes] even later down the line, I would know that I want one of my socks to be polka-dotted. (P8)*

This aligns with prior findings about designers wanting to save "stories" behind images and artifacts [75]. The extent to which the search trails alone would enable these stories to be re-created in a more distant future is a question for future work.

5.7.3 Discomfort with Continual Tracking

While participants enthusiastically embraced the tool for its benefits, some participants did feel uncomfortable with the idea of a tool watching their every step. One participant continued to feel slight discomfort even at the end of the study:

> *It felt a bit odd, that the extension was following me, in a way. Like it was monitoring what I'm doing [...] its tracking me. I know that's the point. But it was a little weird. (P4, final interview)*

A few other participants felt uncomfortable at first but got more comfortable with the system after they used it for a while. For example, one participant was initially hesitant to have so many images stored, and spent the first week of the study trying to manage the size of their collection via selective inspection:

> *One thing that I realized is that I am clicking less, because I don't want everything to go to database [...] I don't want every image to be recorded, even though I search for*

> *image. Because sometimes the thumbnail is very small and I want to see the bigger picture to see if I really like it. But if I click the thumbnail, it gets saved automatically. And I don't want that. (P7, mid-study interview)*

After the mid-session interview, this participant decided to give it a try and began seeing advantages:

> *...after last week, when I did not want to click everything, because it actually gets saved, I actually tried to not mind so much for that and it actually helped me, because after I started clicking on everything that I liked, even if I don't like it so much, I ended up using [Prism] as an extended folder [...] for reference [images] that I would not have saved otherwise. (P7, final interview)*

Participants liked the fact that they could disable tracking when doing personal image searches, but this feature had the disadvantage that they would then sometimes forget to enable it when starting a work task.

5.8 Discussion and Future Work

Our study findings validate a number of principles behind Prism's design. Our findings also reveal unanticipated ways in which designers used search trails collected during their visual research to promote effective exploration and communication of their design ideas. In this section, we reflect on our findings and some potential avenues of future research.

5.8.1 Alternative Image Inspection Mechanics

There is a rich design space surrounding detecting when a designer expresses interest in an image. For example, some of our field study participants suggested using image hover

instead of image click. However, as P7's quote in the previous section indicates, there will likely be tensions between retaining images more liberally and the size of the resulting image collections. One direction for future research is to investigate and compare the implications of different image inspection mechanics for the designers and their image collections.

5.8.2 Image Collections as Interactive Sketchbooks

Our findings suggest that designers used their Prism collections and its filtering capabilities to help guide their exploration through the design space, communicate ideas with others, and synthesize higher-level properties of initial ideas and concepts. Many of these activities have been emphasized in prior design literature (e.g., [3,7,18]), which has often advocated approaches like sketching to enable rapid idea generation, collection and communication [7]. Our findings indicate that effectively indexed and lightweight image collections from visual research can potentially act as a type of interactive image-based "sketchbook" for design alternatives.

5.8.3 Continual Tracking

Our findings suggest that collecting all images helps designers with visual research, however, some participants felt uncomfortable being continuously tracked by our tool. Anticipating such a scenario, we implemented a method to turn tracking off for when a participant is doing research that they do not want to be tracked by Prism. However, many participants forgot to turn the tracking on for work sessions. This led them to miss out on some of the images in their collections and explains some of the usage variability we saw in our log data. We also saw at least one participant being somewhat hesitant to click on images for further inspection, which was their way to curate their collection. Future

research should, therefore, consider tradeoffs associated with requiring the user to give a stronger “interest” signal, for those users who want greater curation control. While providing a feature to disable tracking was partly important in the context of data collection for a study, future research will also have to consider ways to make the system status salient, perhaps also considering targeted reminders (when appropriate) of the system’s status.

Future research could also investigate ways to harmonize continuous tracking and the retention of personal search information. For example, instead of a permanent on/off switch for tracking, Prism could prompt designers whether they want their history to be tracked at the beginning of each image-seeking session. This session-based control strategy could help designers avoid the loss of information due to human error (e.g., forgetting to turn the system on or off), although constant prompts from the system might become frustrating over time. Alternatively, Prism could allow designers to curate their trails, for example by allowing them to remove unwanted images and queries from Prism. The ability to remove personal search information from Prism can make designers more comfortable with tracking, but excessive amounts of curation might break the continuity of transitions in search trails, making it more difficult to re-create one’s thought process.

5.8.4 Additional Search Trails Annotations

Our participants saw a lot of value in capturing search trails as a supplement to their images, however, at least one participant wanted to add additional notes about their rationale behind image collection. While adding such a feature would be a straightforward extension, it also raises the question of what other forms of information could potentially be captured and displayed along images. For example, it might be possible to also collect properties of the

designer's workspace, such as files open or screenshots of designs being generated in certain relevant software applications.

5.9 Summary

This chapter investigated a way of retaining valuable information about graphic designers' visual research via interactive search trails. We developed Prism, a system that captures the evolution of a graphic designers' thought process during visual research by retaining all the search queries the designer has used and all the images the designer considered when researching a specific design direction on Google Images.

After evaluating Prism in a two-week-long field study with 11 professional graphic designers, we discovered that Prism helped our participants to capture their inspiration seeking far beyond the collections of images that they found most inspirational. Specifically, Prism helped our participants to capture serendipitous inspiration for future projects and to retain and to reflect on the alternative design ideas that they had considered but never saved. Our findings also speak to a number of unanticipated benefits of using Prism in everyday work, such as facilitating constructive feedback from colleagues and establishing a common vocabulary with clients.

Chapter 6

Supporting Graphic Designers' Inspiration Seeking Practices with Software Tools

Chapters 3 to 5 demonstrate examples of how software tools can support inspiration seeking practices within two different domains of graphic designers' activity: software learnability and ideation. In this chapter, we reflect on our findings across these two domains and synthesize a set of design considerations for supporting graphic designers' inspiration seeking practices.

6.1 Support Capture and Re-Finding of Accidental Inspirational Discoveries

Our findings suggest that graphic designers' inspiration seeking practices often rely on accidental inspirational discoveries which happened in the past. Serendipitous inspiration was a common theme across all of our system evaluation studies. For example, our

participants who used Switter mentioned accidentally stumbling across a tutorial which was not necessarily satisfying their information need at that time, but which they knew would be useful some time in future. Similarly, some of our participants from the Prism evaluation mentioned that they often stumble across images which might not be useful for their current project, but which give them inspiration for potential future projects. From the words of our participants, they often struggle re-finding these accidentally discovered inspirational images because they usually do not save them at the time of discovery.

Our findings suggest that software systems can help designers to use already discovered inspiration by capturing and aiding with re-discovery of previously accidentally discovered inspirational materials. For example, Prism's automated image capturing helped our participants to save unexpectedly discovered inspirational images even if they were not directly related to the participant's task. Prism then allowed our participants to use their memories about their process to re-find these images later (e.g., by remembering which day they found an image or what exactly they were searching for when they found it).

6.2 Build a Software Tool to Augment a Designer's Expert Judgement

Our findings speak to the difficulty in knowing what a designer wants to achieve when seeking inspiration. For example, our initial investigation of ad-libitum exploration in Chapter 3 showed that exploration could be considered successful even if a designer did not find anything new to learn, because the designer might perceive it as a validation of their knowledge level. Similarly, our investigation of Instory in Chapter 4 showed that a designer might deliberately seek inspiration from specific images that the designer already

knows about, even though prior work tells us that drawing inspiration from novel images might lead to more creative results. In both cases, inferring a designer's goals based on general criteria might conflict with the designer's actual inspiration seeking objective and might lead to frustration and the loss of trust in the system.

Our findings suggest that software could help designers to achieve their diverse inspiration seeking objectives by augmenting the designers' experience in those areas where a computational approach is more efficient than human effort. For some tasks, such as searching by keywords, saving information to a local storage, or extracting information from software tutorials, computers are often faster and more reliable than humans. Thus, one way in which systems can support graphic designers' inspiration seeking practices is by automating a set of routine sub-tasks within a designer's process, so that the designer can focus on the creative components of their process. Below we describe one example of such an approach, which we implemented in our systems (Switter, Instory, and Prism).

Recording and Displaying Obscured Objective Information About Inspirational Materials or About a Designers' Process

Our investigations of ad libitum exploration in Chapter 3 and of visual research "stories" in Chapter 5 suggest that designers could benefit from a system that reveals some obscured information about inspirational materials or about a designer's process. For example, in the case of ad libitum exploration, command summaries that Switter displayed under each tutorial helped our participants to assess the inspirational value of the tutorials. In the case of visual research, Prism automatically captured and displayed search trails, which helped our participants to keep track of alternative exploration paths. Both detailed command summaries and search trails represent obscured information – information which might not

be easily accessible without additional tools like Switter and Prism. This obscured information was also objective in nature, meaning that it was a characteristic of a process (e.g., as in the case of Prism) or of a specific inspirational material (e.g., as in the case of Switter) and did not carry any assumptions about a designers' goal.

Our findings suggest that revealing obscured objective information about a designer's process or about inspirational materials is one of the ways in which a software tool can augment a designer's expertise in a non-intrusive way.

6.3 Respond to Designers' Individualistic Goals by Supporting a Variety of Exploration Strategies

Every participant in our field evaluations of Switter and Prism used these systems for their own projects and tasks with unique design goals. One of the reasons why our systems allowed our participants to fulfill their individualistic goals was because both Switter and Prism provided a variety of ways to explore the information captured and displayed by the system. In the case of Switter, for example, our participants were able to explore tweets via a chronological feed, via the interface replica, and via activity indicators. Combining these three exploration methods allowed our participants to adjust Switter to their unique goals, such as finding use cases for a specific Photoshop command, juxtaposing various Photoshop tools to compare their use, finding synergies between different commands, etc. Similarly, Prism supported our participants' individualistic goals by allowing them to easily switch between to custom mood boards for when they wanted to focus on a specific a design direction and to a holistic overview of their entire process for when they wanted to see the bigger picture.

Our findings suggest that providing flexibility in how designers can access, organize, and explore the information provided by a software tool could make the tool more suitable for supporting a variety of individualistic inspiration seeking goals.

6.4 Investigate the Variety of Novel Inspiration Seeking Goals Enabled by a System

Our findings suggest that a system providing graphic designers with new ways to explore inspirational material might inspire these designers to pursue novel inspiration seeking goals, which would not be possible without using the system. For example, Switter enables graphic designers to explore Photoshop tutorials from the perspective of Photoshop commands used in each tutorial. During the evaluation, we observed our participants using Switter to pursue learning goals which we did not anticipate, such detecting and to improving personal weaknesses with Photoshop and discovering synergies between Photoshop commands. Similarly, using Prism motivated some of our participants to reflect on their design process together with their colleagues and clients.

Investigating these novel goals and unexpected use cases might help researchers uncover new insight about graphic designer's practices and discover new ways in which they could support these practices. For example, our findings from the Switter evaluation suggest that one of the ways in which graphic designers discover new software techniques is by investigating synergies between various software commands. Researchers can therefore leverage this finding to design strategies for helping graphic designers compare and invent new software techniques.

6.4.1 Designing System Evaluations to Investigate Unexpected Use Cases

Our experience suggests that one of the ways to capture novel inspiration seeking goals and unexpected use cases is to increase ecological validity of system evaluation studies. Arguably one of the main reasons why our Switter and Prism evaluations revealed unexpected use cases was the fact that our evaluations were open-ended, longitudinal, and were conducted in the field. We speculate that by giving our participants freedom in choosing their own design task, we increased our participants' intrinsic motivation for using a new tool. Similarly, we speculate that by giving our participants time to experiment with various ways to use our systems, our participants managed to discover unique ways in which our systems fit in and augment their personal work practices. Seeing how our systems fit into the designers' work practices allowed us to learn about those inspiration seeking goals which graphic designers want to pursue but lack the means to do so.

Our experiences also suggest the value of using mixed-method designs for soliciting participants' feedback about unexpected use cases. For example, both Switter and Prism evaluations were designed as explanatory mixed-method studies: during the study we collected data on how our participants interact with the system and discussed various interaction patterns which emerged from this data during mid-study and post-study interviews. Using interaction data allowed us to help our participants to recall specific challenges they had in their work and to narrate how our systems helped (or did not help) them to solve these challenges.

6.5 Summary

This chapter presents a synthesis of our findings and experiences into a set of general suggestions on how to build and study inspiration-support software systems for graphic designers. Our findings suggest that when building such a system, researchers should consider supporting capture and re-discovery of serendipitous inspiration, as graphic designers often rely on their past serendipitous discoveries when seeking inspiration for their projects. Our findings also suggest that in order to avoid interfering with a designer's inspiration-seeking goals, the system should aim to augment the designer's expert judgement, for example by recording and displaying obscured information about the designers' process. Finally, our experiences suggest that enabling a variety of ways in which graphic designers can explore information provided by a system could help researchers to better understand the types of graphic designers' inspiration-seeking goals.

This chapter also shares a retrospective on our experience evaluating our systems. Specifically, we argue that increasing the ecological validity of evaluations (e.g., by conducting open-ended longitudinal field studies) and using a mixed-method evaluation design might help researchers to learn about novel inspiration seeking goals enabled by introducing a system.

Chapter 7

Conclusions and Future Work

Online inspiration-seeking is an integral part of a graphic designer's work. By exploring images on the Internet, graphic designers deepen their knowledge of their target domain and get inspiration for design ideas. Additionally, graphic designers pro-actively seek to improve their expertise with complex design software, such as Photoshop, Illustrator, Sketch, etc. Both practices can be characterized as exploratory search [56] aimed to teach a designer something new, but without having upfront knowledge of the learning goal.

The goals of this thesis were to explore the challenges that graphic designers face when seeking inspiration on the Internet and to investigate the potential for software systems to support graphic designers' inspiration seeking practices. This chapter summarizes the contributions of this thesis to the Human-Computer Interaction (HCI) field and discusses promising avenues for future research.

7.1 Thesis Contributions

This thesis makes empirical and artifact contributions to the HCI field in the areas of software learnability and ideation, as well as contributions to the design and evaluation of inspiration-seeking support systems.

7.1.1 Contributions to the Software Learnability Domain

Our investigation into the area of software learnability firstly contributes a description of *ad libitum exploration* – an inspiration-seeking strategy that involves habitual monitoring of software-learning materials posted on social media. Chapter 3 contributes the insight that designers engage in ad libitum exploration to seek unexpected and novel ways to use design software, such as Photoshop or Sketch. Our findings also demonstrate that designers engage in this type of inspiration-seeking notwithstanding their low expectations for discovering inspirational software techniques. Finally, our results speak to a number of challenges faced by designers who engage in ad libitum exploration, such as the difficulty of assessing the inspirational value of software learning materials, the challenge of keeping up with the fast information pace on social media, and the lack of means to filter information.

Our work in the software learnability domain also contributes the design and evaluation of Switter – a system which we designed and built to support ad libitum exploration on Twitter (Chapter 3). Switter's design contributes a novel method of browsing and filtering software learning materials via an interface replica of a target design software (in our case – Photoshop). Our field evaluation demonstrated how tools like Switter can be integrated into graphic designers' work practices. Additionally, our results showed that augmenting

designers' ad libitum exploration capabilities with software systems can motivate new inspiration-seeking goals and therefore facilitate the discovery of new software techniques.

7.1.2 Contributions to the Ideation Domain

Our exploration of inspiration-seeking practices in the ideation domain first describes the practice of *immediate reflection* on images saved during visual research – a type of reflection that takes place during or immediately after a visual research session. Our formative study results from Chapter 4 demonstrate that graphic designers engage in immediate reflection to assess and refine future exploration directions, to gauge the idea saturation of an image collection, and to check for potential fixation effects due to designers' confirmation bias. Our results contribute insight into the nature of confirmation bias in the context of this type of visual research, as well as how designers approach overcoming it. Chapter 4 also contributes the design of Instory – a system we built to use as a technology probe in our exploratory study. Instory implements a unique approach for exposing relevant search information alongside inspirational images saved on Pinterest. Our study results indicate a number of potential ways in which this search information could assist designers with inspiration-seeking related challenges, such as correcting fixated image collections and task resumption.

Chapter 5 contributes the design and evaluation of Prism – a system we built to help designers retain the types of information about their visual research that they commonly forget to save. Some of the examples of such information are the alternative design directions explored during visual research, variations of attempted search queries, and serendipitously discovered inspirational images. Building on findings from our exploratory study in Chapter 4, Prism contributes the novel approach of automatically recording and

displaying designers' relevant search and exploration activity. Our field evaluation showed that Prism allows designers to remember their train of thought, including their motivations and goals for exploring certain design directions. Our evaluation results also demonstrated that Prism can help designers resume visual research after a break and re-find serendipitously discovered inspirational images.

7.1.3 Contributions to the Design and Evaluation of Inspiration-Seeking Support Systems

Based on our experience in designing and evaluating systems in two different areas of designers' practice, Chapter 6 contributes a set of general guidelines for designing and evaluating systems which support graphic designers' inspiration-seeking practices. Specifically, we suggest that systems should aim to support the capture and re-discovery of serendipitous inspiration, augment a designer's expertise, and respond to designers' individualistic inspiration-seeking goals. We also suggest that system evaluations should aim to not only validate the approach implemented in a system, but also to investigate the variety of novel inspiration-seeking goals that might be enabled by the system.

7.2 Thesis Limitations

This thesis explores three different inspiration-seeking practices and investigates ways to support them with software tools. While our choice of research methods and participant recruitment strategies allowed us to answer our open-ended research questions, they also introduced certain tradeoffs which we acknowledge in this section.

Potential effects of sampling bias. Our strategy for recruiting professional graphic designers to participate in our studies heavily relied on social media, such as Reddit and

Twitter. While this strategy increased our reach to the specialized target audience, it also might have introduced some sampling bias. Specifically, since the majority of our participants were active users of social media, more investigation needs to be done to verify whether our findings generalize to a broader audience.

Breadth vs. depth of exploration. This thesis explores three different inspiration-seeking practices that relate to design ideas and software techniques. This breadth-focused approach allowed us to juxtapose approaches for supporting different types of designers' inspiration search and to synthesize a set of general design considerations for how software can support online inspiration-seeking. However, the tradeoff of such an approach is that a more in-depth evaluation of our systems was outside of the thesis scope. For example, one can imagine a series of follow-up studies for Prism that would investigate the utility of individual features of the system as well as compare Prism's impact to that of more traditional image-collecting methods, such as saving images to bookmarks.

Qualitative vs. quantitative evaluations. This work predominantly relies on formative studies and qualitative data analysis methods. These methods allowed us to explore the types of use cases that our participants found for our tools (which aligned with our research questions). However, they did not allow us to compare the utility of different features within our tools and to measure the efficacy of the inspiration-seeking strategies enabled by our systems. Future work could investigate different ways to quantify a tool's support for inspiration-seeking practices, so that one could compare the utility of different inspiration-seeking support approaches or the features within the same tool.

7.3 Future Research Directions

7.3.1 Quantifying Support for Inspiration-Seeking Practices

Due to the open-ended nature of our research questions, this thesis focused on the qualitative evaluation of inspiration-seeking support systems. These evaluations helped us assess the feasibility of supporting inspiration-seeking practices with software tools and to investigate the types of support that these tools can provide. The next step in in this line of research would be to evaluate the performance of such systems, for example, to find optimal approaches for supporting specific inspiration-seeking practices. However, the variety of designers' unique inspiration-seeking goals makes it difficult to find a quantitative metric for measuring how well a tool supports an inspiration-seeking practice. This difficulty aligns with general challenges of quantifying the performance of creativity support tools [76,77] and exploratory search support systems [85]. Although the HCI literature describes methods for measuring support for creative tasks (e.g., Creativity Support Index [13]), the goals of inspiration-seeking tasks that we investigate in this thesis focus on acquiring knowledge rather than on creating an end-result. Finding ways to adjust methods like the Creativity Support Index to measure success of supporting inspiration-seeking tasks remains an open question.

7.3.2 Tailoring Support to a Designer's Goals by Incorporating Designers' Feedback

Given that designers can have different inspiration-seeking goals, a system could better adjust to these goals by incorporating a designer's feedback. For example, future work could investigate how to prioritize exploration paths based on a designer's personal goals and preferences. In the case of Switter (Chapter 3), the system could potentially allow

designers to set up pre-defined command filters and to switch between them. This could allow designers to quickly focus their exploration on those command combinations which they are most interested in. Reflecting on these command combinations could provide designers with additional insight about their software knowledge. Similarly, Prism could allow its users to adjust the sensitivity of image capturing mechanics based on specific websites or specific inspiration seeking goals. For example, a designer could configure Prism to capture only explicitly saved images when seeking for a specific image, and to capture all the observed images when doing a broad exploration of the design space.

7.3.3 Augmenting Inspiration-Seeking with Social Trails

This thesis investigated ways to support designers' inspiration-seeking practices using characteristics of the inspirational material and of designers' own inspiration-seeking strategies. Future research could extend this approach by allowing designers to share their inspirational experiences and discoveries with each other. For example, one could imagine Switter showing designers the command synergies that other designers found when using Switter. Similarly, systems like Instory could suggest alternative exploration paths learnt from other designers to facilitate exploration of unfamiliar images. Such augmentations with social trails could help designers to explore unfamiliar areas of the design space, which could potentially decrease fixation effects. Despite the potential merits of enabling exploration via this social dimension, it is not clear what information designers would find useful for exploring inspirational material and which parts of their inspiration-seeking strategies and experiences they would be willing to share.

7.3.4 Supplementing Designers' Expertise with Machine Intelligence

Our results lend support to an approach that leaves the high-level thinking to the designer and has a healthy respect for expertise – computers are particularly good at tracking, whereas the designers are most capable of assessing the suitability of information for their goals. At the same time, there are a number of ways that embedding machine intelligence could potentially further augment designers' inspiration-seeking practices. For example, systems like Instory could potentially use computational vision algorithms to highlight particularly distinct images inspected for a given query (e.g., an image showing an animal among the images that only show plants). Informing designers about similarities, and grouping images with similar elements, might provide designers with additional cues about potentially fixating image collections. With systems akin to Prism, a similar approach could be used to detect and highlight serendipitously discovered inspirational images, as they might stand out among other images retrieved from the respective search query. As another example, a system could use machine learning to provide more advanced filtering capabilities, such as filtering by the types and the number of objects shown in images, or filtering by the image composition.

7.3.5 Generalizing to Long-Term Use and Other Design Domains

This thesis investigated the use of inspiration-seeking support systems over the period of one and two weeks. However, we do not know whether our results would generalize to long-term use, such as several months or years. In addition to the generalization of our results, long-term use raises a number of new open questions, such as whether inspiration-seeking support systems can be used for a long-term reflection on one's inspiration seeking strategies. For example, when used over a long period of time, inspiration-seeking support

systems could allow designers to investigate which levels of specificity of search queries tend to result in finding more inspirational images, or what types of search strategies allow them to capture more serendipitous inspiration.

Finally, we have evaluated the utility of inspiration-seeking support tools in graphic design, however more research is needed to investigate whether our approaches for supporting inspiration-seeking practices would also be beneficial in other design domains, such as animation, interior design, character design, etc.

References

1. Lada A. Adamic, Jun Zhang, Eytan Bakshy, and Mark S. Ackerman. 2008. Knowledge Sharing and Yahoo Answers: Everyone Knows Something. In *Proceedings of the 17th International Conference on World Wide Web*, 665–674.

2. Nikola Banovic, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2012. Waken: Reverse Engineering Usage Information and Interface Structure from Software Videos. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*, 83–92.

3. Mark Baskinger. 2008. Pencils Before Pixels. *ACM Interactions* 15, 2: 28–36.

4. Scott Bateman, Jaime Teevan, and Ryen W White. 2012. The Search Dashboard: How Reflection and Comparison Impact Search Behavior. In *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems*, 1785–1794.

5. Nathalie Bonnardel. 1999. Creativity in Design Activities: The Role of Analogies in a Constrained Cognitive Environment. In *Proceedings of the 3rd Conference on Creativity & Cognition*, 158–165.

6. Andrea Bunt, Patrick Dubois, Ben Lafreniere, Michael A. Terry, and David T. Cormack. 2014. TaggedComments: Promoting and Integrating User Comments in Online Application Tutorials. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, 4037–4046.

7. Bill Buxton. 2010. *Sketching User Experiences: Getting the Design Right and the Right Design*. Morgan Kaufmann, San Francisco, CA, USA.

8. Tony Campbell, Justin Matejka, and George Fitzmaurice. 2010. Chronicle: Capture, Exploration, and Playback of Document Workflow Histories. In *Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology*, 143–152.

9. Carlos Cardoso and Petra Badke-Schaub. 2011. The Influence of Different Pictorial Representations During Idea Generation. *Journal of Creative Behavior* 45, 2: 130–146.

10. Carlos Cardoso and Petra Badke-Schaub. 2011. Fixation or Inspiration: Creative Problem Solving in Design. *Journal of Creative Behavior* 45, 2: 77–82.

11. John M. Carroll and Mary Beth Rosson. 1987. *Paradox of the Active User*. MIT Press.

12. Joel Chan, Steven P. Dow, and Christian D. Schunn. 2015. Do the best design ideas

(really) come from conceptually distant sources of inspiration? *Design Studies* 36, C: 31–58.

13. Erin Cherry and Celine Latulipe. 2014. Quantifying the Creativity Support of Digital Tools through the Creativity Support Index. *ACM Transactions on Computer-Human Interaction* 21, 4: 1–25.

14. Parmit K. Chilana, Andrew J. Ko, and Jacob O. Wobbrock. 2012. LemonAid: Selection-Based Crowdsourced Contextual Help for Web Applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1549–1558.

15. Evangelia G Chrysikou and Robert W Weisberg. 2005. Following the wrong footsteps: Fixation effects of pictorial examples in a design problem-solving task. *Journal of Experimental Psychology: Learning Memory and Cognition* 31, 5: 1134–1148.

16. Tao Dong, Mira Dontcheva, Diana Joseph, Karrie Karahalios, Mark Newman, and Mark Ackerman. 2012. Discovery-Based Games for Learning Software. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2083–2086.

17. Brian Dorn and Mark Guzdial. 2010. Learning on the Job: Characterizing the Programming Knowledge and Learning Strategies of Web Designers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 703–712.

18. Steven P. Dow, Julie Fortuna, Dan Schwartz, Beth Altringer, Daniel L. Schwartz, and Scott Klemmer. 2011. Prototyping Dynamics: Sharing Multiple Designs Improves Exploration, Group Rapport, and Results. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2807–2816.

19. Patrick Dubois, Volodymyr Dziubak, and Andrea Bunt. 2017. Tell me more! Soliciting reader contributions to software tutorials. In *Proceedings of the 43rd Graphics Interface Conference*, 16–23.

20. Volodymyr Dziubak and Andrea Bunt. 2018. Prism: Enhancing Graphic Designers' Visual Research with Interactive Search Trails. In *Proceedings of the ACM Conference on Designing Interactive Systems*, 1283–1294.

21. Volodymyr Dziubak, Patrick Dubois, Andrea Bunt, and Michael Terry. 2016. Switter: Supporting Exploration of Software Learning Materials on Social Media. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, 1209–1220.

22. Charles M. Eastman. 1969. Cognitive Processes and Ill-Defined Problems: a Case Study From Design. In *Proceedings of the International Joint Conference on Artificial Intelligence: IJCAI*, 669–690.

23. Michael Ekstrand, Wei Li, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2011. Searching for software learning resources using application context. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, 195–204.

24. Jennifer Fernquist, Tovi Grossman, and George Fitzmaurice. 2011. Sketch-Sketch Revolution: An Engaging Tutorial System for Guided Sketching and Application Learning. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, 373–382.

25. James Fogarty, Desney Tan, Ashish Kapoor, and Simon Winder. 2008. CueFlik: Interactive Concept Learning in Image Search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 29–38.

26. Adam Fourney, Ben Lafreniere, Parmit Chilana, and Michael Terry. 2014. InterTwine: Creating Interapplication Information Scent to Support Coordinated Use of Software. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, 429–438.

27. Adam Fourney, Ben Lafreniere, Richard Mann, and Michael Terry. 2012. "Then Click OK!" Extracting References to Interface Elements in Online Documentation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 35–38.

28. Adam Fourney, Richard Mann, and Michael Terry. 2011. Query-Feature Graphs: Bridging User Vocabulary and System Functionality. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, 207–216.

29. Bill Gaver and Heather Martin. 2000. Alternatives: Exploring Information Appliances Through Conceptual Design Proposals. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, 209–216.

30. Tovi Grossman and George Fitzmaurice. 2010. ToolClips: An Investigation of Contextual Video Assistance for Functionality Understanding. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1515–1524.

31. Tovi Grossman, George Fitzmaurice, and Ramtin Attar. 2009. A Survey of Software Learnability: Metrics, Methodologies and Guidelines. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 649–658.

32. Yasunari Hashimoto and Takeo Igarashi. 2005. Retrieving Web Page Layouts using Sketches to Support Example-based Web Design. In *Proceedings of 2nd Eurographics Workshop on Sketch-Based Interfaces and Modeling*.

33. Scarlett R. Herring, Chia-Chen Chang, Jesse Krantzler, and Brian P. Bailey. 2009. Getting Inspired! Understanding How and Why Examples are Used in Creative Design Practice. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 87–96.

34. Scarlett R. Herring, Brett R. Jones, and Brian P. Bailey. 2009. Idea Generation Techniques among Creative Professionals. In *Proceedings of the IEEE Hawaii International Conference on System Sciences*, 1–10.

35. Idris Hsi and Colin Potts. 2000. Studying the Evolution and Enhancement of Software Features. In *International Conference on Software Maintenance*, 143–151.

36. Hilary Hutchinson, Benjamin B. Bederson, Allison Druin, Catherine Plaisant, Wendy Mackay, Helen Evans, Heiko Hansen, Stephane Conversy, Michel Beaudouin-Lafon, Nicolas Roussel, Loïc Lacomme, Björn Eiderbäck, Sinna

Lindquist, Yngve Sundblad, and Bosse Westerlund. 2003. Technology Probes: Inspiring Design for and with Families. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 17–24.

37. David G. Jansson and Steven M. Smith. 1991. Design Fixation. *Design Studies* 12, 1: 3–11.

38. Leah Kaufman and Brad Weed. 1998. Too Much of a Good Thing? Identifying and Resolving Bloat in the User Interface. In *Conference Summary on Human Factors in Computing Systems.*, 207–208.

39. Caitlin Kelleher and Randy Pausch. 2005. Stencils-Based Tutorials: Design and Evaluation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 541–550.

40. Adnan Alam Khan, Volodymyr Dziubak, and Andrea Bunt. 2015. Exploring Personalized Command Recommendations based on Information Found in Web Documentation. In *Proceedings of the 20th International Conference on Intelligent User Interfaces*, 225–235.

41. Juho Kim, Phu Nguyen, Sarah Weir, Philip J. Guo, Robert C. Miller, and Krzysztof Z. Gajos. 2014. Crowdsourcing Step-by-Step Information Extraction to Enhance Existing How-to Videos. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, 4017–4026.

42. Nicholas Kong, Tovi Grossman, Björn Hartmann, George Fitzmaurice, and Maneesh Agrawala. 2012. Delta: A Tool For Representing and Comparing

Workflows. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1027–1036.

43. Chinmay Kulkarni, Steven P. Dow, and Scott R. Klemmer. 2012. Early and Repeated Exposure to Examples Improves Creative Work. In *Proceedings of the Annual Meeting of the Cognitive Science Society*.

44. Ben Lafreniere, Andrea Bunt, Matthew Lount, and Michael Terry. 2013. Understanding the Roles and Uses of Web Tutorials. In *Seventh International AAAI Conference on Weblogs and Social Media*, 303–310.

45. Amanda LeClerc. 2010. Seeking Serendipity: The Inspiration Hunt of a Creative Professional. *Faculty of Information Quarterly* 2, 3: 1–8.

46. Brian Lee, Savil Srivastava, Ranjitha Kumar, Ronen Brafman, and Scott R. Klemmer. 2010. Designing with Interactive Example Galleries. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2257–2266.

47. Wei Li, Tovi Grossman, and George Fitzmaurice. 2012. GamiCAD: A Gamified Tutorial System For First Time AutoCAD Users. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*, 103–112.

48. Wei Li, Tovi Grossman, and George Fitzmaurice. 2014. CADament: A Gamified Multiplayer Software Tutorial System. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, 3369–3378.

49. Wei Li, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2011. TwitApp: In-product Micro-Blogging for Design Sharing. In *Proceedings of the 24th Annual*

ACM Symposium on User Interface Software and Technology, 185–194.

50. J. S. Linsey, I. Tseng, K. Fu, J. Cagan, K. L. Wood, and C. Schunn. 2010. A Study of Design Fixation, Its Mitigation and Perception in Engineering Design Faculty. *Journal of Mechanical Design* 132, 4: 041003.

51. Frank Linton, Andy Charron, and Debbie Joy. 2000. OWL: a recommender system for organization-wide learning. *AAAI Educational Technology & Society* 3, 1: 62–76.

52. Matthew Lount and Andrea Bunt. 2014. Characterizing Web-Based Tutorials: Exploring Quality, Community, and Showcasing Strategies. In *Proceedings of the 32nd ACM International Conference on The Design of Communication CD-ROM*.

53. Andrés Lucero. 2012. Framing, Aligning, Paradoxing, Abstracting, and Directing: How Design Mood Boards Work. In *Proceedings of the Designing Interactive Systems Conference*, 438–447.

54. Sylvain Malacria, Joey Scarr, Andy Cockburn, Carl Gutwin, and Tovi Grossman. 2013. Skillometers. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*, 321–330.

55. Lena Mamykina, Bella Manoim, Manas Mittal, George Hripcsak, and Björn Hartmann. 2011. Design Lessons from the Fastest Q&A Site in the West. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2857–2886.

56. Gary Marchionini. 2006. Exploratory Search: From Finding to Understanding.

Communications of the ACM 49, 4: 41–46.

57. Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2011. Ambient help. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 2751–2760.

58. Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2011. IP-QAT: In-Product Questions, Answers & Tips. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, 175–184.

59. Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2013. Patina: Dynamic Heatmaps for Visualizing Application Usage. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 3227–3236.

60. Justin Matejka, Wei Li, Tovi Grossman, and George Fitzmaurice. 2009. CommunityCommands: Command Recommendations for Software Applications. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology*, 193–202.

61. Philip Mendels, Joep Frens, and Kees Overbeeke. 2011. Freed: A System for Creating Multiple Views of a Digital Collection during the Design Process. In *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems*, 1481–1490.

62. Scarlett R. Miller and Brian P. Bailey. 2014. Searching for inspiration: An in-depth look at designers example finding practices. In *ASME 2014 International Design Engineering Technical Conferences and Computers and Information in Engineering*

Conference, 1–10.

63. Dan Morris, Meredith Ringel Morris, and Gina Venolia. 2008. SearchBar: A Search-Centric Web History for Task Resumption and Information Re-finding. In *Proceeding of the Twenty-Sixth Annual CHI Conference on Human Factors in Computing Systems*, 1207–1216.

64. Raymond S. Nickerson. 1998. Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology* 2, 2: 175–220.

65. Bernard A. Nijstad and Wolfgang Stroebe. 2006. How the Group Affects the Mind: A Cognitive Model of Idea Generation in Groups. *Personality and Social Psychology Review* 10, 3: 186–213.

66. Bernard A. Nijstad, Wolfgang Stroebe, and Hein F.M. Lodewijkx. 2002. Cognitive Stimulation and Interference in Groups: Exposure Effects in an Idea Generation Task. *Journal of Experimental Social Psychology* 38, 6: 535–544.

67. Matti Perttula and Pekka Sipilä. 2007. The Idea Exposure Paradigm in Design Idea Generation. *Journal of Engineering Design* 18, 1: 93–102.

68. Suporn Pongnumkul, Mira Dontcheva, Wilmot Li, Jue Wang, Lubomir Bourdev, Shai Avidan, and Michael F. Cohen. 2011. Pause-and-Play: Automatically Linking Screencast Video Tutorials with Applications. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, 135–144.

69. A. T. Purcell and J. S. Gero. 1992. Effects of Examples on the Results of a Design Activity. *Knowledge-Based Systems* 5, 1: 82–91.

70. John Rieman. 1996. A Field Study of Exploratory Learning Strategies. *ACM Transactions on Computer-Human Interaction* 3, 3: 189–218.

71. Daniel Ritchie, Ankita Arvind Kejriwal, and Scott R Klemmer. 2011. d.tour: Style-Based Exploration of Design Example Galleries. In *Proceedings of the ACM Conference on User Interface Software and Technology*, 165–174.

72. Joey Scarr, Andy Cockburn, Carl Gutwin, and Philip Quinn. 2011. Dips and Ceilings: Understanding and Supporting Transitions to Expertise in User Interfaces. In *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems*, 2741–2750.

73. Donald A Schön. 1982. *The Reflective Practitioner*. Basic Books.

74. Moushumi Sharmin and Brian P Bailey. 2013. ReflectionSpace: An Interactive Visualization Tool for Supporting Reflection-on-action in Design. In *Proceedings of the 9th ACM Conference on Creativity & Cognition*, 83–92.

75. Moushumi Sharmin, Brian P. Bailey, Cole Coats, and Kevin Hamilton. 2009. Understanding Knowledge Management Practices for Early Design Activity and Its Implications for Reuse. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 2367–2376.

76. Ben Shneiderman. 2007. Creativity Support Tools: Accelerating Discovery and Innovation. *Communications of the ACM* 50, 12: 20–32.

77. Ben Shneiderman, Gerhard Fischer, Mary Czerwinski, Mitch Resnick, Brad Myers, Linda Candy, Edmonds Ernest, Mike Eisenberg, Elisa Giaccardi, Tom Hewett,

Pamela Jennings, Bill Kules, Nakakoji Kumiyo, Jay Nunamaker, Randy Pausch, Ted Selker, Elisabeth Sylvan, and Michael Terry. 2006. Creativity Support Tools: Report From a U.S. National Science Foundation Sponsored Workshop. *International Journal of Human-Computer Interaction* 20, 2: 61–77.

78. Pao Siangliulue, Joel Chan, Krzysztof Z. Gajos, and Steven P. Dow. 2015. Providing Timely Examples Improves the Quantity and Quality of Generated Ideas. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*, 83–92.

79. Anselm Strauss and Juliet Corbin. 1990. *Basics of Qualitative Research*. SAGE Publications.

80. Endel Tulving. 1972. Episodic and semantic memory. In *Organization of memory*. 381–403.

81. Luis A. Vasconcelos and Nathan Crilly. 2016. Inspiration and Fixation: Questions, Methods, Findings, and Challenges. *Design Studies* 42: 1–32.

82. Laton Vermette, Parmit Chilana, Michael Terry, Adam Fourney, Ben Lafreniere, and Travis Kerr. 2015. CheatSheet: A Contextual Interactive Memory Aid for Web Applications. In *Proceedings of the 41st Graphics Interface Conference*, 241–248.

83. Andy Warr and Eamonn O'Neill. 2005. Understanding Design as a Social Creative Process. In *Proceedings of the 5th Conference on Creativity & Cognition*, 118.

84. Ryen W. White and Steven M. Drucker. 2007. Investigating Behavioral Variability in Web Search. In *Proceedings of the 16th International Conference on World Wide*

Web, 21–30.

85. Ryen W. White, Gary Marchionini, and Gheorghe Muresan. 2008. Evaluating Exploratory Search Systems. Introduction to Special Topic Issue of Information Processing and Management. *Information Processing and Management* 44, 2: 433–436.

86. Ka-Ping Yee, Kirsten Swearingen, Kevin Li, and Marti Hearst. 2003. Faceted Metadata for Image Search and Browsing. In *Proceedings of the Conference on Human Factors in Computing Systems*, 401.

87. Tom Yeh, Tsung-Hsiang Chang, and Robert C Miller. 2009. Sikuli: Using GUI Screenshots for Search and Automation. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology*, 183.

88. Barry J. Zimmerman. 1989. Models of Self-Regulated Learning and Academic Achievement. In *Self-Regulated Learning and Academic Achievement* (1st ed.), Barry J. Zimmerman and Dale H. Schunk (eds.). Springer New York, New York, NY, 1–25.

Appendices

Appendix A

TCPS 2: Core Certificate

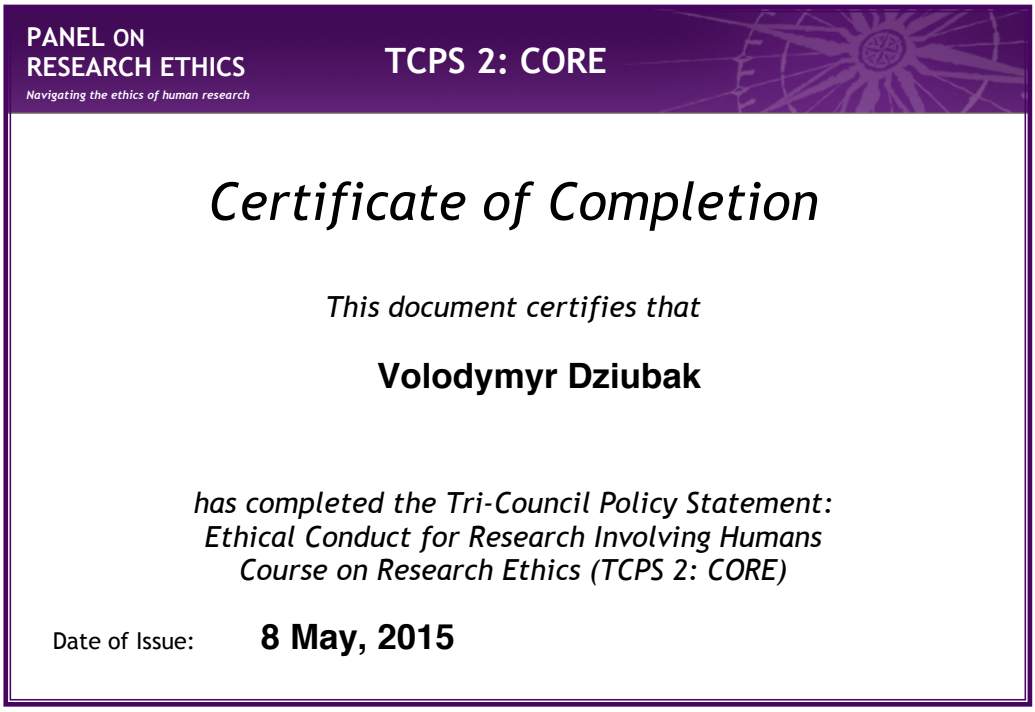

PANEL ON RESEARCH ETHICS
Navigating the ethics of human research

TCPS 2: CORE

Certificate of Completion

This document certifies that

Volodymyr Dziubak

has completed the Tri-Council Policy Statement:
Ethical Conduct for Research Involving Humans
Course on Research Ethics (TCPS 2: CORE)

Date of Issue: **8 May, 2015**

Appendix B

Ad Libitum Exploration Study

Ethics Approval

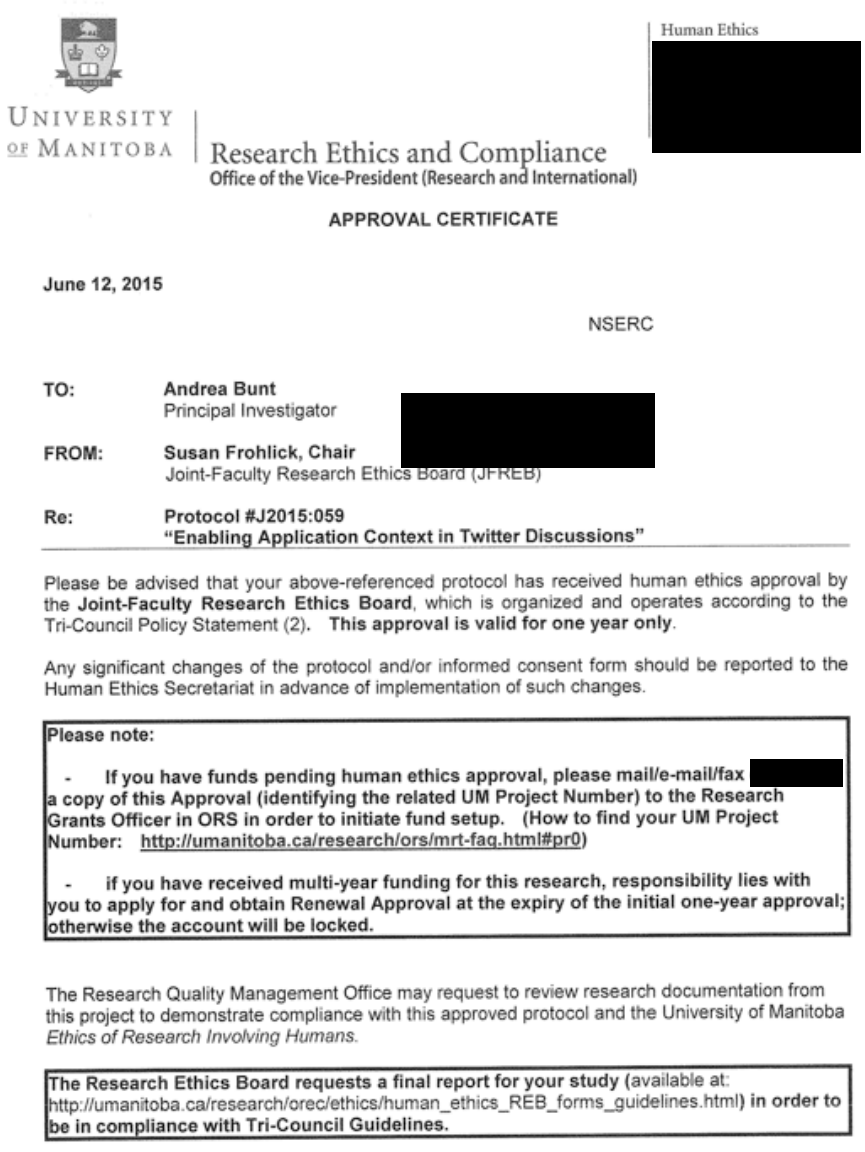

UNIVERSITY OF MANITOBA | Research Ethics and Compliance
Office of the Vice-President (Research and International)

Human Ethics [redacted]

APPROVAL CERTIFICATE

June 12, 2015

NSERC

TO: **Andrea Bunt**
Principal Investigator

FROM: **Susan Frohlick, Chair** [redacted]
Joint-Faculty Research Ethics Board (JFREB)

Re: **Protocol #J2015:059**
"Enabling Application Context in Twitter Discussions"

Please be advised that your above-referenced protocol has received human ethics approval by the **Joint-Faculty Research Ethics Board**, which is organized and operates according to the Tri-Council Policy Statement (2). **This approval is valid for one year only**.

Any significant changes of the protocol and/or informed consent form should be reported to the Human Ethics Secretariat in advance of implementation of such changes.

Please note:

- **If you have funds pending human ethics approval, please mail/e-mail/fax** [redacted] **a copy of this Approval (identifying the related UM Project Number) to the Research Grants Officer in ORS in order to initiate fund setup. (How to find your UM Project Number: http://umanitoba.ca/research/ors/mrt-faq.html#pr0)**

- **if you have received multi-year funding for this research, responsibility lies with you to apply for and obtain Renewal Approval at the expiry of the initial one-year approval; otherwise the account will be locked.**

The Research Quality Management Office may request to review research documentation from this project to demonstrate compliance with this approved protocol and the University of Manitoba *Ethics of Research Involving Humans*.

The Research Ethics Board requests a final report for your study (available at: http://umanitoba.ca/research/orec/ethics/human_ethics_REB_forms_guidelines.html) **in order to be in compliance with Tri-Council Guidelines.**

umanitoba.ca/research

UNIVERSITY OF MANITOBA | Research Ethics and Compliance
Office of the Vice-President (Research and International)

Human Ethics
[redacted]

AMENDMENT APPROVAL

June 23, 2015

TO: **Andrea Bunt**
Principal Investigator

FROM: **Susan Frohlick, Chair**
Joint-Faculty Research Ethics Board (JFREB)

Re: **Protocol #J2015:059**
"Enabling Application Context in Twitter Discussions"

This will acknowledge your Amendment Request dated June 18, 2015 requesting amendment to your above-noted protocol.

Approval is given for this amendment. Any further changes to the protocol must be reported to the Human Ethics Secretariat in advance of implementation.

Call for Participation

$15

Actively working with Photoshop or AutoCAD?

Help us with our study and get $15 Starbucks gift card!

[redacted]

Research approved by University of Manitoba Faculty Research Ethics Board

Photoshop/AutoCAD study [redacted]

Photoshop/AutoCAD study [redacted]

Photoshop/AutoCAD study [redacted]

Photoshop/AutoCAD study [redacted]

Photoshop/AutoCAD study [redacted]

Photoshop/AutoCAD study [redacted]

Photoshop/AutoCAD study [redacted]

Photoshop/AutoCAD study [redacted]

Photoshop/AutoCAD study [redacted]

Recruitment Script

We are currently conducting a user study involving users of Adobe Photoshop and Autodesk AutoCAD. The purpose of this study is to explore the option of discussing specific commands of the software using Twitter.

Participation in the study will involve interacting with web-based prototype that imitates original application interface (Photoshop or AutoCAD) followed by a semi-structured interview. The entire session will take from 30 to 60 minutes of your time. Participating in the study will be rewarded with a $15 gift card.

If you are interested or need more information, please contact Volodymyr Dziubak

[redacted]

For further inquiries about the study, please contact Dr. Andrea Bunt

[redacted]

Consent Form

UNIVERSITY OF MANITOBA

DEPARTMENT OF COMPUTER SCIENCE

Winnipeg, Manitoba [redacted]

Research Project Title: Enabling Application Context in Twitter Discussions

Researchers: Dr. Andrea Bunt, [redacted], Volodymyr Dziubak [redacted]

Please take the time to read this carefully and to ensure you understand all the information.

This consent form, a copy of which will be left with you for your records and reference, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information.

You are invited to participate in a research interview about using Twitter to learn about specific features of complex software applications. The goal of the study is to 1) gather feedback from users of feature-rich software applications on how they learn about new commands in the software and 2) build understanding of how Twitter features (e.g. networking and sharing via retweets) could improve current learning practices. If you have any questions or concerns at this time or any time during the interview, please feel free to ask the on-site researcher for clarification.

Participation in this study is voluntary, and will take approximately 30 minutes to one hour of your time.

You will receive a $15 gift card for your participation. We wish to audiotape the session. The audio will serve as a reference point in our data analysis, allowing us to review the study session in detail. Any information you choose to disclose is completely confidential and will be used for anonymized research analysis. We may use anonymized quotes for purposes of dissemination; your name will not be included or in any other way associated with the data presented in the results of this study.

Data collected during this interview will be retained for a period of maximum seven years in a locked cabinet or password-protected computer in a locked office in the EITC building, University of Manitoba, to which only researchers associated with this interview (Dr. Andrea Bunt, Volodymyr Dziubak) have access. In addition, the University of Manitoba may look at research records to see that the research is being done in a safe and proper way. Once published, results of the interview will be made available to the public for free at http://home.cs.umanitoba.ca/~bunt/. Again, no personal information about your involvement will be included.

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a subject. By doing this you also confirm that you are of the age of majority in Canada (18 years or more). In no way does this waive your legal rights nor release the researchers, sponsors, or involved institutions from their legal and professional responsibilities. You are free to withdraw from the study at any time, and /or refrain from answering any questions you prefer to omit, without prejudice or consequence. Your continued

participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation. You may withdraw at any time throughout the interview and you will still receive your full compensation of $15 gift card unconditionally. You can also ask to withdraw your data any time between now and two weeks from today's date. Your data will be destroyed if you choose to withdraw from the study.

This research has been approved by the University of Manitoba Joint Faculty Research Ethics Board. If you have any concerns or complaints about this project you may contact Dr. Andrea Bunt at [redacted] or the Human Ethics Coordinator (HEC) at [redacted] A copy of this consent form has been given to you to keep for your records and reference.

☐ **For purposes of research analysis, we wish to record the interview.** By checking this box, you agree that you understand this and that we may use the audio record for data analysis purposes.

Participant's Signature _____ Date _____

Researcher's Signature _____ Date _____

Demographics Questionnaire

Demographics questionnaire

Participant Number: _____

Gender:

Male / Female

Age:

18-25 26-35 36-45 46 and over

What is your experience with AutoCAD (Photoshop)?

0 Years 1-3 Years 3-6 Years 6+ Years

How often do you use Twitter?

I don't use Twitter

Once or twice per month

Couple of times per week

Couple of times per day

Sample Interview Questions

1. Have you ever explored the menu of software to see what kinds of commands are there? Imagine that during the exploration you found a command that does something cool, but you do not know when and how it could be applied in an actual task. How would you imagine looking for such information?
2. Imagine you've just discovered a command that significantly simplified your workflow (e.g. snapping to background grid). How do you imagine sharing your finding with the others?
3. Which functionality you think would be useful apart from tweeting about specific commands?
4. Imagine you found a tweet with useful information about how to use the command you know well in a better way. How do you imagine using Twitter features (e.g. retweet or reply to the tweet) to share this information with others?
5. Do you have other designers/architects in your Twitter network? If your connections were also using this system, would you prefer to have tweets from your connections to be promoted in some way?

Appendix C

Switter Evaluation

Ethics Approval

UNIVERSITY OF MANITOBA | Research Ethics and Compliance
Office of the Vice-President (Research and International)

Human Ethics [redacted]

AMENDMENT APPROVAL

August 19, 2015

TO: **Andrea Bunt**
Principal Investigator

FROM: **Susan Frohlick, Acting Chair** [redacted]
Joint-Faculty Research Ethics Board (JFREB)

Re: **Protocol #J2015:059**
"Enabling Application Context in Twitter Discussions"

This will acknowledge your Amendment Request dated August 14, 2015 requesting amendment to your above-noted protocol.

Approval is given for this amendment. Any further changes to the protocol must be reported to the Human Ethics Secretariat in advance of implementation.

umanitoba.ca/research

Recruitment Script

We are currently conducting a user study involving users of Adobe Photoshop. The purpose of this study is to evaluate usability of an experimental prototype that allows people to discuss specific commands of the software using Twitter.

Participation in the study will involve interacting with a web-based prototype that imitates the original application interface (Photoshop). Participants will be asked to interact with the system for at least 30 minutes each day for the period of 1 week. At the end of each day we will ask participants to describe the experience of using the system in a quick online questionnaire. We will also conduct two interviews: one in the middle of the study and one at the end. Each interview will take from 30 to 60 minutes long and will be held either in person or online via Skype.

In appreciation of participation, we will offer give a $75 CAD gift card.

If you are interested or need more information, please contact Volodymyr Dziubak [redacted]

For further inquiries about the study, please contact Dr. Andrea Bunt [redacted]

Consent Form

Research Project Title: Enabling Application Context in Twitter Discussions

Researchers: Dr. Andrea Bunt, [redacted], Volodymyr Dziubak [redacted]

Please take the time to read this carefully and to ensure you understand all the information.

This consent form, a copy of which will be left with you for your records and reference, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information.

You are invited to participate in a research study about using Twitter to learn about specific features of complex software applications. The goal of the study is to 1) gather feedback from users of feature-rich software applications on how they learn about new commands in the software and 2) build understanding of how Twitter features (e.g. networking and sharing via retweets) could improve current learning practices. If you have any questions or concerns at this time or any time during the interview, please feel free to ask the on-site researcher for clarification.

Participation in this study is voluntary. We will ask you to use an experimental web-browser based system for at least 30 minutes each day for the total duration of one week. At the end of each day we will ask you to describe your experience with the system in a short online questionnaire. We will also conduct 2 interview sessions: one in the middle of the study (end of day 3) and one at the end of the study. Each interview will take approximately 30 to 60 minutes of your time.

As gratitude for participation, you will receive a $25 CAD gift card at the beginning of the study. Upon the completion of the study you will additionally receive another gift card worth $50 CAD.

We wish to audiotape the session. The audio will serve as a reference point in our data analysis, allowing us to review the study session in detail. Any information you choose to disclose is completely confidential and will be used for anonymized research analysis. We may use anonymized quotes for purposes of dissemination; your name will not be included or in any other way associated with the data presented in the results of this study.

Data collected during this interview will be retained for a period of maximum seven years in a locked cabinet or password-protected computer in a locked office in the EITC building, University of Manitoba, to which only researchers associated with this interview (Dr. Andrea Bunt, Volodymyr Dziubak) have access. In addition, the University of Manitoba may look at research records to see that the research is being done in a safe and proper way. Once published, results of the interview will be made available to the public for free at http://home.cs.umanitoba.ca/~bunt/. Again, no personal information about your involvement will be included.

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a subject. By doing this you also confirm that you are of the age of majority in Canada (18 years or more). In no way does this waive your legal rights nor release the researchers, sponsors, or involved institutions from their legal and professional responsibilities. You are free to withdraw from the study at any time, and /or refrain from

answering any questions you prefer to omit, without prejudice or consequence. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

You may withdraw at any time throughout the study. In this case you will not receive the final $50 CAD gift card. However, you will keep the initial $25 CAD gift card unconditionally. You can also ask to withdraw your data any time between now and two weeks from today's date. Your data will be destroyed if you choose to withdraw from the study.

This research has been approved by the University of Manitoba Joint Faculty Research Ethics Board. If you have any concerns or complaints about this project you may contact Dr. Andrea Bunt at [redacted] or the Human Ethics Coordinator (HEC) at [redacted]. A copy of this consent form has been given to you to keep for your records and reference.

☐ **For purposes of research analysis, we wish to record the interview.** By checking this box, you agree that you understand this and that we may use the audio record for data analysis purposes.

Participant's Signature _____ Date _____

Researcher's Signature _____ Date _____

Demographics Questionnaire

Demographics questionnaire

Participant Number:

Age:

18-25 26-35 36-45 46 and over

What is your experience with Photoshop?

1-3 Years 4-6 Years 7-10 Years 10+ Years

How often do you check Twitter?

I don't use Twitter

Once or twice per month

Couple of times per week

Couple of times per day

Sample Questions for the Daily Journal

1. Which command(s) do you remember reading about today? Why did these commands catch your attention? Did you know how to use or have you used these commands before?
2. Describe the tweet(s) that you found useful. What was/were the tweet about? What made this information interesting to you?
3. Which functionality do you think should be added to the system? Please, describe what you were trying to achieve when you realized that the described functionality was missing.
4. Did you discover any bugs when working with the system? If so, please describe what went wrong and what should be done to reproduce the bug.
5. (free-form feedback) Please, provide any other feedback about the system that does not fit into the categories above. For example, what you think might be implemented in a better way, or any suggestions regarding design, positioning of user interface elements, etc.

Appendix D

Instory Evaluation

Ethics Approval

APPROVAL CERTIFICATE

June 9, 2016

TO: **Andrea Bunt**
Principal Investigator

FROM: **Lorna Guse, Chair**
Joint-Faculty Research Ethics Board (JFREB)

Re: **Protocol #J2016:048 (HS19778)**
"Supporting Inspiration Seeking Among Professional Digital Designers"

Please be advised that your above-referenced protocol has received human ethics approval by the **Joint-Faculty Research Ethics Board**, which is organized and operates according to the Tri-Council Policy Statement (2). **This approval is valid for one year only and will expire on June 9, 2017.**

Any significant changes of the protocol and/or informed consent form should be reported to the Human Ethics Secretariat in advance of implementation of such changes.

Please note:

- **If you have funds pending human ethics approval, please mail/e-mail/fax [redacted] a copy of this Approval (identifying the related UM Project Number) to the Research Grants Officer in ORS in order to initiate fund setup. (How to find your UM Project Number: http://umanitoba.ca/research/ors/mrt-faq.html#pr0)**

- **if you have received multi-year funding for this research, responsibility lies with you to apply for and obtain Renewal Approval at the expiry of the initial one-year approval; otherwise the account will be locked.**

The Research Quality Management Office may request to review research documentation from this project to demonstrate compliance with this approved protocol and the University of Manitoba *Ethics of Research Involving Humans.*

The Research Ethics Board requests a final report for your study (available at: http://umanitoba.ca/research/orec/ethics/human_ethics_REB_forms_guidelines.html) **in order to be in compliance with Tri-Council Guidelines.**

umanitoba.ca/research

UNIVERSITY OF MANITOBA | Research Ethics and Compliance
Office of the Vice-President (Research and International)

Human Ethics
[redacted]

AMENDMENT APPROVAL

August 5, 2016

TO: **Andrea Bunt**
Principal Investigator
[redacted]

FROM: **Lorna Guse, Chair**
Joint-Faculty Research Ethics Board (JFREB)

Re: **Protocol #J2016:048 (HS19778)**
"Supporting Inspiration Seeking Among Professional Digital Designers"

This will acknowledge your Amendment Request dated July 25, 2016 requesting amendment to your above-noted protocol.

Approval is given for this amendment. Any further changes to the protocol must be reported to the Human Ethics Coordinator in advance of implementation.

UNIVERSITY OF MANITOBA
Research Ethics and Compliance
Office of the Vice-President (Research and International)

AMENDMENT APPROVAL

August 23, 2016

TO: **Andrea Bunt**
Principal Investigator [redacted]

FROM: **Lorna Guse, Chair**
Joint-Faculty Research Ethics Board (JFREB)

Re: **Protocol #J2016:048 (HS19778)**
"Supporting Inspiration Seeking Among Professional Digital Designers"

This will acknowledge your Amendment Request dated August 22, 2016 requesting amendment to your above-noted protocol.

Approval is given for this amendment. Any further changes to the protocol must be reported to the Human Ethics Coordinator in advance of implementation.

Recruitment Script

We are currently conducting a user study to understand designers' experience when collecting inspirational design examples. We are looking for professional designers to provide feedback on a prototype system designed to improve designer's experience of creating a mood board – a collection of images that designers collect and reference for inspiration. The prototype system displays collected images accompanied by widgetsthat reveal how and why specific images were collected. Our user study investigates how showing this additional information about a mood board impacts a designer's awareness of the number and diversity of explored ideas.

Participation in the study will take about one to two hours of your time and will consist of 1) completing a sample mood board collecting task using the prototype and 2) a semi-structured interview, where we will discuss your experience with the system and collect your feedback on its features and functionality. Participants will be rewarded with a $40 gift card.

If you are interested or need more information, please contact Volodymyr Dziubak [redacted]

For further inquiries about the study, please contact Dr. Andrea Bunt [redacted]

Consent Form

UNIVERSITY OF MANITOBA

DEPARTMENT OF COMPUTER SCIENCE

Winnipeg, Manitoba [redacted]

Research Project Title: Supporting inspiration seeking among professional digital designers

Researchers:

Dr. Andrea Bunt,
Associate Professor, Department of Computer Science, University of Manitoba, [redacted]

Volodymyr Dziubak,
Ph.D. student, Department of Computer Science, University of Manitoba, [redacted]

Research Sponsored by Natural Sciences and Engineering Research Council of Canada

Please take the time to read this carefully and to ensure you understand all the information.

This consent form, a copy of which will be left with you for your records and reference, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information.

You are invited to participate in a research study on supporting inspiration seeking among professional digital designers. We will ask you to interact with a prototype system designed to support creating a mood board – a collection of images that designers collect and reference for inspiration. The prototype system displays collected images accompanied by widgets that reveal how and why specific images were collected. Our user study investigates how this additional information about a mood board impacts awareness of the number and diversity of explored ideas. If you have any questions or concerns at this time or any time during the study, please feel free to ask the on-site researcher for clarification.

Participation in this study is voluntary, and will take approximately one to two hours of your time. During this time, we will give you a description of a sample design project and ask you to create a mood board of images you would reference for inspiration during the design process. After completing the task, we would like to discuss your experience with the system in a form of an interview. During the interview, we would collect your feedback on the system's features and functionality. Afterwards there will be a debriefing session, where we will provide further detail on the purpose of the study, and give you an opportunity to ask any further questions.

You will receive a $40 gift card for your participation. We wish to audio-record and video-record/ screen-capture the session. With your permission, we may ask to take pictures or screenshots of designs and inspirational materials we discuss in the study. None of these materials will be used for any purpose other than non-commercial research. Recordings will serve as a reference point in our data analysis, allowing us to review the study session in detail. Any information you choose to disclose is completely confidential and will be used for anonymized research analysis. We may use anonymized quotes and pictures (or screenshots) or inspirational materials and designs for purposes of dissemination; your name will not be included or in any other way associated with the data presented in the results of this study.

There are no known risks associated with participation in the study. As a personal benefit, you might learn about current practices of inspiration collection among professional designers.

Data collected during this interview will be retained for a period of maximum seven years in a locked cabinet or password-protected computer in a locked office in the EITC building, University of Manitoba, to which only researchers associated with this interview (Dr. Andrea Bunt, Volodymyr Dziubak) have access. In addition, the University of Manitoba may look at research records to see that the research is being done in a safe and proper way. Once published, results of the interview will be made available to the public for free at http://hci.cs.umanitoba.ca/. Again, no personal information about your involvement will be included.

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a subject. By doing this you also confirm that you are of the age of majority in Canada (18 years or more). In no way does this waive your legal rights nor release the researchers, sponsors, or involved institutions from their legal and professional responsibilities. You are free to withdraw from the study at any time, and /or refrain from answering any questions you prefer to omit, without prejudice or consequence. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation. You may withdraw at any time throughout the interview and you will still receive your full compensation of $40 gift card unconditionally. You can also ask to withdraw your data any time between now and one month from today's date. Your data will be destroyed if you choose to withdraw from the study.

This research has been approved by the University of Manitoba Joint Faculty Research Ethics Board. If you have any concerns or complaints about this project, you may contact Dr. Andrea Bunt at [redacted] or the Human Ethics Coordinator (HEC) at [redacted] A copy of this consent form has been given to you to keep for your records and reference.

☐ **I allow researchers to record audio and video (screen-capture) of the study session.** I give my permission to use anonymized recorded data for analysis and to include anonymized transcribed quotes in publications and presentations.

☐ **I allow researchers take pictures or screenshots of designs and inspirational materials discussed in the study.** I give my permission to use anonymized collected pictures (screenshots) for data analysis and to include anonymized pictures (screenshots) in publications and presentations.

Participant's Signature _____ Date _____

Researcher's Signature _____ Date _____

Demographics Questionnaire

Participant ID: _____

Age:

18-24 25-31 32-38 39-45 46+

How many years of professional design experience you have?

1-2 Years 3-4 Years 5-6 Years 7+ Years

What are your specialties within the design area? (e.g. web design, logo design, etc...)

List (up to three) design software programs you use most often in your work

---

Sample Interview Questions

1. How did you approach the task? What were your first design ideas?
2. How many design ideas do you think you explored while collecting mood board for the sample project?
3. Did you try to explore several completely different options or did you focus on a specific idea and iterated over it?
4. Which ideas out of the ones you explored do you think were most interesting?
5. When adding an image to your collection, the interface prompted you to specify a reason. Would you be willing to indicate the reasoning behind adding each image or not? Why?
6. While searching for images, did you ever get distracted by non-related images? If so, did you feel you had difficulty getting back to your initial search?
7. Did you have trouble remembering what you have already searched for?
8. How confident do you think you are in terms of exploring a wide range of ideas? Do you feel there might be ideas you have missed?
9. How would you rank the additional information provided by the prototype by it's usefulness? (from the most useful information to the least useful information)
10. How much did additional information distract you from collecting a mood board? Which

Appendix E

Prism Evaluation

Ethics Approval

UNIVERSITY OF MANITOBA | Research Ethics and Compliance

Human Ethics [redacted]

AMENDMENT APPROVAL

March 9, 2017

TO: **Andrea Bunt**
Principal Investigator

FROM: **Kevin Russell, Chair**
Joint-Faculty Research Ethics Board (JFREB) [redacted]

Re: **Protocol #J2016:048 (HS19778)**
"Supporting Inspiration Seeking Among Professional Digital Designers"

Joint-Faculty Research Ethics Board (JFREB) has reviewed and approved your Amendment Request received on **March 8, 2017** to the above-noted protocol. JFREB is constituted and operates in accordance with the current *Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans.*

This approval is subject to the following conditions:

1. Approval is given for this amendment only. Any further changes to the protocol must be reported to the Human Ethics Coordinator in advance of implementation.
2. Any deviations to the research or adverse events must be submitted to JFREB as soon as possible.
3. Amendment Approvals do not change the protocol expiry date. Please refer to the original Protocol Approval or subsequent Renewal Approvals for the protocol expiry date.

Recruitment Script

We are currently conducting a user study to understand designers' experience when collecting inspirational design examples. We are looking for professional designers to test our prototype system for managing and organizing inspirational and reference images found on the Internet. We will ask you to use our system in your normal design workflow for a duration of 2 weeks. We will have two interview sessions (in the middle of the study and at the end of the study), during which we will discuss your experience with the system and collect your feedback on its features and functionality. Participants will enter a raffle for winning one of two $100 gift cards.

If you are interested or need more information, please contact Volodymyr Dziubak [redacted]

For further inquiries about the study, please contact Dr. Andrea Bunt [redacted]

Consent Form

UNIVERSITY OF MANITOBA

DEPARTMENT OF COMPUTER SCIENCE

Winnipeg, Manitoba [redacted]

Research Project Title: Supporting inspiration seeking among professional digital designers

Researchers:

Dr. Andrea Bunt, Associate Professor, Department of Computer Science, University of Manitoba [redacted]

Volodymyr Dziubak, Ph.D. student, Department of Computer Science, University of Manitoba [redacted]

Research Sponsored by Natural Sciences and Engineering Research Council of Canada

Please take the time to read this carefully and to ensure you understand all the information.

This consent form, a copy of which will be left with you for your records and reference, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information.

You are invited to participate in a research study on supporting inspiration seeking among professional digital designers. For the duration of two weeks, we will ask you to use a prototype system that allows you to capture and record your image-seeking behavior on Google Images. Specifically, the prototype system allows you to record search queries that you type and images that you inspect when exploring images. The system then visualizes the recorded process: it shows investigated images together with additional cues about how and when you saw these images. The tool also allows you to arrange investigated images into collections for further reference. Our user study investigates whether such a tool could promote reflection on the process of collecting inspirational images and how the information provided by the tool affects designers' inspiration seeking practices. If you have any questions or concerns at this time or any time during the study, please feel free to contact the researcher for clarification.

The study involves two interview sessions (one after the first week and one after the second week), during which we will contact you and ask you about your experience using the prototype system and about whether it affected your workflow.

We wish to audio-record and video-record/ screen-capture each interview session for data analysis. With your permission, we may ask to take pictures or screenshots of designs and inspirational materials we discuss in the study. None of these materials will be used for any purpose other than non-commercial research. Recordings will serve as a reference point in our data analysis, allowing us to review the study session in detail. Any information you choose to disclose is completely confidential and will be used for anonymized research analysis. We may use anonymized quotes and pictures (or screenshots) or inspirational materials and designs for

purposes of dissemination; your name will not be included or in any other way associated with the data presented in the results of this study.

As gratitude for participation, you will receive a $25 CAD gift card at the beginning of the study. Upon the completion of the study you will additionally receive another gift card worth $50 CAD.

There are no potential risks associated with participation in the study. As a personal benefit, you might learn about current practices of inspiration collection among professional designers.

Data collected during this interview will be retained for a period of maximum seven years in a locked cabinet or password-protected computer in a locked office in the EITC building, University of Manitoba, to which only researchers associated with this interview (Dr. Andrea Bunt, Volodymyr Dziubak) have access. In addition, the University of Manitoba may look at research records to see that the research is being done in a safe and proper way. Once published, results of the interview will be made available to the public for free at http://hci.cs.umanitoba.ca/. Again, no personal information about your involvement will be included.

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a subject. By doing this you also confirm that you are of the age of majority in Canada (18 years or more). In no way does this waive your legal rights nor release the researchers, sponsors, or involved institutions from their legal and professional responsibilities. You are free to withdraw from the study at any time, and /or refrain from answering any questions you prefer to omit, without prejudice or consequence. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation. You may withdraw at any time throughout study. In this case, you will keep the initial $25 CAD gift card unconditionally. You can also ask to withdraw your data any time between now and one month from today's date. Your data will be destroyed if you choose to withdraw from the study.

This research has been approved by the University of Manitoba Joint Faculty Research Ethics Board. If you have any concerns or complaints about this project, you may contact Dr. Andrea Bunt at [redacted] or the Human Ethics Coordinator (HEC) at [redacted] A copy of this consent form has been given to you to keep for your records and reference.

☐ **I allow researchers to record audio and video (screen-capture) of the interview.** I give my permission to use anonymized recorded data for analysis and to include anonymized transcribed quotes in publications and presentations.

☐ **I allow researchers take pictures or screenshots of designs and inspirational materials discussed in the study.** I give my permission to use anonymized collected pictures (screenshots) for data analysis and to include anonymized pictures (screenshots) in publications and presentations.

Participant's Signature _____ Date _____

Researcher's Signature _____ Date _____

Demographics Questionnaire

Participant ID: _____

Age:

18-24 25-31 32-38 39-45 46+

How many years of professional design experience you have?

1-2 Years 3-4 Years 5-6 Years 7+ Years

What are your specialties within the design area? (e.g. web design, logo design, etc...)

List (up to three) design software programs you use most often in your work

---

Sample Interview Questions

1. How many times did you use the prototype tool?
2. Could you, please, guide me through one of your image-seeking sessions? What did you search for? When did you stop searching? How did you arrive to the set of images that you ended up using (as references or for inspiration)?
3. Describe how, if at all, you used the tool in one of your image-seeking session. What did you use it for? How did you use it?
4. Did you discover something about your process/search habits that you did not know before? If so, what was it?
5. How would you approach the same task without the tool? Which parts of your workflow would be different?
6. Do you see yourself using this kind of a tool in your regular workflow? What would you use it for?
7. Are there any other features that you feel the tool is missing? What do you think such a feature would help you with?