Prism: Enhancing Graphic Designers' Visual Research with Interactive Search Trails

Volodymyr Dziubak
University of Manitoba,
Winnipeg MB, Canada
vdziubak@cs.umanitoba.ca

Andrea Bunt
University of Manitoba,
Winnipeg MB, Canada
bunt@cs.umanitoba.ca

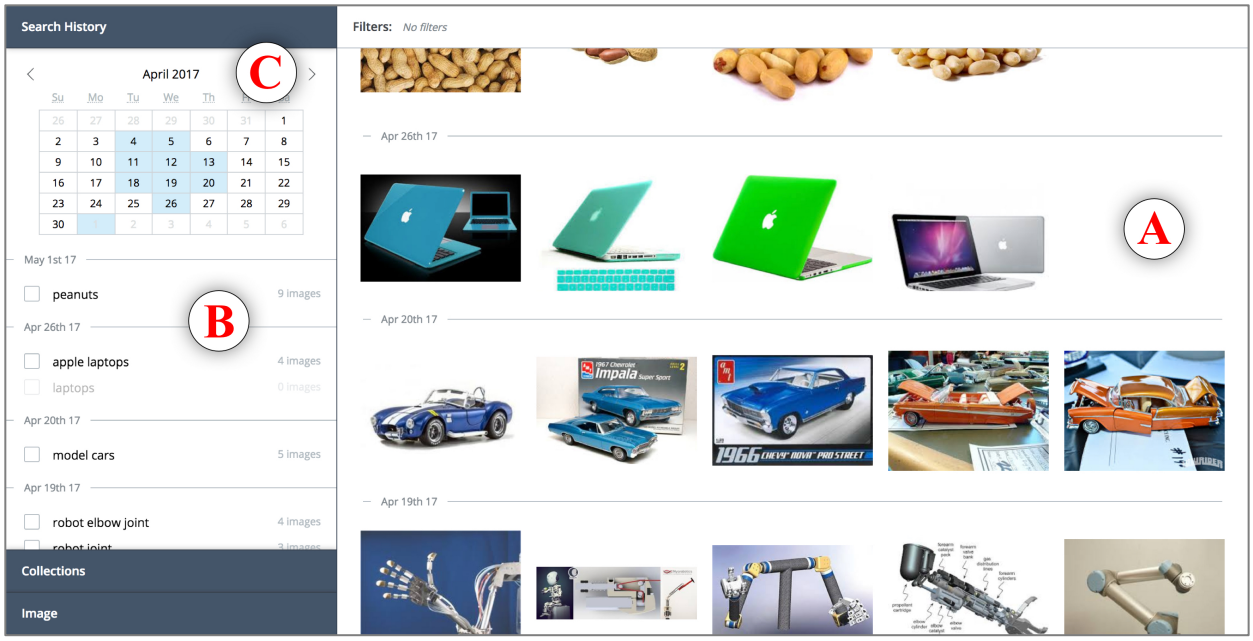

Figure 1. Screenshot of Prism. (A) Extended image library; (B) Index into the library via search queries; (C) Calendar highlighting days when Prism captured search trails.

ABSTRACT

Graphic designers often use the Web to collect images to use as inspiration and references for their work. Their resulting collections of images, however, typically do not retain important aspects of their visual research, such as their thought process when searching and all explored design avenues. Guided by an exploratory study with 14 expert graphic designers, we developed Prism – a system that supports a graphic designer's visual research on the Internet by automatically capturing all inspected images and annotating them with the designer's search trails. We evaluated Prism through a two-week field study with 11 expert designers. Our findings suggest that Prism's capture and display capabilities helped the designers to reify their design thinking, to better reflect on and compare alternative design ideas, and to collaborate with their colleagues and clients.

DIS '18, June 9–13, 2018, , Hong Kong

ACM ISBN 978-1-4503-5198-0/18/06□$15.00
https://doi.org/10.1145/3196709.3196756

AUTHOR KEYWORDS

Creativity support tools; Visual research

ACM CLASSIFICATION KEYWORDS

H.5.m. Information interfaces and presentation (e.g., HCI)

INTRODUCTION

Online visual research is an important part of the graphic design process. In various stages of their projects, graphic designers and artists forage the Web for inspiration and reference images (e.g., [3,9,16]). This practice helps graphic designers to explore the design space and to get a better understanding of the problem.

Traditionally, graphic designers keep track of their visual research by collecting images that inspire them, for example, by saving them to their computer. These image collections

provide the designers with a visual framework for comparing and evaluating their design ideas, as well reference points for later stages of the design process [9]. However, there are numerous other aspects of this visual research that are as important to designers as the images themselves, but that are not reflected in these image collections [9,16,24]. One example of such additional information are the "stories" behind the saved images [23,24] – a combination of information on how the designer found each image and why they decided to save it. Another example includes the design directions that the designer intended to explore, but did not explore because they ended up following an alternative path [9].

In this work, we investigate the idea of enhancing graphic designers' visual research workflows via automatically captured search trails. The motivation behind using search trails is to help capture the evolution of a graphic designer's 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.

After exploring this concept by having 14 expert graphic designers interact with a prototype technology probe, we designed and developed Prism (Figure 1) – a system that automatically collects all images that a 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. After using Prism for two weeks in their everyday work, participants were excited about how the system supported their visual research and spoke to a number of important use cases. They found that the system allowed them to reify their design thinking, to better reflect on alternative design ideas, and to establish better communication and collaboration with their colleagues and clients.

Overall, this paper makes the following contributions:

- Introducing the idea of automatically captured search trails within the context of visual research for graphic design.
- Prism, a prototype system that embodies this approach.
- Qualitative findings from a two-week field study with 11 expert graphic designers, demonstrating support for the general approach and suggesting promising avenues for further research.

RELATED WORK

Our coverage of related work focuses on prior work characterizing and supporting visual research. We also briefly overview other work on leveraging search query information for improved user experience.

Visual research

From the perspective of ideation theory, visual research helps designers build analogies, which is one of the driving forces behind a creative process [3]. Active and passive visual research are among the primary techniques for idea generation [10]. Collecting design examples gives designers a visual framework for evaluating and communicating their ideas to other stakeholders [9]. Designers also often collect and store inspirational design examples to reflect on their process in the future. Such reflection allows them to analyze the flow of their design process and use it as a template for future projects [24]. In what follows, we discuss some of the challenges that designers encounter when conducting visual research and prior work on supporting designers with this task.

Supporting exploration and retrieval of design examples

Existing research on visual research in design indicates that designers often struggle to articulate their abstract design ideas into keywords to use in a search engine [9,16,24]. Prior work in Human-Computer Interaction (HCI) has investigated a variety of alternative search engines to support visual research. For example, Hashimoto and Igarashi developed a method of retrieving design examples using user-generated sketches as queries [8]. Yee *et al.* [27] allowed designers to browse images and design examples by the image metadata (e.g., media, themes, location, date, shapes). Lee *et al.* [12] 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.* [21] moved this approach a step further by introducing techniques for style-based searching, recommendation, and filtering. In contrast, our approach focuses on enhancing the information surrounding images once designers retrieve them via a search engine. We implement our approach for a traditional keyword search, but it could also be extended to these types of advanced image-retrieval methods.

Information loss in visual research

Designers often lose information when conducting visual research [9,24]. For example, they forget to save or bookmark their search and have to re-do the search again later [9]. They can lose their train of thought as they explore alternative ideas sparked by the images that they see [9]. Sharmin *et al.* also found that storing images does not capture "stories" behind them (e.g., the rationale, the thought process, and the reasons why they saved them) – information that is as important to the designers as the images themselves [24]. Our work builds on these findings, aiming to help designers retain some of the information lost during visual research.

Supporting reflection on visual research

It is common for designers to reflect on their past design processes (reflection-on-action, as defined by Schön [22]). Such reflection helps them to communicate their design process to their clients and to gauge the progression of their process and style [24].

Various knowledge management systems have been designed to help designers reflect on their practice by capturing aspects of their process. For example, the Freed system [15] allows designers to manually organize their image collections into spatial views and to indicate relations between them. Rich bookmarks provide a designer with the means for reflection through manual curation, such as spatial positioning, color, and translucence [25]. These systems aim to provide an interactive environment for curation, whereas Prism aims to enhance the visual research experience by automatically capturing the designer's search trails.

ReflectionSpace [23] automatically organizes images found on the designer's computer using metadata from the images, such as creation date and file name. Our Prism system also automatically organizes the designer's inspiration and reference images. However, we focus on capturing aspects of the search activity itself, rather than leveraging properties retained by the stored images.

Design fixation

Design fixation is another well-researched challenge that emerges during visual research in design practice [11,18–20]. Design fixation happens when a designer becomes influenced by existing design solutions. This form of unconscious bias towards existing design solutions has been shown to hinder creativity [19]. Prior work has found that harmful effects of exposure to external inspiration sources might be reduced by using examples that are neither too far, nor too close to the target problem [13], by increasing the diversity of presented examples [18], and by using examples that illustrate ideas novel to the designer [19,20]. Our work does not specifically aim to solve the challenges caused by design fixation, however, Prism could potentially help visual designers to reflect on the diversity of explored venues.

Using browsing history to enhance search tasks

Prior research has used browsing history to investigate how people use web search to satisfy their information needs. For example, similar to Prism, White and Drucker also used the concept of a search trail, which in their case was defined a segment of browsing history that started with a search query and ended when the system assumed that the user completed their information-seeking activity [26]. Our work extends the idea of search trails to the context of visual research.

Prism is also motivated by prior work on using search trails to improve user's general searching capabilities or their searching efficiency. For example, the Search Dashboard system [2] 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 [17] 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. Our work brings the idea of capturing search information to the domain of creativity and visual research, which, to our knowledge, has not yet been investigated.

ENHANCING VISUAL RESEARCH WITH INTERACTIVE SEARCH TRAILS

Existing research outlines a number of challenges that designers face during visual research. The highly subjective nature of the domain, however, makes it difficult to find a computational approach that would support designers' visual research without interfering with their design goals and stylistic preferences. For example, the literature suggests that drawing inspiration from the design examples that are diverse and new to the designer could lead to more innovative and creative designs [18, 19, 20]. However, our informal conversations with experienced graphic designers suggested that a system that tries to push designers towards exploring diverse and novel ideas might be perceived as questioning their expertise. To avoid the danger of a system that assumes expertise beyond that of the designer, we decided to build our approach around the idea of providing the designer with objective information about their process, which would enable them to assess their process and react if they consider necessary.

Motivated by prior research on using search trails, 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.

TECHNOLOGY PROBE

To investigate how to capture and represent search trails to support graphical designers' visual research, we developed an initial prototype system to serve as a technology probe in an exploratory study with expert designers. The goal of the technology probe was to elicit further feedback from designers regarding the general utility of search trails for their process and potential ways to organize and visualize their search information.

When defining the features of our initial technology probe, we focused on search data that could be easily captured automatically and on representations that did not involve complex processing. To this end, to help designers remember how they found the images they saved, our technology probe automatically captured the search query for each manually saved image. The prototype then annotated the resulting collection of images with an interactive list of search queries, (Figure 2 A), ordered by the number of images saved from the query. Clicking a query on the side would filter the image collection, showing only those images that were collected from it.

To provide further data on the designer's search process (to aid with recall), our technology probe also tracked how

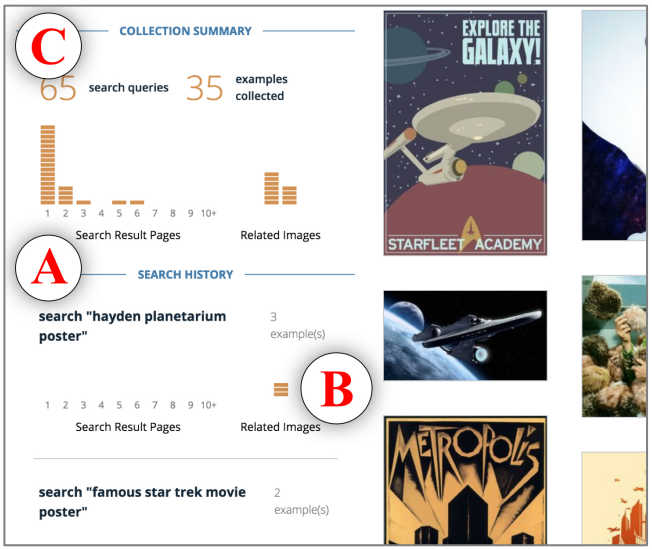

Figure 2. Screenshot of the technology probe used in our exploratory study. (A) Full list of search queries used in visual research ordered by the number of saved images per query; (B) Origin of the image (e.g., which page of search results or related images); (C) Summary of search trails for the collection

"deeply" the designer had explored each search query. Specifically, our technology probe tracked how deeply the designer dived into search results pages to find the images they liked, or if they found those images by browsing related images lists. The prototype then showed this depth of exploration in a histogram below each search query (Figure 2 B).

The prototype also showed the designer a higher-level summary of his/her visual research session in the form of search statistics, including the total number of queries and overall depth of exploration (Figure 2 C).

EXPLORATORY USER STUDY

We recruited 14 expert graphic designers (4 female) on social media (Reddit) and asked them to use our prototype for visual research for a sample design task. This design task involved creating three mood boards (collections of images that convey a general feel and style for a chosen design direction [14]) for a poster advertising a Star Trek-themed event. To help ensure that our study task resembles a real-world design scenario, a local graphic designer helped refine the phrasing and the scope of the task.

In the study, the participants created the mood boards, and then took part in a semi-structured interview. During the interview, we aimed to get a more in-depth understanding of methods and challenges of conducting visual research, and to investigate the designer's initial attitude towards the interactive search trails collected via our technology probe.

On average, the study lasted approximately one hour. Each participant received $40 in gift cards as appreciation for their time.

Findings

All participants liked the idea of using their search trails to enhance visual research. Specifically, participants felt that seeing their search trails would enable them to access their prior thought processes, which according to the literature would help them reflect on their practices [24]:

> *...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)*

Our participants provided valuable feedback about collecting and visualizing search trails. Specifically, all participants suggested that it would be much more valuable to capture detailed information about the image collection process rather than detailed information on saved images. For example, the participants said they wanted search trails to capture images beyond those that they have explicitly saved, as illustrated in the following quote:

> *So I've shown interest in this image, so I open it and then closed it. So it should be in the system. Like, there is an album or you did something it shows me every image that I open its not (P5)*

The participants did not see much value in seeing their depth of exploration for each search query, as greater depth was often simply looking for a desired perspective of composition. Similarly, the participants did not see much value in the summary of search trails provided by our prototype. They felt that this information could be easily estimated from looking at the collected images and the list of search queries.

Many participants also suggested changing the order of search queries in the sidebar to a chronological one, as they felt could help them remember their thought process, e.g., how they had iterated through search keywords to arrive at the final variation. Finally, a few participants suggested including additional filtering options for the collected images, namely filtering by colors and image resolution, as they often use these filtering options when seeking images online. Using this and other participant feedback, we designed and implemented Prism, which we describe next.

PRISM

Prism is a system that supports a graphic designer's visual research by automatically capturing all images that caught a designer's attention, and annotating them with the designer's search trails (Figure 1). 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.

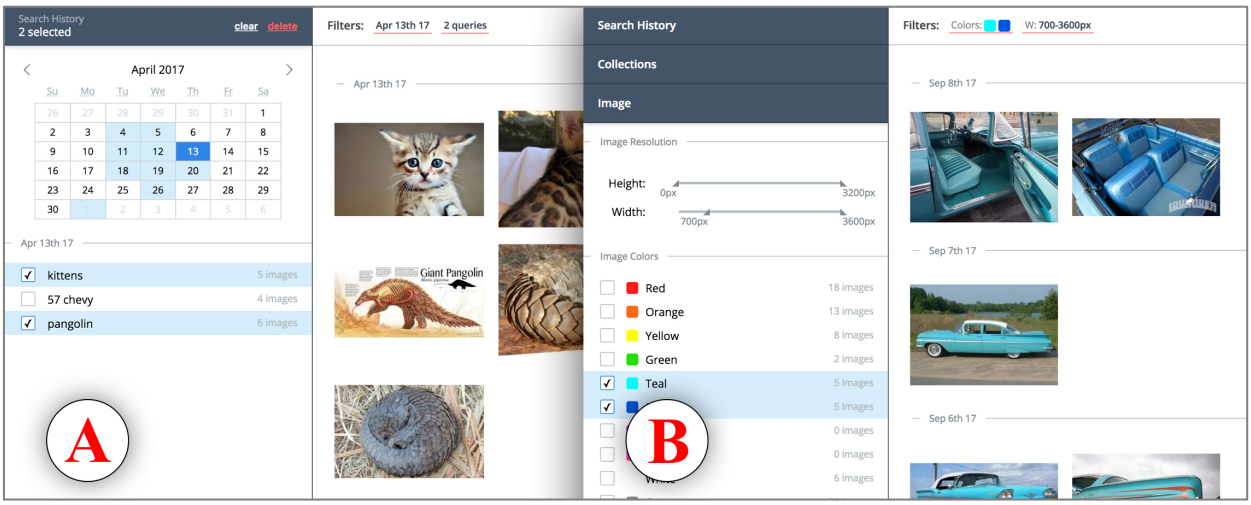

Figure 3. Examples of filtering in Prism. (A) Filtering by calendar and search queries; (B) filtering by image resolution and colors

Automating image capture via click-based inspection

In automatically capturing images, Prism aims to store images for which the designer expresses an interest. To detect a designer's "interest" in an image, Prism relies on the mechanics of clicking. Many online image repositories (e.g., Google Images and Pinterest), allow their users to click on a small image thumbnail to brings up a detailed view for that image. We decided to rely on click-based "interest" detection after we observed participants in our exploratory study clicking images that had caught their attention. 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 our discussion.

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 exploratory study suggest that designers often lose their train of thought [9]. The literature and our study results also suggest that over the long term, designers often do not remember why they saved a certain image [24]. Prism attempts to address these issues by automatically tracking *all* search queries that designers use in their visual research and annotating the collected images with these search trails.

Collection overview and filtering

In its default view, Prism displays a chronological list of all searches, split by days on the left-hand side of the image library (Figure 1 B). Such a list provides a designer with a summary of their process, and shows them how their thoughts evolved over time. Clicking on an image in the Prism library expands it and shows additional information, 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.

Given that Prism will naturally lead to designers collecting more images, the system provides a range of navigation and filtering options (Figure 3). 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 3 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 3 A for an example).

Finally, following the participants feedback from our exploratory study, we included features of filtering by image color and resolution (see Figure 3 B for an example).

Manual grouping

Prism also allows designers to manually group images into 'collections' (by accessing the "collections" tab in Figure 1). This feature was added in case designers were not satisfied with Prism's automated organizations.

Implementation

The current implementation of Prism works with Google Images, but ideally such a system could work with various search engines. 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 4).

Figure 4. Status icon for our Google Chrome extension. Left: tracking is on (red dot blinks); Right: tracking is off.

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.

FIELD STUDY

We evaluated Prism through a two-week long field study with 11 professional graphic 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.

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 graphic 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 exploratory study.

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. We did not provide any further instructions on how and when to use Prism to see how participants might naturally adopt such a tool.

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). Our interviews were open-ended in nature: our two main questions were to ask participants to describe how they had been using Prism and how their workflow with Prism was different from that without Prism. Additional probing questions were tailored to participant responses to these questions.

Data collection and analysis

We audio recorded the interviews and logged all interactions with Prism. Interviews were transcribed in full. We analyzed data from the transcripts by creating affinity diagrams using a bottom-up inductive approach [5]. We used these affinity diagrams to jointly interpret the data and to extract common themes.

FINDINGS: USAGE DATA

Throughout the study, participants collected a total of 1377 images over 466 search queries. On average, each participant 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) collected 397 images over 106 search queries, while the least active (P4) collected only 13 images across 20 search queries. More than half of the participants (6 of the 11) used Prism to collect over 100 images.

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-based 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.

FINDINGS: VALIDATING THE 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 (e.g., Pinterest). 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.

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-tail prince", but I don't like the fairy-tail 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-tail 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 trials 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)

Serendipitous or 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.

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 very 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)

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.

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 [6]. Design ideas are also often best presented by considering the space of alternatives in combination [4,7]. The fact that Prism supports these practices in a light-weight manner is encouraging.

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 very difficult to understand what their clients mean. For example, if their clients says "brick wall", they might be imagining a picture that is very 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.

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.

FINDINGS: WEAKNESSES AND OPEN ISSUES

While participants were generally very 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.

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 Freepic 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.

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 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:

> *like 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 [24]. 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.

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 very 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.

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 trials 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.

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.

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., [1,4,6]), which has often advocated approaches like sketching to enable rapid idea generation, collection and communication [4]. 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.

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. 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.

We also saw evidence of 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.

Additional search trails annotations

Our participants saw a lot of value in capturing search trials 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.

Supplementing designers' expertise with machine intelligence

Participant comments 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 their collections for their design tasks. At the same time, there are a number of ways that embedding machine intelligence could potentially further augment designers' visual research. For example, a tool 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). Highlighting such images might help designers keep track of serendipitous inspiration, or prompt them to make notes about why they inspected these conceptually different images. 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 composition.

Generalizing to long-term use and other design domains

Our findings showed that search queries and dates were sufficient for our participants to navigate their collections from two weeks of usage. However, future evaluations should also explore the degree to which search queries would help designers recount their design process after longer periods of time than covered in our study (e.g., months or even years).

Finally, this work has evaluated the utility of search trails for visual research in the domain of graphic design. More research is needed to investigate whether the approach implemented in Prism would also be beneficial in other design domains, such as animation, interior design, character design, etc.

SUMMARY

This paper investigates the idea of extending designers' visual research with automatically captured interactive search trails. We instantiated the idea in Prism – a prototype tool that automatically captures images that a designer inspects on Google Images and contextualizes them in the designer's search trails. Findings from our two-week field evaluation suggest the utility of our approach and outline promising directions for future research.

ACKNOWLEDGEMENTS

We would like to thank the National Sciences and Engineering Research Council (NSERC) and the University of Manitoba's Graduate Fellowships program for funding this research.

REFERENCES

1. Mark Baskinger. 2008. Pencils Before Pixels: A Primer In Hand-Generated Sketching. ACM Interactions, March/April, 28–36
2. Scott Bateman, Jaime Teevan, and Ryen W White. 2012. The Search Dashboard: How Reflection and Comparison Impact Search Behavior. Proceedings of the ACM Conference on Human Factors in Computing Systems, 1785–1794
3. Nathalie Bonnardel. 1999. Creativity in Design Activities: the Role of Analogies in a Constrained Cognitive Environment. Proceedings of the ACM Conference on Creativity & Cognition, 158–165
4. ill Buxton. 2007. Sketching User Experiences: Getting the Design Right and the Right Design. Morgan Kaufmann Publishers Inc.
5. Juliet Corbin and Strauss Anselm. 2008. Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory. SAGE Publications, Inc.
6. Steven P Dow, Julie Fortuna, Dan Schwartz, Beth Altringer, Daniel Schwartz, and Scott Klemmer. 2011. Prototyping Dynamics: Sharing Multiple Designs Improves Exploration, Group Rapport, and Results. Proceedings of the ACM Conference on Human Factors in Computing Systems, 1–10
7. Bill Gaver and Heather Martin. 2000. Alternatives: Exploring Information Appliances Through Conceptual Design Proposals. Proceedings of the ACM Conference on Human Factors in Computing Systems, 209–216
8. Yasunari Hashimoto and Takeo Igarashi. 2005. Retrieving Web Page Layouts Using Sketches to Support Example-Based Web Design. Eurographics Workshop on Sketch-Based Interfaces and Modeling
9. 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. Proceedings of the ACM Conference on Human Factors in Computing Systems, 87–96
10. Scarlett R. Herring, Brett R. Jones, and Brian P. Bailey. 2009. Idea Generation Techniques Among Creative Professionals. Proceedings of the IEEE Hawaii International Conference on System Sciences, 1–10
11. David G. Jansson and Steven M. Smith. 1991. Design Fixation. Design Studies 12(1), 3–11
12. Brian Lee, Savil Srivastava, Ranjitha Kumar, Ronen Brafman, and Scott R. Klemmer. 2010. Designing With Interactive Example Galleries. Proceedings of the ACM Conference on Human Factors in Computing Systems, 2257-2266
13. Julie S. Linsey, Ian Tseng, Katherine Fu, Jonathan Cagan, Kristin 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
14. Andrés Lucero. 2012. Framing, Aligning, Paradoxing, Abstracting, and Directing: How Design Mood Boards Work. Proceedings of the Designing Interactive Systems Conference, 438–447
15. Philip Mendels, Joep Frens, and Kees Overbeeke. 2011. Freed: a System for Creating Multiple Views of a Digital Collection During the Design Process. Proceedings of the ACM Conference on Human Factors in Computing Systems, 1481–1490
16. Scarlett R. Miller and Brian P. Bailey. 2014. Searching for Inspiration: an In-Depth Look at Designers' Example Finding Process. ASME 2014 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference
17. Dan Morris, Meredith Ringel Morris, and Gina Venolia. 2008. SearchBar: a Search-Centric Web History for Task Resumption and Information Re-Finding. Proceedings of the ACM Conference on Human Factors in Computing Systems, 1207–1216
18. Bernard A. Nijstad, Michael Diehl, and Wolfgang Stroebe. 2003. Cognitive stimulation and interference in idea generating groups. Group creativity: Innovation through collaboration, 137-159
19. Matti Perttula and Pekka Sipilä. 2007. The Idea Exposure Paradigm in Design Idea Generation. Journal of Engineering Design 18(1), 93–102
20. Terry A. Purcell and John S. Gero. 1991. Effects of Examples on the Results of a Design Activity. Knowledge-Based Systems 5(1), 82-91
21. Daniel Ritchie, Ankita Arvind Kejriwal, and Scott R Klemmer. 2011. D.tour: Style-Based Exploration of Design Example Galleries. Proceedings of the ACM Conference on User Interface Software and Technology, 165–174
22. Donald A Schon. 1984. The Reflective Practitioner: How Professionals Think in Action. Vol. 5126. Basic books
23. Moushumi Sharmin and Brian P Bailey. 2013. ReflectionSpace: an Interactive Visualization Tool for Supporting Reflection-on-Action in Design. Proceedings of the ACM Conference on Creativity & Cognition, 83–92
24. Moushumi Sharmin, Brian P. Bailey, Cole Coats, and Kevin Hamilton. 2009. Understanding Knowledge Management Practices for Early Design Activity and its Implications for Reuse. Proceedings of the ACM Conference on Human Factors in Computing Systems, 2367–2376

25. Andrew M. Webb, Rhema Linder, Andruid Kerne, Nic Lupfer, Yin Qu, Bryant Poffenberger, and Colton Revia. 2013. Promoting Reflection and Interpretation in Education. *Proceedings of the 9th ACM Conference on Creativity & Cognitio*n, 53–62

26. Ryen W. White and Steven M. Drucker. 2007. Investigating Behavioral Variability in Web Search. *Proceedings of the 16th Conference on World Wide Web,* 21–30

27. Ka-Ping Yee, Kirsten Swearingen, Kevin Li, and Marti Hearst. 2003. Faceted Metadata for Image Search and Browsing. Proceedings of the ACM Conference on Human Factors in Computing Systems, 401–408