

# Designing Novel Image Search Interfaces by Understanding Unique Characteristics and Usage

Paul André<sup>1</sup>, Edward Cutrell<sup>2</sup>, Desney S. Tan<sup>2</sup>, and Greg Smith<sup>2</sup>

<sup>1</sup> ECS, University of Southampton, UK  
pa2@ecs.soton.ac.uk

<sup>2</sup> Microsoft Research, Redmond, WA, USA  
{cutrell, desney, gregsmi}@microsoft.com

**Abstract.** In most major search engines, the interface for image search is the same as traditional Web search: a keyword query followed by a paginated, ranked list of results. Although many image search innovations have appeared in both the literature and on the Web, few have seen widespread use in practice. In this work, we explore the differences between image and general Web search to better support users' needs. First, we describe some unique characteristics of image search derived through informal interviews with researchers, designers, and managers responsible for building and deploying a major Web search engine. Then, we present results from a large scale analysis of image and Web search logs showing the differences in user behaviour. Grounded in these observations, we present design recommendations for an image search engine supportive of the unique experience of image search. We iterate on a number of designs, and describe a functional prototype that we built.

**Keywords:** image search, log analysis, design.

## 1 Introduction

Most major search engines treat general Web search and image search very similarly. A searcher types in one or more keywords into a text box and results appear in a paginated view ordered by ranking algorithms that are largely based on textual features. While the experience for image and Web search are similar, we believe that there are important differences that are not well-recognized or exploited by search engine providers. The goals of searchers using image search are likely quite different from those using general Web search. In addition, images occupy a high-dimensional semantic and visual space that could be used to improve retrieval and the presentation of/interaction with search results.

Researchers have introduced many innovations in retrieval for image search (e.g., textual and visual features [2], distance measures [20], clustering [21], query-by-image, etc.), although few of these techniques are used in practice in major image search engines. In contrast to the excellent work in retrieval techniques, there has been little research in understanding image search behaviour, particularly in comparison to Web search. We believe we can contribute more immediate value by recasting

existing technologies deployed for image search within a new interface and interaction model designed for the unique aspects of the image search experience.

In this paper we explore ways to design more compelling image search experiences by sharing two main contributions. First, we conducted informal interviews with researchers, designers, and managers working to develop and deploy a major commercial search engine. From these interviews, we derived four main characteristics that make image search unique from its Web search counterpart. Following this, we also report on a large-scale analysis of query logs to characterise some of the differences between image and Web search in the user activity of a major commercial search engine. Second, based on these observations, we generate design recommendations for interfaces tuned for the needs of image search. To exemplify these recommendations, we describe a functional prototype interface for image search that we designed to support the unique features of image search.

In the following sections, we review the literature and some existing image search engines. We then describe the methodology and results of our interviews and query log analysis. We detail how these results form the basis for the design of a new image search interface. We conclude with recommendations for evaluation and future work.

## **2 Related Work**

### **2.1 Image Search Engines**

The major Web search engines (e.g. Google, Yahoo, Baidu, Microsoft Live, Ask) all provide a similar experience for image search: keyword-based query resulting in a grid of image thumbnails. Various query refinements are available, including image size, aspect ratio, colour and different kinds of content (e.g., illustrations, or images containing a face).

Besides the major search engines, there are a wide range of specialised, original or experimental image search engines on the Web (we counted at least twenty-one and imagine there are many others). Though space precludes a full listing, three of the most interesting examples are: captivating colour search [13], in which images are returned based purely on one or two chosen colours; Viewzi [19], which presents results in a number of ‘views’ – flipping the pages of a cookbook when searching for recipes, a combined photo and tag cloud, an Amazon book view; and Getty Catalyst search [8], which leverages rich ontological and tag metadata to provide easy refinement of searches.

While there are a number of innovative image search engines, they tend to either be purely aesthetic-driven or cater to niche communities, neither of which increases functionality for the average image searcher. Our goal in this work is to explore a design space that allows appealing aesthetic treatment while still supporting greater efficiency in the task-centric behaviour seen in the larger search engines.

## 2.2 Improving Retrieval and Presentation of Images

The majority of academic work on image search has focused on back-end technologies. Such improvements include similarity space projection [16] in which images and queries are projected into a third space such that relevant images are kept close to the corresponding query; a framework to reduce the semantic gap in content-based image retrieval [20]; and automatic image annotation through a domain-specific ontology [7]. Though all are promising, these techniques can face serious challenges to Web scale implementation, and as far as we can tell, none have seen commercial use.

A number of projects utilise new user interfaces to exploit work on the back-end. In many cases, these comprise different visual representations of various clustering techniques [2,3,5,21]. Other innovations in feature extraction have led to different interfaces. Liu et al. [15] display a one page overview of results, with thumbnails cropped to regions of interest based on low-level features, and similar images displayed near each other. In Visual Islands [25], raw features are mapped to a 2D space, and the images are snapped to a grid maintaining spatial layout of the 2D space. EnjoyPhoto [26] presents a subset of 'high-quality' images scraped from forums with rich metadata, and the user interface utilises fish-eye and force transfer implementations for in-place browsing.

CueFlik [6] is an interactive machine-learning interface to create rules for re-ranking images based on image characteristics extracted from representative examples the searcher provides. IntentSearch [4] categorises a query image into one of several predefined intention categories, with a specific similarity measure used inside each category to combine image features for re-ranking. At the end of 2008, Microsoft Live Image Search released a feature allowing a user to search for 'similar images to this'. However, the user interface remained essentially unchanged.

To our knowledge, few papers have focused purely on supporting the image search user experience through novel user interfaces. One exception to this is work by Porta [18], which describes several techniques for presenting all images within a collection in a short time to support users who do not know precisely what to search for.

In summary, although there are many innovative image retrieval technologies in the literature, the major search engines have been slow to adopt them. In order to provide a Web-scalable prototype that performs well for most scenarios of image search (i.e. the long tail), we chose not to focus on complex retrieval algorithms, but instead on presenting current search technologies within a new user interface and interaction setting that is specifically designed for the unique qualities of image search. However, to better understand those unique qualities, it is necessary to understand how users actually use current technologies to search for images.

## 2.3. Image Search Usage

Previous work on understanding image search *strategies* has largely focused on classifying the kinds of queries posed by professionals. Analysis of image requests of journalists [17] found that over half were for a named person or object, and that browsing was an essential strategy, though it was not generally well supported. Simi-

larly, a study of a newspaper image archive [22] indicated that around 40% of requests were for a named person, with a significant number of requests for general objects or themes. In a study of queries to a commercial database by professional users [14], proper nouns accounted for only 9% of queries, whereas nouns accounted for around 50% of searches, and again browsing and exploration were seen to be under-supported.

The most relevant work in understanding image search *behaviour* was undertaken in 2001 by Goodrum and Spink [10] using the Excite search engine. They found that image queries contained on average 3.74 words (though at least one of these was the image request term, e.g. ‘picture’ or ‘image’). They also reported a high percentage of unique terms, and that modification of queries was common (accounting for 60% of queries). An analysis of image search patterns [9] found that the most commonly occurring patterns were rapid browsing of thumbnails and individual images. Hung [12] analysed search activities of journalism students and found that most activities in all tasks were related to browsing and enlarging images. Finally, a recent query log analysis [26] showed that more than 20% of image search queries are related to location, nature, and daily life categories.

The above work underlines the importance of flexible query modification and rapid browsing in the practice of image search. Like general Web search, image search appears to be used by a wide variety of people in support of many different tasks. However, much of the detailed query analysis was done on search logs that were collected at the dawn of the Web. Current commercial Web search engines now have dedicated image search verticals, and the number and variety of people using these services is dramatically larger than it was in 2000. As a result, we believe it would be helpful to explore how Web search and image search differ with current interfaces.

### **3 Understanding Characteristics and Usage of Image Search**

#### **3.1 Informal Interviews to Characterise Unique Goals and Intent**

To better understand the fundamental characteristics of image search, we spoke with 8 researchers, designers, and managers associated with Live Image Search. From these discussions, we derived 4 characteristics of image search that may be relevant for interface design.

- 1) The search results themselves may fulfil a user’s needs. Image search, like general Web search, can be both exploratory (information about art at the Louvre) or goal-directed (looking for a specific fact or image). In some cases, the ‘answer’ to the query can be found in the results page for either image or Web search and a user need never click on any result. One might find the answer to the query, “When was the Louvre built?” in the text of result snippets in Web search. Similarly, the query “What does the pyramid at the Louvre look like?” could be found just by examining a page of thumbnails returned by an image search.
- 2) Image search is often more exploratory than Web search. Searchers may be looking for an image with a particular visual ‘style’, with a predefined ‘type’ in mind, or with certain characteristics, and they may not be able to express those re-

quirements until the desired image is found. This type of exploratory search within the rich visual feature space of images is not well supported today.

- 3) Image search is often used purely for entertainment to play or explore in visual space with no end goal. We see examples of this in Flickr's 'Interesting Photos' page, or the playful image search sites that allow search by something abstract such as colour [13]. The inherently visual nature of image search lends itself particularly well to the artful presentation of results.
- 4) Image search encourages tangents. Related to the previous point, the visual nature of image search makes it very easy to become sidetracked when something else of interest catches their eye, even if the initial query is task-focused. And if a searcher is browsing for entertainment, this kind of tangent is actually *desired*.

### 3.2 Search Log Analysis to Understand User Behaviour

Building on findings in the interviews, we conducted large scale analysis of search logs from the same search engine. Results presented in this paper are derived from a one day sample (October 10<sup>th</sup> 2008), allowing a broad view of over 55 million queries. The logs contain anonymised information about the query, the number of result pages seen, the number and location of clicks, and time spent looking at results. To remove variability caused by geographic and linguistic variation in search behaviour, we filtered to only queries generated in the English speaking United States ISO locale.

We examined search activity at two different levels of granularity: the query level (i.e., an individual query with a number of result page views), and the session level (i.e., multiple queries within a single session). We demarcated sessions by 30 minutes of no activity (as used in [24]). For both image and Web search, we explored similar questions at each level of granularity.

Table 1 shows the query level statistics comparing image and Web search. The statistics for a single query incorporate all interactions within that query (e.g., clicking a result, clicking to see the next page), and are reset when a new query is issued. First, we see that the average number of result clicks per query is higher for image search than Web search. Second, the average depth (number of consecutive results pages viewed for a single query) for image search is more than twice that of Web search<sup>1</sup>. If we filter to look only at the queries that had result clicks, thereby excluding all queries in which the user just looked at results on the results page without clicking, then we see that 30% of all queries for image search contained a result click, significantly lower than the 43% for Web search. However, when a query did contain a click, image queries had more clicks than Web queries.

Overall, the dwell time (i.e., the amount of time a user spends viewing a page of search results) was significantly higher for image than Web queries.

---

<sup>1</sup> It should be noted that in Live Image Search, pagination is implemented through a continuous scrolling mechanism, so total depth may be somewhat inflated over search engines with a strict click-based pagination. However, the general trend of increased average depth confirms intuitions in discussions and the difference is quite dramatic.

**Table 1.** Comparison of search log statistics at the query level, for image and Web searches.

<i>Query level</i>	<b>Image</b>	<b>Web</b>
Avg. depth	3.47	1.25
Median depth	2	1
Avg. clicks	0.73	0.63
Median clicks	0	0
% with > 0 clicks	30%	43%
Avg. clicks	2.5	1.5
Median clicks	1	1
Avg. dwell time (s)	85	52

Table 2 details the session-level statistics for the same measures shown in Table 1. A session consists of a set of queries issued by a user in near temporal proximity. The session ends when there is a period of inactivity of at least 30 minutes between queries for a given (anonymous) ID. At this level we are able to gain a broader view of user behaviour, focusing on the number of queries issued within a session, and which of those queries contained clicks on search results.

At the session level, we see that, on average, more queries are issued and more results are clicked in an image search than a Web search session. Second, even though a higher proportion of image search sessions resulted in no search result clicks, the average number of clicks for image search was significantly higher than for Web search among sessions with a click. Finally, echoing findings at the query level, the average dwell time for a session was significantly longer for image than Web search.

**Table 2.** Comparison of search log statistics at the session level, for image and Web searches.

<i>Session level</i>	<b>Image</b>	<b>Web</b>
Avg. queries/session	2.45	1.59
Median queries/session	1	1
Avg. clicks per session	1.87	0.96
Median clicks/session	0	0
% with > 0 clicks	35%	42%
Avg. clicks	5.4	2.3
Median clicks	3	1
Avg. dwell time (s)	312	109

There are several reasons that multiple queries may appear in a session, including different kinds of query refinements and consecutive unrelated searches. In order to understand this behaviour, we manually examined query text from around 1000 sessions chosen randomly from both image and Web search logs that contained sessions of two or more queries. We term these ‘query trails’. We categorised them into 3 groups, those containing *unrelated*, *tangentially* related, or obviously *related* queries. Though there is some overlap between these categories and the classification is subjective, this scheme provides a reasonable estimate of the progression of search terms within a session. Table 3 displays the breakdown of classification, with examples.

**Table 3.** Classification and examples of search log queries.

<i>Query trail classification</i>	<b>Image</b>	<b>Web</b>	<b>Example</b>
Unrelated	20%	35%	reebok → family guy; fitness → deal or no deal
Tangential	10%	5%	horned lizard → horned lizard desert → desert color
Related	70%	60%	spiders → hawaiian spiders → hawaiian sugarcane spiders

The majority of query sessions for both image and Web search comprise a trail of *related* queries, although there are somewhat more for image than for Web search. In contrast, Web search has a larger proportion of trails with *unrelated* queries. This suggests that when searchers turn to image search, they tend to explore a single line of questioning or topicality, iterating on related queries. This is less true for Web search, where searchers are more likely to engage in unrelated queries within a session.

### 3.3 Summary and Analysis of Results

The findings from our interviews and query log analysis suggest several interesting implications for the design of interfaces for Web-based image search. In both query- and session-level search statistics, image searchers view more pages of search results (average depth), they spend more time looking at those pages of search results, and they click on more results than Web searchers. This implies that either query relevance is not as important for image as for Web search or that query relevance is simply much worse for image search than Web search. We hypothesize that the increased click average is mainly due to two factors. The first is that there is often no definitive answer to an image query – the sought after image could be one of many, and it is a subjective result. The second is that we believe image thumbnails offer visually appealing objects that searchers may click on solely for the aesthetic value.

Compared to Web searches, image searches contain a greater number of queries that have no result clicks. However, when image searches do contain clicks, they contain significantly more. We hypothesize that the reason there are more queries with no clicks is that the answer (or lack of) is contained in the results page itself (i.e. low relevance is obvious and no result clicks are necessary, or a question such as ‘what does a passion fruit look like’ is answered immediately), or the user is merely browsing for enjoyment. The greater number of tangential query progressions implies that there is a diverse range of results returned that inspire a new search and/or image searchers are more conducive to going ‘off-task’ than Web searchers. The large number of related query trails shows that refinement and iteration are important in both image and Web search. However, we believe that a facility for such refinement is particularly important in image search where a goal (a mental picture of the desired image) may be far more difficult to specify with traditional keyword queries.

In the following section we consider how the findings from our log analysis, as well as from interviews and related work, can be used for design recommendations.

## 4 Design Suggestions

Based on our findings, we derive design suggestions that range from pure aesthetics to specific functionality in order to better support new image search experiences.

### 4.1 Supporting Exploration

**Exploratory interaction paradigm.** Our analysis suggests that image searches are more exploratory than Web search: more search pages viewed, 80% of query trails directly or tangentially related, and looking for a non-definitive answer within a multi-featured visual and semantic space. The existing page-based grid view of images, though clean and simple, does not visually or conceptually convey the user's path through this space. A literal representation of such a space would likely be overwhelming (e.g., lists of features mapped onto a multi-dimensional representation). However, any simple metaphors or concepts to change the interaction and interface in a way that supports the exploration of associated terms, or image-based queries, would be an advantage. Because image searchers tend to view many more results (more queries per session, more page views per query, and more clicks per session), support for rapid browsing would also be desirable. A highly interactive system that immediately displays more results is very important.

**Entry point.** Entry points for the major image search engines are virtually indistinguishable from Web search: a query box on a clean, blank page. Such a practice reinforces a task-centric approach to image search; a searcher is offered a blank slate that they must fill with a keyword expression. Exploratory search is better served by allowing searchers to frame a search in various ways. One way to do this is to pre-populate the opening page with sample images, just as some “experimental” image search sites do. Such images may be popular pictures, things taken from current news, or based on former searches (e.g., the new tab shown by the Google Chrome Web browser [11]). A searcher can easily launch a new query, but these opening images also give a starting point for non-directed browsing. We believe this would provide a more compelling visual experience, an opportunity to learn about current news stories or popular searched for topics, and increase the chance of serendipitous findings [1].

### 4.2 Fun and Aesthetics

In contrast to Web search which is dominated by text, image search is overtly visual. We believe that the plain grid view of current search engines fails to adequately exploit this fact. Because the search results themselves are visual artefacts, there are a variety of visual layout techniques that, while not necessarily enhancing functionality *per se*, can provide a beautiful aesthetic and a more compelling search experience.

As noted in section 3.1, a common use of image search is pure image *browsing*, whether for general entertainment or for specific information needs. As such, aesthetics are likely very important compared to other search engines that are used for more utilitarian tasks.

### 4.3 Query Refinement

The image search logs show a significant amount of query reformulation towards a specific goal (around 70% of all queries). Such reformulation supports our belief that an image searcher generally has a good idea of the type or style of the desired result prior to issuing their query. Though this may also be true of Web search, in image search, it is often much more difficult to express this textually. For example, an image searcher may be looking for the ‘mood’ or framing of an image, for the dominant colour or perhaps a particular pose or position of a subject.

Some requirements can be expressed using existing refinements, such as aspect ratio, or black and white. Others require more complex features of the image (e.g., edge histograms, colour histograms, face detection, etc.). Some existing image search sites allow query by image, or a ‘more like this’ query. Because of the mental model that searchers have of images, and the extensive reformulation as seen in the logs, such features should be considered essential refinement options.

Refinement based on semantics should also be considered. In this context, semantics of an image could be acquired from the labels it has in the image database, the query terms that led to clicks on that image, or user-generated tags. Since tags can be at different granularities or conceptual levels, they allow for another type of refinement: broadening, narrowing, or even ‘side-stepping the current query. They may also be used alongside visual features in a ‘more like this’ query.

**Query and Refinement History.** The easy refinement proposed above addresses the concern that current image search engines do not allow users to fully *explore* the rich visual and semantic space that images inhabit. However, such simple refinement and exploration can quickly lead to a complex history that is hard to retrace. A searcher may well ask, “How did I get here and how do I find my way back to that good content I saw a few minutes ago?” Therefore we need to design a history that details the current path as well as a means to navigate that history. Within a session, this would represent explicit queries (i.e. text typed by a user) and all the refinements that may have been added (e.g. metadata filtering, more like this image). Between sessions this could act in a similar way to Web browser history and allow revisitation of previous queries.

### 4.4 View and Save Images

Since about twice as many image search results are clicked in a session than Web search results, it is worth considering how one might support easily viewing many images (as discussed above in exploratory search), and saving them for later viewing. Though we have spent time considering the fun and aesthetics involved in image search, there is still a core productivity-oriented component. For such tasks, where the end goal is important, it may be useful and efficient to allow saving of particular images from a result set, allowing the user to later revisit the saved set. Such a feature could be integrated with the within- and between-session histories.



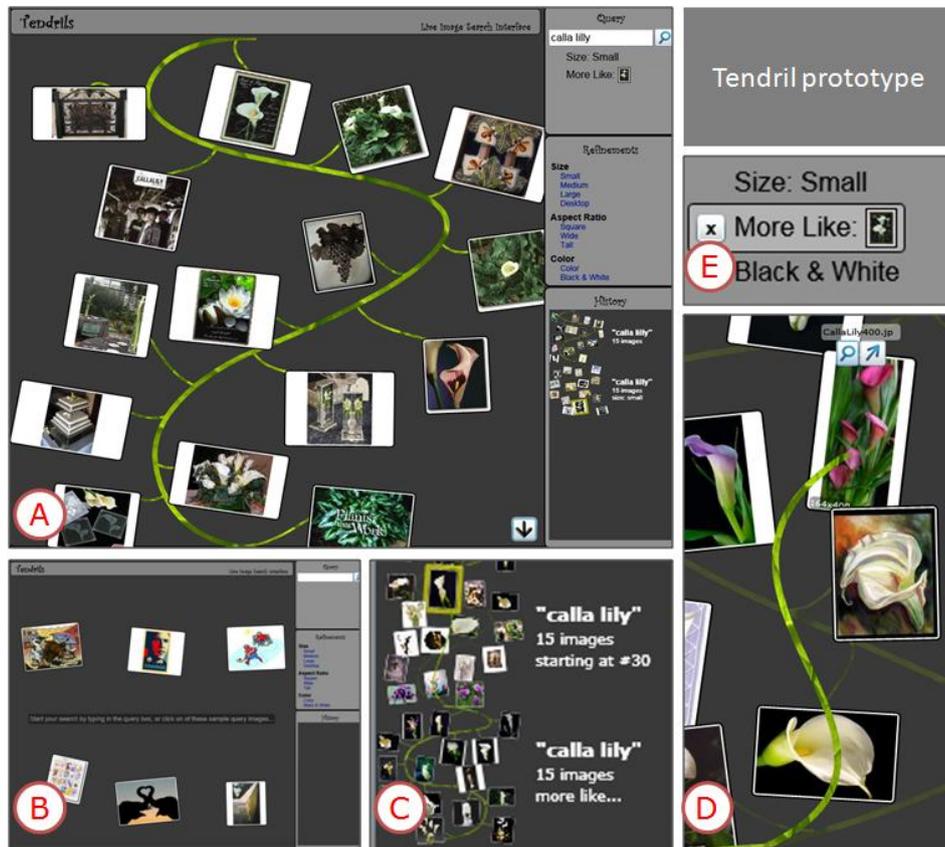
## 5.1 Final Design

The design we chose to prototype incorporates many of the features distilled from related work, interviews, and from log behaviour. Of course, this design is not intended to be ‘complete,’ but an interpretation or instantiation of the design suggestions in section 4. We highlight some of the limitations and future work in section 6. Our design emphasizes the importance placed upon query modification seen in previous designs and highlights *exploration* of visual space, which unlike most of the previous designs is also alluded to in the presentation of results. Nicknamed ‘Tendrils’, it uses a metaphor of tendrils (or ivy vines) for an organic display of images. The user seeds a search, and explores the resulting path (tendrils) through the space. The overall space that a tendrils could ‘grow’ into is infinite, and so we can imagine a true exploration of that space. Alterations or modifications to the query result in branching off the initial path. Figure 2 shows the interface and selected details from the working prototype, available (to reviewers) at <http://thegdp.net/tendrils>.

**Layout and Launch Page.** The right hand side of the screen is used to display the search box and current query, history, and refinement possibilities. In the main area, a selection of images on the front page is used to highlight the exploratory (and not pure keyword) search experience, as well as to facilitate serendipitous browsing (Fig. 2-B).

**Result Display.** To emphasise the exploratory nature of the search engine, and to offer a distinctive and compelling experience, the result display takes the form of a ‘tendrils’ extending across the screen, leading the searcher through the imagined vast result space (Fig. 2-A). Conceptually this is grounded in exploration of unseen semantic and visual space, something that in later iterations we hope to convey more obviously with the camera panning out to see the ‘world view’ before zooming in again on a particular part.

**Context Menu and Query Modification.** Hovering over an image presents a context menu with familiar actions such as ‘view’, and also enables query modification so prevalent in related work and in our log analysis. Query modification is a form of pivoting, and could be used to find an image ‘more like this’, or to use tags or other metadata to broaden/narrow/sidestep the current search. We have chosen to expose the ability to search for ‘more like this’ (since it was available through Live Image Search XML), and we have specifically designed the user interface to promote and take advantage of such a feature. The user can select to search for more like this from the context menu, or (soon) drag an image to the search box. Such a feature is especially useful in those cases where users may have an idea of what they are looking for but are unable to express it in words. Semantic exploration in the form of keywords (either taken from an image database, or from Flickr or a similar service) is also desirable. While examples such as Getty and Flickr have this metadata, current search indexes do not, or do not expose it in an efficient manner. In the future, we would like to include it in the manner described in Fig. 1-B/C.



**Fig. 2.** The Tendril design concept. (A): a full page with query and results. (B): the launch page with a selection of random images. (C): enlarged view of history. (D): example of branching after querying for 'more like this'. (E): enlarged view of query area. (Prototype at <http://thegdp.net/tendril>)

**Search Box.** Along with the initial query, this area acts as a control area for query modifications (Fig 2-E), where refinements or 'similar images' added to the query are shown. This is also the region into which related images are dragged for refinement.

**Branching.** As modifications are made to the query, the previous path through the search space is no longer valid, and so a new path is 'branched' off. An example of this is seen in Fig. 2-D, where a new query has been issued based on an image.

**History.** With the ability to easily explore and refine, a complex path could soon be built. The history allows a user to see the path to their current search results, and to 'step back' if, for instance, the latest modification was not to their liking. A miniature version of the tendril on that page (each tendril is unique) provides a representation that could be used to aid navigation back to a previously seen page (Fig. 2-C).

## 5.2 Implementation Details

The Tendril prototype is implemented as a Silverlight 2 application running in the browser and an associated Windows Communication Foundation Web service, both hosted on an ASP.NET Website. Queries and refinement clauses entered into the interface are translated into query options and submitted asynchronously back to the Web service on the Tendrils Website. To utilise existing functionality (notably the ‘search for similar images’ option) in Live Image Search, the service builds the appropriate query string in Live Image Search syntax and submits the query to `search.live.com/images`, asking for a response in XML format. The XML response is returned to the Silverlight application and parsed for the individual search results. The application then submits additional asynchronous requests to the Tendrils Web service for a thumbnail for each result, and inserts the thumbnails into graphical templates in the Silverlight interface as they are returned. It is worth noting that designing this implementation as a novel front-end that sits on a fully functional search engine allows us to tap into the vast and growing index of images, as well as the functionality offered by that engine.

## 6 Limitations and Conclusions

In this paper we have considered how understanding the unique characteristics and usage of image search can contribute to the design of a focused and compelling image search experience. We offer two main contributions. First, we offer a qualitative and quantitative analysis of differences in behaviour between image and Web search. Second, we use those insights to ground recommendations for the design of image search experiences. Additionally, we present five designs grounded in these recommendations including a full working prototype.

While Web log analysis provides a broad description of what people do, we do not have the context of their experience. In the future, we would like to observe and interview a variety of individual image searchers to better understand their needs and behaviour in specific contexts of use.

We continue to refine our prototype, in terms of both performance and aesthetics. In the future we would like to take one of our other design concepts to prototype, and do a formal evaluation of use against each other, and against commercial search engines, to determine exactly if and how our design suggestions improve the image search experience and goals therein. These concepts are only a few examples in a world of potential new designs in image search. We hope that the analysis and design recommendations in this paper will inspire others to consider improving holistic experience as well as retrieval in their own designs for image search on the Web.

**Acknowledgments.** We would like to thank Dan Liebling for data gathering and analysis assistance, and Steven Drucker, Gonzalo Ramos, Gang Hua, and the many others we talked to for various helpful and informative discussions.

## References

1. André, P., Teevan, J. & Dumais, S. T. From X-Rays to Silly Putty via Uranus: Serendipity and its Role in Web Search. In CHI '09. (In Press)
2. Cai, D., He, X., Li, Z., Ma, W., & Wen, J. Hierarchical clustering of WWW image search results using visual, textual & link information. In MM '04, pp.952—959. (2004)
3. Chen, Z., Wenyin, L., Hu, C., Li, M., & Zhang, H. iFind: a Web image search engine. In SIGIR '01, p.450. (2001)
4. Cui, J., Wen, F., & Tang, X. Real time google and live image search re-ranking. In MM '08, pp.729—732. (2008)
5. Ding, H., Liu, J., & Lu, H. Hierarchical clustering-based navigation of image search results. In MM '08, pp.741—744. (2008)
6. Fogarty, J., Tan, D., Kapoor, A., & Winder, S. CueFlik: interactive concept learning in image search. In CHI '08, pp.29—38. (2008)
7. Gao, Y., Luo, H., & Fan, J. Searching and browsing large scale image database using keywords and ontology. In MM '06, pp.811—812. (2006)
8. Getty Catalyst Search, <http://www.gettyimages.com/Catalyst/Default.aspx>
9. Goodrum, A. A., Bejune, M. M. & Siochi, A. C. A State Transition Analysis of Image Search Patterns on the Web. In CIVR '03, LNCS 2728, pp.281-290. (2003)
10. Goodrum, A. & Spink, A. Image searching on the Excite Web search engine. *Inf. Process. Manage.* 37, 2 (Mar. 2001), 295-311.
11. Google Chrome Web Browser, <http://www.google.com/chrome/>
12. Hung, T-Y. Search Moves and Tactics for Image Retrieval in the Field of Journalism: A Pilot Study. *J. of Educational Media & Library Sciences.* 42, 3, 329-346 (2005).
13. Idée Colour Search, <http://labs.ideeinc.com/multicolr/>
14. Jørgensen, C. & Jørgensen, P. Image querying by image professionals. *Journal of the American Society of Information Science & Technology.* 56, 12, 1346-1359 (2005).
15. Liu, H., Xie, X., Tang, X., Li, Z., & Ma, W. Effective browsing of Web image search results. In MIR '04, pp.84—90. (2004)
16. Liu, Y., Qin, T., Liu, T., Zhang, L., & Ma, W. Similarity space projection for Web image search and annotation. In MIR '05, pp.49—56. (2005)
17. Markkula, M. & Sormunen, E. End-User Searching Challenges Indexing Practices in the Digital Newspaper Photo archive. *Information Retrieval.* 1, 4, 259-285 (2000)
18. Porta, M. Browsing large collections of images through unconventional visualization techniques. In AVI '06, pp.440—444. (2006)
19. Viewzi Search Engine, <http://viewzi.com>
20. Wang, C., Zhang, L., & Zhang, H. Learning to reduce the semantic gap in Web image retrieval and annotation. In SIGIR '08, pp.355—362. (2008)
21. Wang, S., Jing, F., He, J., Du, Q., & Zhang, L. IGroup: presenting Web image search results in semantic clusters. In CHI '07, pp.587—596. (2007)
22. Westman, S. & Oittinen, P. Image Retrieval by End-Users and Intermediaries in a Journalistic Work Context. In IliX '06, pp.171—187. (2006)
23. Westman, S., Lustila, A., & Oittinen, P. Search strategies in multimodal image retrieval. In IliX '08, vol. 348. (2008)
24. White, R.W. & Drucker, S.M. Investigating behavioral variability in Web search. In WWW '07, pp.21—30. (2007)
25. Zavesky, E., Chang, S., & Yang, C. Visual islands: intuitive browsing of visual search results. In CIVR '08, pp.617—626. (2008)
26. Zhang, L., Chen, L., Jing, F., Deng, K., & Ma, W. EnjoyPhoto: a vertical image search engine for enjoying high-quality photos. In MM '06, pp.367—376. (2006)