

SEARCH BEHAVIOUR ON PHOTO SHARING PLATFORMS

Silviu Maniu*

Neil O’Hare, Luca Maria Aiello, Luca Chiarandini, Alejandro Jaimes

University of Hong Kong
Hong Kong SAR, China
smaniu@cs.hku.hk

Yahoo! Research
Barcelona, Spain
{nohare, alucca, chiarluc, ajames}@yahoo-inc.com

ABSTRACT

The behaviour, goals, and intentions of users while searching for images in large scale online collections are not well understood, with image search log analysis providing limited insights, in part because they tend only to have access to user *search* and *result click* information. In this paper we study user search behaviour in a large photo-sharing platform, analyzing all user actions during search sessions (i.e. including post result-click pageviews). Search accounts for a significant part of user interactions with such platforms, and we show differences between the queries issued on such platforms and those on general image search. We show that search behaviour is influenced by the query type, and also depends on the user. Finally, we analyse how users behave when they reformulate their queries, and develop URL class prediction models for image search, showing that query-specific models significantly outperform query-agnostic models. The insights provided in this paper are intended as a launching point for the design of better interfaces and ranking models for image search.

1. INTRODUCTION

Photo sharing platforms such as Flickr or Instagram are increasingly popular and, similarly to online social networks, they support activities such as sharing their photos with friends and forming common-interest *groups* in which user can usually join freely to share multimedia content with the other members. Such platforms also support image search; previous work showed that over 2% of page-views in Flickr are accounted for by searches [1], and effective search performance is arguably important for the long-term success of such platforms.

If the goals of users in general web image search are not well understood, they are even less understood on photo sharing platforms, where there little work on user search behaviour has been published. On the other hand, the server logs of such platforms give us access to *entire* user search sessions, including all post search interactions, not just the *search* and *result click* interactions available in search engine logs. This

gives us the opportunity to come to a deeper understanding of what users do after issuing a search.

In this paper, we study the search behaviour of users of a large online photo sharing platform, namely Flickr. We study the typical types of search conducted on such platforms, and note some differences from general image search. We look at the entire user session after an initial keyword search, with a view to uncovering behaviour patterns that go beyond simple “search and click on result” events. Modeling browsing behaviour using *search trees*, we show how search behaviour is influenced by *query type* and by *user*, in that certain types of query or users show exhibit behaviour. We go on to show that URL-class prediction models trained on different categories of searches – such as query types and user types – perform better than prediction models trained on the entire data, emphasizing that these behavioural differences can have a predictive power.

We review the related work in search log analysis and image search behaviour in Section 2. Then we describe our dataset and how it is processed in Section 3. The taxonomy of queries is presented in Section 4, followed by an analysis of search behaviour, focusing in particular on how search behaviour varies according to query type and user type, in Section 5. Finally, we conclude the paper in Section 6.

2. RELATED WORK

There has been much work on analysing the logs of commercial web search engines, uncovering relationships between queries [2] and using log analysis to improve search engine rankings [3]. Broder [4] proposed three distinct types of queries based on user intent: informational, navigational and transactional. Other work has automatically classified queries within this taxonomy [5].

Studies of user behaviour using web server logs are often limited by the fact that the logs only record interactions with the search engine itself, with subsequent actions not recorded in the logs. White & Drunker [6] circumvented this problem by inviting users to install a browser plug-in which logged all their browsing activities, and analysed the entire search sessions of over 2,000 participants, characterizing users based on *search trails*, similar to our *search trees*. The availability of

*This work was carried out while Silviu Maniu was visiting Yahoo! Research Barcelona for a research internship.

tabbed browsing on modern web browsers means web browsing sessions are rarely linear, and models for tabbed browsing have been proposed by Chierichetti et al. [7]. Much of the work on understanding image search behaviour has focused on professional users, using a combination of qualitative methods and automatic analysis of search logs [8]. Such studies tend to show that a variety of search strategies are used, and that browsing and exploration are often important strategies [9].

Studies of image search web server logs include Jansen et al [10], who analysed audio, video and image searches from the Alta Vista search engine. Andre et al. [11] analyse a large image search log and note that, compared with general web search sessions, image search sessions have greater average depth (number of results pages clicked for a query), they have more results clicked, and users spend more time looking at results pages, inferred from this that image search is more exploratory than web search. Other work has studied taxonomies for image search, attempting to adapt Broder’s [4] taxonomy of web search to image search [12].

3. DATASET AND PREPROCESSING

From the Flickr web server logs, we take as a sample the pageviews of a very large set of anonymous, randomly selected users, during 2011. We split a user’s page views into sessions when the inactivity between two page views is longer than 25 minutes or when the user leaves the Flickr website. Since we are interested in studying search behaviour, we focus on sessions containing at least one search action. To generate a linguistically and culturally homogeneous dataset, we only consider sessions from US IP addresses. We also remove extremely long sessions whose length above the 99th percentile giving a dataset of approximately 1 Million sessions.

The logs contain a record for each HTTP request received by the server, and include *userId*, *timestamp*, *url* and a *referral url*. Often, multiple URLs can map to exactly the same page “layout” (e.g., in Flickr, “display photo” and “display photo stream” correspond to specific page layouts). As in previous work, we manually created a set of regular expressions to classify the URLs into 96 different classes [1]. The highest frequency URL classes within these search sessions are listed in Table 1. Searches and photo views account for approximately 50% of page views in search sessions, suggesting that, to fully understand user interactions during search sessions, we need to look at the other 50% of page views, which describe user interactions which diverge from search and result-click.

3.1. Search Trees

Since sessions are not strictly linear in nature, due to backtracking (use of the ‘back’ button) and branching (use of tabs) behaviour, we represent search sessions as *search trees*. The first search action in a session is the root of the tree and, for each subsequent pageview in the session, we create a node

Class	Description
search, search/next photo	Search for photos and next result page Display full-page photo (except in the context of a photostream)
user search_people & search_people/next photostream	Display the photos of an user on a grid Search for people and next result page Sequential display of user’s photos (or display of single photo in photostream context)
group	Display the profile of a group

Table 1. The most common URL classes in search sessions in our log sample.

	Overall	Trees	Chains
Sessions	1,071,954	-	-
Total	-	1,017,037	1,622,329
Avg. width	-	1.815	-
Avg. depth	-	1.575	3.129
Unique types	-	109,693	108,255
Trees/session	-	1.053	-
Chains/tree	-	-	1.513

Table 2. Events, search trees and search chains in the dataset.

representing its URL class and add it as a child of the node representing its referrer URL. In the resulting tree, any leaf represents a termination of a browsing branch; this does not necessarily mean the end of a search session, as other branches can occur later. Although in some cases a single session can contain more than one tree, in the remainder of this paper, for simplicity, we will use the terms *search sessions* and *search trees* interchangeably, i.e. by *search session* we refer to a subtree within a session corresponding to search activities. To create a more compact representation, we collapse non-branching sequences of nodes of the same class with the same URL parameters into a single node, ignoring differences in URL parameters in the following circumstances: the page number parameter for *search/next* nodes is ignored, and the photo id parameter is ignored for *photostream* nodes when two photos belong to the same photostream, and for *photo* nodes when two photos belong to the same set or group pool (indicating that the user is browsing within the same photostream, set, or pool). In this representation, we identify *search chains* (similar to *search trails* [6]) as the paths in a search tree that start at the root of the tree and end at a leaf.

3.2. Dataset Statistics

Table 2 summarises some statistics about the *search trees* and *search chains* in our corpus. The *search tree* representation gives over 100,000 unique search trees, 95% of which have a depth at most 3 and width less or equal to 4, while 95% of chains also have a length of 4 or less. For the remainder of the paper, we will refer to distinct search trees as *tree types*.

In Figure 1 we plot the cumulative distribution of repetitions for several URL Classes (i.e. how often a view of a certain page type is followed by a view of the same page




Fig. 1. CDF plot of repetitions for several URL classes.




Fig. 2. Most frequent search tree types.

type). We can see that, in search sessions *photo* views are followed by other *photo* views less than 15% of the time, whereas *photostream* and *user_photos* nodes appear much more often one after the other. This suggests that, when in a *photostream* view, a user is likely to browse photos in this photostream. When a user enters the *photo* view after a search, however, they are unlikely to view other photos within the same set or pool: this is likely to be an artifact of the Flickr user interface at the time of this study, which defaults to browsing a photostream, with options for browsing related sets or pools receiving less prominence in the interface. These patterns suggests that the user is browsing the results after the search, e.g., the user viewing a sequence of individual photos (*photostream*), or a sequence of thumbnails (*user*).

Figure 2 shows the 12 most common search trees, which give a succinct summary of the main user activities following search. The two most common search trees correspond to a search followed by no further action (*t1*), and a search followed by clicking on a single result (*t2*), which between them account for over 43% of the trees. *T1* trees may represent searches where the user is “satisfied” with the first page of thumbnails; alternatively, they could be “failed searches”. Search reformulation is quite common (*t3,t7,t12*), as are browsing photostreams via a single photo view (*t4*) and searching groups (*t5*). Branching is relatively infrequent, as only 1 out of the top 12 (and 10 out of the top-50) cannot be represented as chains.

Class	Prop.	Subclass	Examples	Prop.
Specific	35.7%	places	san francisco	14.1%
		events	burning man,	9.9%
		products	iphone 4, geektool	6.1%
		people	steve jobs, lady gaga	5.0%
		organisation	nypd, lafd, fdny	0.6%
General	47.2%	objects	trees, mountains, tiger	27.5%
		concepts	fashion, sports	19.7%
Photography	12.8%	photo equipment	fuji x100, nikon d7000	6.5%
		photo techniques	bokeh, depth of field	5.5%
		events	bc33, bc34	0.8%
Meta	4.3%	user/group names	-	3.4%
		other	api key	0.9%

Table 3. Taxonomy of annotated queries.

4. TAXONOMY OF IMAGE SEARCH

Query taxonomies for image search differ from those used for web search. Some work [12] has attempted to adapt Broder’s [4] taxonomy of intent for web search, while others have classified queries based on the type of objects and concepts the query refers to. Enser [13] distinguishes between unique (e.g. specific people) and non-unique queries, while Westman & Oittinen [8] follow the scheme of Shatford [14], and classify queries as queries for general objects, specific objects and abstract queries. We broadly follow those taxonomies, and distinguish between general and specific queries, and introduce 2 categories that are specific to photo sharing platforms:

- **Specific Queries**, which correspond to *unique search*, represent searches for a known-items, subcategorised by type: places, events, people, organisations and products.
- **General Queries**, which correspond to *non-unique search*, represent searches for items belonging to a certain category. As in Westman & Oittinen [8], we further sub-classify these as either being *objects* or *concepts*.
- **Photography Queries** are specific to photo sharing platforms, and include searches for photo equipment and techniques, and for photography related events¹.
- **Meta Queries** include searches for specific usernames and groups, and for site-specific Flickr features.

We manually annotated the 1000 most frequent queries from our corpus into this taxonomy. Queries that were ambiguous, or that do not clearly belong to this taxonomy, were labelled as “unknown”, leaving 974 queries with known categories. From Table 3, we can see that 35.7% of queries are *specific*, 47.2% are *general*, 12.6% are *photography* and 4.3% are *meta* queries. There are a less *general* queries than is reported by Jansen [15], although that work focused on all queries, not just the most popular queries. Searches for people

¹Mainly comprised of photography “bootcamps” – events in which photographers meet for training purposes.

Category	photo	search / reform.	search / next	user	group	end
Overall	30%	17.1%	5.4%	5.3%	2.2%	28.1%
Annotated	36.3%	9.3%	8.0%	4.6%	4.5%	27.8%
Specific	38.9%	7.3%	9.2%	5.8%	3.1%	28.3%
event	40.5%	5.0%	16.9%	2.5%	2.9%	28.2%
place	44.1%	7.7%	6.8%	4.2%	3.7%	27.2%
people	32.5%	7.0%	4.7%	10.7%	1.4%	31.4%
General	35.2%	12.6%	7.9%	3.4%	2.7%	30.3%
object	38%	11.5%	7.3%	2.0%	1.7%	32.4%
concept	32%	14.8%	8.8%	4.9%	3.6%	26.7%
Photographic	36.7%	5.1%	6.1%	0.7%	23.0%	11.4%

Table 4. Proportion of page views for common URL Classes on the first click after search, grouped query type (due to space, we only show query types with > 5% of the global proportion).

are much less important on photo sharing platforms than has previously been reported, both for general web image search [16, 15] and in a journalistic context [8]. It is also noteworthy *photography* accounts for 12.8% of popular searches, and that *meta* queries, which may not even be true image searches, account for over 4.3% of popular queries.

5. SEARCH BEHAVIOUR

In this section we take a closer look at user search behaviour on the Flickr platform, firstly focusing in particular on query-dependent and user-dependent behaviour, and then presenting query-type and user-type based URL class prediction models.

5.1. Query Type Based Variation

In Table 4, we show the distribution of URL classes for the first click after a search, for each query type. Comparing *all queries* with the subset of *annotated queries*, we can see that the annotated queries, which are the most popular queries, are reformulated much less frequently, and they lead to clicks on photos more often, suggesting that these most popular queries are “easier” than tail queries, which is to be expected.

Specific queries are followed by a photo click more often than general queries are: 39% of the time, compared to 35%. For some of the sub-categories, the difference is even greater, for example 44% of *place* queries lead directly to photo clicks, compared to only 32% for *concept* queries. On the other hand, *person* queries have much less photo clicks than other *specific* queries, and *object* queries have much more photo clicks than other *general* queries. Another important difference between *specific* and *general* queries is that specific queries are reformulated much less than general queries, with *event* queries reformulated the least (6.9%) and *concept* queries reformulated the most (14.8%). It is noteworthy that for *event* queries, different from other *specific* queries, users conduct a deeper exploration of the result list using *search/next*. This likely corresponds to browsing the set of results, suggesting that, unlike other *specific* queries, users conduct a deeper exploration of




Fig. 3. Reformulation, result click and search session end probabilities versus search tree depth.

results for *event* queries. Apart from the *event* queries, however, we see that the other *specific* query types lead to slightly less exploration of the results than *general* searches.

All of this suggests, as one would expect, that *general* queries are more exploratory in nature, appear to be more difficult to satisfy, and lead to slightly more complex post-search behaviour. Although these differences may appear obvious, we are not aware of any previous work that has documented and quantified this difference in the context of image search.

Figure 3 takes a closer look at query reformulation, showing the probability of transition from a query reformulation to another query reformulation at a given query depth (i.e. after 0,1,2, etc. previous reformulations). For all queries we see that the probability of reformulating a query increases with the reformulation depth, i.e. the more the query has been reformulated, the more likely that it will be reformulated again. Conversely, clicking on a search result is *less* likely as the query is reformulated more often. This behaviour is much less pronounced for specific queries, with the reformulation probability leveling off after 3-4 reformulations. We are not aware of any previous work documenting such behaviour, either for image search or for general web search.

Finally, when searching for photo equipment, the behaviour of users changes radically. In particular, such searches lead to clicks on group pages 23% of the time, have the least amount of query reformulation, and the session ends immediately after the initial search much less often. Based on this, it seems that when users of photo sharing platforms conduct photography-related searches they are often searching for groups dedicated to their search topic.

5.2. User Variation

To investigate user-based variations in behaviour, we take all users in our dataset who have conducted at least 10 search sessions. We use their distribution of search tree types to create a feature vector to represent each user, where each of the 50 most frequent trees is a feature, with all the other types grouped together in a 51st, “other”, feature. We use the X-Means clustering algorithm, an extension of K-Means that estimates the optimal number of clusters by maximising the Bayesian Information Criterion [17]. We obtain 4 clusters, the centroids of which are shown in Table 5. As the cluster centroids actually

Tree type	global	cl. 1	cl. 2	cl. 3	cl. 4
User Prop.	100%	7.77%	38.31%	35.55%	18.37%
t1	0.2593	0.0784	0.2479	0.0798	0.7068
t2	0.0773	0.017	0.1462	0.0366	0.0378
t3	0.0282	0.0165	0.0398	0.0126	0.0392
t4	0.014	0.0052	0.025	0.0079	0.0063
t5	0.0297	0.2056	0.0082	0.0266	0.006
oth	0.4127	0.2126	0.3285	0.7101	0.0976

Table 5. Global and per-cluster centroids. For clarity, we only show centroid values for only the top-5 trees, and "other" trees.

Category	Overall	Cl. 1	Cl. 2	Cl. 3	Cl. 4
<i>Specific</i>	35.6%	33.7%	35.2%	26.6%	54.4%
event	9.9%	2.5%	7.5%	5.5%	8.9%
places	14.1%	8.0%	12.2%	10.2%	12.4%
products	6.1%	3.0%	9.6%	6.4%	12.6%
people	5.5%	20.0%	5.7%	3.4%	16.2%
organisation	0.6%	0.2%	0.2%	1.1%	4.2%
<i>General</i>	47.2%	20.8%	52.4%	45.3%	37.6%
objects	27.5%	10.9%	30.5%	26.6%	21.3%
concepts	19.7%	9.9%	21.9%	18.7%	16.3%
<i>Photography</i>	12.8%	16.5%	10.1%	21.4%	6.2%
<i>Meta</i>	4.3%	29.1%	2.4%	6.8%	1.9%

Table 6. Proportion of query types for each user cluster.

represent the search tree distributions, unsurprisingly the four clusters thus identified differ significantly in their search behaviour as represented by these tree types. Table 6 shows the query type distributions for each of the clusters.

We identify the following groups:

- **Cluster 1** contains 8% of users, who often use ‘social’ search types, like *search_groups* or *search_people*. As shown in Table 6, the members of this cluster search mainly for persons, usernames and groups.
- **Cluster 2** is the most populous, and contains users whose behaviour is the most similar to global average, both in terms of search tree types and query types.
- **Cluster 3** contains 38% of users, whose search behaviour is in the long tail, with over 70% of search trees are outside the top-50. They conduct more photography-related searches, and are characterised by longer sessions, with more clicks per session.
- **Cluster 4** contains users who are more likely to issue *specific* queries, with a lot of *people* queries. They have the shortest sessions, with the fewest clicks per session. A large proportion of searches (70%) have no post-search clicks.

The differences in the distribution of query types in the user clusters also emphasise the modeling power of the search trees in capturing user behaviour: clustering user based on only the 50 most frequent search trees gives user clusters which show significant variation in the types of queries they issue.

Group	t[sec]	clicks	sec/click	photos	users	groups	pstreams
global	306	4.46	68.60	0.79	0.37	0.15	0.22
all	768	7.98	96.24	1.66	0.64	0.43	0.36
cl. 1	688	7.78	88.43	0.66	0.77	0.69	0.26
cl. 2	743	7.11	104.50	1.98	0.51	0.23	0.36
cl. 3	912	9.83	92.77	1.72	0.87	0.78	0.47
cl. 4	586	6.54	89.60	1.30	0.44	0.11	0.22

Table 7. Search session statistics for user clusters: time, and clicks on photos, user pages, group pages and photostreams.

Table 7 summarises the statistics of search sessions, grouped by user cluster. Unsurprisingly, since we only keep users who participated in at least 10 search sessions, the clustered users are more active, for all measures, than the overall population. The ‘social users’ in Cluster 1 click on fewer photos, and on more user group profiles. Cluster 3 users are significantly more active in terms of time spent and clicks performed, while Cluster 4 users are the least active.

5.3. Predicting User Search Behaviour

We now investigate whether these query-based and user-based differences in search behaviour have an impact on the task of web page type prediction, which can have important applications such as pre-fetching of web pages. We adapt the higher order Markov chain models studied by Chierichetti et al. [18]. A k th order Markov chain is a probabilistic chain in which the next state transition depends on the k previous states. In our setting, states are sequences of URL classes visited during a search session, and we predict the URL class of the next page that a user will visit. To apply these chain-based models, in this evaluation we do not consider search trees, but make the simplifying assumption that sessions can be represented as sequences. We train models search chain models of orders $\{1, \dots, 4\}$ for the entire corpus (global model), for the 3 main query types (general, specific and photo), and for the 4 user clusters, using 80% of the data for training and 20% for testing.

We evaluate the models based on the accuracy of the predictions, and the results are shown in Figure 4. We can see that, consistent with previous results for web pages [18], higher order models have more predictive power. More importantly, using different prediction models for different query types leads to a considerable increase in accuracy: training with all the data leads to a maximum accuracy of 0.36 (4th order models), but when using only data corresponding to specific queries the accuracy is increased to 0.50 for the specific query types. The average accuracy for the query specific models is of 0.46, again way above the global models. It is also worth noting that the global model is trained on a much larger dataset, since a relatively small subset of the queries are annotated.

Models trained for specific users also show an improvement, although not for all user types, with user-specific models for *Cluster 3* performing worse than the global model, and with precision generally lower for this class of users, which




Fig. 4. URL class prediction for query type and user type.

is explainable by the fact that these users behaviour is less typical, and therefore harder to predict. Again, the average of the user-specific models is still more predictive than the global model, although the difference is quite small.

Apart from potential applications of the models, these results, in particular the query-based results, emphasise that the differences in user search behaviour that we have been studying are significant enough that they have strong predictive power.

6. CONCLUSIONS AND FUTURE WORK

In this paper we use web server logs to study user search behaviour in a large photo sharing platform, namely Flickr. Our study uses logs of user behaviour during entire search sessions, as opposed to only the search and result click data that are available on standard search logs. Using a taxonomy of image search to describe the main categories of search performed on this platform, we note differences with previous results on general image search, and image search in journalism. We represent search sessions as trees, and show important query-based and user-based differences in search behaviour. We go on to show, for query-based differences in particular, that these differences can have a strong predictive value.

In future work, we plan to conduct a detailed comparison of image search behaviour on photo sharing platforms with general web image search, and to conduct a deeper analysis to better understand the intent behind image search queries.

7. ACKNOWLEDGEMENTS

This research is partially supported by European Community's Seventh Framework Programme FP7/2007-2013 under the ARCOMEM and SOCIAL SENSOR projects, by the Spanish Centre for the Development of Industrial Technology under the CENIT program, project CEN-20101037 (www.cenitsocialmedia.es), "Social Media.", and by Grant TIN2009-14560-C03-01 of the Ministry of Science and Innovation of Spain.

8. REFERENCES

[1] L. Chiarandini, M. Trevisiol, and A. Jaimes, "Discovering social photo navigation patterns," in *ICME*, 2012, pp. 31–36.

[2] P. Boldi, F. Bonchi, C. Castillo, D. Donato, A. Gionis, and S. Vigna, "The query-flow graph: model and applications," in *CIKM*, 2008.

[3] B. Xiang, D. Jiang, J. Pei, X. Sun, E. Chen, and H. Li, "Context-aware ranking in web search," in *SIGIR*, 2010.

[4] Andrei Broder, "A taxonomy of web search," *SIGIR Forum*, vol. 36, no. 2, pp. 3–10, 2002.

[5] Bernard J. Jansen, Danielle L. Booth, and Amanda Spink, "Determining the informational, navigational, and transactional intent of Web queries," *Information Processing and Management*, vol. 44, pp. 1251–1266, 2008.

[6] Ryen W. White and Steven M. Drucker, "Investigating behavioral variability in web search," in *WWW. 2007*, pp. 21–30, ACM.

[7] Flavio Chierichetti, Ravi Kumar, and Andrew Tomkins, "Stochastic models for tabbed browsing," in *WWW. 2010*, pp. 241–250, ACM.

[8] Stina Westman and Pirkko Oittinen, "Image retrieval by end-users and intermediaries in a journalistic work context," in *IiX. 2006*, pp. 102–110, ACM.

[9] Stina Westman, Antti Lustila, and Pirkko Oittinen, "Search strategies in multimodal image retrieval," in *IiX. 2008*, pp. 13–20, ACM.

[10] Bernard J. Jansen, Abby Goodrum, and Amanda Spink, "Searching for multimedia: analysis of audio, video and image web queries," *World Wide Web*, vol. 3, no. 4, pp. 249–254, 2000.

[11] Paul André, Edward Cutrell, Desney S. Tan, and Greg Smith, "Designing novel image search interfaces by understanding unique characteristics and usage," in *Proceedings of the 12th IFIP TC 13 International Conference on Human-Computer Interaction: Part II. 2009*, pp. 340–353, Springer-Verlag.

[12] Mathias Lux, Christoph Kofler, and Oge Marques, "A classification scheme for user intentions in image search," in *CHI '10 Extended Abstracts on Human Factors in Computing Systems*, New York, NY, USA, 2010, CHI EA '10, pp. 3913–3918, ACM.

[13] P. G. B. Enser, "Query analysis in a visual information retrieval context," *Journal of Document and Text Management*, vol. 1, no. 1, pp. 25–52, 1993.

[14] Sara Shatford, "Analyzing the subject of a picture: a theoretical approach," *Cataloging and Classification Quarterly*, vol. 6, no. 3, pp. 39–62, Spring 1986.

[15] Bernard J. Jansen, "Searching for digital images on the web," *Journal of Documentation*, vol. 64, no. 1, pp. 81–101, 2008.

[16] Dian Tjondronegoro, Amanda Spink, and Bernard J. Jansen, "A study and comparison of multimedia web searching: 1997-2006," *J. Am. Soc. Inf. Sci. Technol.*, vol. 60, no. 9, pp. 1756–1768, 2009.

[17] Dan Pelleg and Andrew W. Moore, "X-means: Extending k-means with efficient estimation of the number of clusters," in *ICML*, 2000, pp. 727–734.

[18] Flavio Chierichetti, Ravi Kumar, Prabhakar Raghavan, and Tamas Sarlos, "Are web users really markovian?," in *Proceedings of the 21st international conference on World Wide Web. 2012*, pp. 609–618, ACM.