

Search or browse? Casual information access to a cultural heritage collection

Robert Villa, Paul Clough, Mark Hall, Sophie Rutter
Information School
University of Sheffield
Sheffield, UK
S1 4DP

{r.villa, p.d.clough, m.mhall, sarutter1} @sheffield.ac.uk

ABSTRACT

Public access to cultural heritage collections is a challenging and ongoing research issue, not least due to the range of different reasons a user may want to access materials. For example, for a virtual museum website users may vary from professionals or experts, to interested members of the public visiting on a whim. In this paper, we are interested in the latter user: a user who visits a cultural heritage website without a clear goal or information need in mind. In the user study reported here, carried out within the context of the interactive task at CLEF (interactive CHiC), 20 participants explored a subset of Europeana with no explicit task provided using a custom-built interface that offered both search and browse functionalities. Results suggest that browsing is used considerably more by the majority of users when compared to text search (all participants used the category browser before carrying out a text search). This highlights the need for cultural heritage search interfaces to provide browsing functionality in addition to conventional text search if they wish to support casual search tasks.

General Terms

Design, Experimentation, Human Factors.

Keywords

Cultural heritage, virtual museums, information access.

1. INTRODUCTION

Providing public access to cultural heritage is an ongoing and challenging area of research. Previous work suggests that visitors to online cultural heritage collections (e.g. virtual museum visitors) are not necessarily motivated by an explicit task, and that interacting with cultural heritage collections is exploratory in nature [8, 9]. Recent work in the area of ‘casual search’ [10] has also investigated situations where users are driven by the pleasure of the search process itself, rather than an explicit information need.

The focus for this paper is how individuals explore a cultural heritage collection when given no task. The results may be used both to contrast with studies which have used explicit tasks, and to motivate changes to cultural heritage systems to better support a diverse range of user tasks.

The work reported here is based on initial results from the Interactive CHiC (Cultural Heritage in CLEF) track of CLEF¹ as run at Sheffield University. The interactive CHiC track is based on the CHiC Europeana data set as used in 2011 and 2012 [1]. An early prototype of an evaluation framework was used [2] which allowed the interactive experiment to be semi-automated. In this work, our focus is on how users explored the collection and in particular how search and browse were used in this exploration. We consider three research questions:

RQ1. How do participants initiate their exploration?

RQ2. Do participants use browse or search in their exploration of the collection?

RQ3. How do participants decide to search or browse, when given no explicit task?

With RQ1 we are particularly interested whether users start their exploration by browsing categories, or by search. RQ2 then considers how users access the collection over their whole session. For RQ3 we will present some initial qualitative data from our lab-based interactive study, where the aim is to identify reasons for the use of either the search or browse functions.

2. PREVIOUS WORK

A general review of museum informatics is provided in [3], although the more specific area of museum visitor studies, investigating why and how individuals visit museums, has a long history [4]. More recent work has focused on visitors to digital museums [5-7]. In [6] the information seeking behavior of cultural heritage experts was studied through interviews, finding that complex information gathering was required for the majority of search tasks. In contrast [7] studied virtual museum visitors, inspired by the work of [8] and [9] which suggest that museum visitors are exploratory in their information seeking. This work [7] found that search occurred far more often than browse behavior for three of the four tasks used in the study, the exception being an open and broad task where browsing occurred to a greater degree.

Museum visitors can, in some respects, be considered as examples of “casual leisure” searchers, as outlined in [10], where examples were found of “need-less” browsing (based on a diary study, and analysis of Tweets, both outside the domain of cultural heritage). Darby and Clough [11] investigated the information seeking

Presented at EuroHCIR2013. Copyright © 2013 for the individual papers by the papers’ authors. Copying permitted only for private and academic purposes. This volume is published and copyrighted by its editors.

¹ <http://www.promise-noe.eu/unlocking-culture>

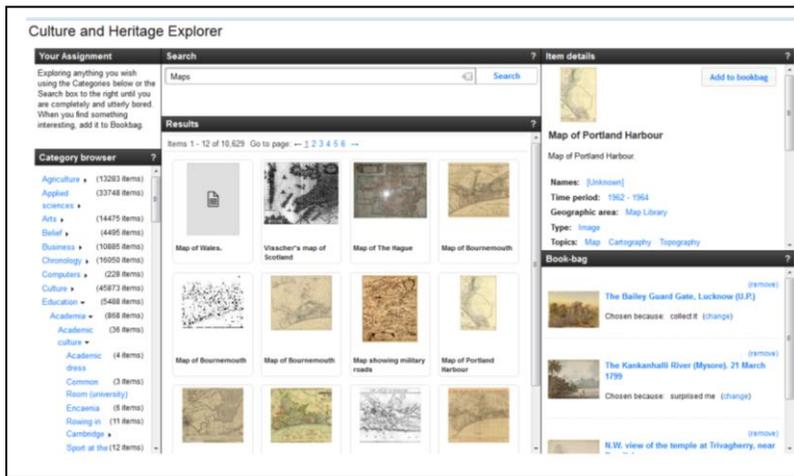


Figure 1: Screenshot of the Interactive CHiC interface

behavior of genealogists, with an emphasis on the behavior of amateurs and hobbyists, rather than professionals. In [12] a review of three digital libraries projects is carried out, from the point of view of Ingwersen and Järvelin's Information Seeking and Retrieval framework [13]. Similar to [10], it points out that information behavior by end users may be the "end in itself".

The study reported here uses a conventional lab-based protocol. However, unlike in previous work, such as [7], the participants were not given an explicit task: the underlying aim being to model a situation closer to that investigated in [10], where there is no explicit information need.

3. INTERACTIVE CHiC

A screenshot of the CHiC interactive system is shown in Figure 1. The interface is split into five main areas, clockwise from left to right: a category browser, search box, item display, bookbag, and search results. The search box operates in the conventional manner, allowing free text queries with search results being displayed as a grid below. When a result is clicked, it is displayed in the "item display" on the right. This information will typically include a small thumbnail, textual description, and the item's associated metadata. Metadata is clickable, e.g. if an item is listed as being owned by the British Library, clicking on the field will search for British Library objects. At the bottom of the item display is a "more like this", which displays the images of up to eight similar objects, which can be viewed three at a time.

On the left of the interface is the "category browser", which allows the user to browse the Europeana collection through a hierarchy of categories. This hierarchy is automatically generated, and is based on the work of [14]. The technique combines the Wikipedia category hierarchy with topics derived from Wikipedia articles into which items are mapped. When a category is clicked, the main results are updated to list the category contents. Small right arrows beside each non-leaf category allows the viewing of sub-categories. The user can therefore search and browse the collection in three main ways: using a text query, selecting a category, or selecting item metadata or "more like this".

On the bottom right of the interface is the bookbag, into which items can be placed. Book-bagged items are kept listed on the display, and can be removed and redisplayed as required. The

underlying search system is based on Apache Solr², which provides the text search, spelling checker, and the "more like this" suggestions (determined using Solr's standard more-like-this functionality). The data set used was the same as that used in interactive CHiC, a dump of the Europeana data set³.

4. EXPERIMENTAL SETUP

The search and browse interface was embedded into an IR evaluation system, which automatically administered pre- and post-questionnaires, and displayed the experimental system. All data reported here is from an in-lab study. This allowed a follow up interview to be carried out, during which each participant reviewed his or her search session. To enable this reviewing, Morae screen recording software was used to record the user's activity, and during the interview, an audio recording was made of user's comments.

An important aspect of the interactive CHiC experimental design was that no explicit task was provided to users. Instead instructions asked the user to explore freely as they wished, until they were bored. Users were informed after they had been active for 10 minutes, and could then continue for a further 5 minutes if they wished, at which point they would be asked to stop (these timings were carry out by hand, and were approximate). Once this was finished, the user's search session would be replayed to them, and an interview conducted to investigate the user's search process. Participants were paid 10 pounds for taking part.

In total 20 participants were recruited for the study, 11 male and 9 female. Eight participants were in the 18-25 year age band, nine in the 26-35 band; the other 3 between 36-45. The majority were students (13), with 5 employed, one unemployed, and one "other". 13 had completed a higher education degree, while six were currently studying an undergraduate degree.

5. RESULTS

5.1 Initiation of exploration

RQ1 asks how users initiate their exploration of the collection. To investigate this, we first looked at how users started their session, and in particular, their searching. For example, did they select a category or enter a query?

Over the whole data set four different actions were used by participants to initiate their session (Table 1, column 2). For the majority of users, the first action was to select one of the categories (15 out of the 20 users). It should be noted that the interface, on startup, showed a set of default results to all users. For three users, the first action was to display one of these default results, another user clicked the "next page" to view the next page of default results, while the final user's first action was to bookmark one of the default result items.

We also investigated the logs to find out each user's first search or browse action, which could be one of category select, text query, or metadata/more like this select. As shown in Table 1 (column 3), for all users this was a category select. In addition to counting the first actions, we also investigated how long each user spent before either clicking the interface, or starting a new

² <http://lucene.apache.org/solr/>

³ <http://www.europeana.eu/>

search/browse using the three previously listed methods. These results are shown in Table 2, along with the overall length of time of each session.

Table 1: Number of users whose first action/first search or browse action were as column one.

Action	#Users first action	#Users first search/browse action
Category select	15	20
Display item	3	-
Next search result page	1	-
Add to bookbag	1	-

Table 2: Time to first action, time to first search/browse action, and overall session time (all times in seconds)

	Min	1 st Qu.	Median	Mean	3 rd Qu.	Max
First action	7.00	19.00	25.00	30.50	38.75	90.00
First search/browse	7.00	22.75	38.00	57.50	81.75	204.0
Total time	129	631.8	783.5	787.8	918.0	1544

There was a considerable variance in the length of time users spent on the task. The median time taken by users was 783.5 seconds (just over 13 minutes), with an interquartile range of 286.2 seconds (approximately 4 minutes, 45 seconds). The minimum time was 129 seconds, and maximum 1544 seconds (over 25 minutes).

Most users spent some time at the start of their session before either clicking on an interface element (median time 25 seconds) or initiating a search (median 38 seconds).

5.2 Search vs. browse

RQ2 asks whether participants use search or browse. Figure 2 presents query and category counts across all users (i.e. counts of how often either text queries were executed or categories selected). Item select and the “more like this” functionality is not included here, due to the relative rarity of these events (across the whole data set this functionality was used only 15 times, by 7 different users).

A non-parametric Wilcoxon rank-sum test indicated that there was a significant difference between queries executed and categories selected ($W = 50.5, p \leq 0.001$). As can be seen from the boxplots, categories were selected far more than queries entered, the median number of queries executed being 2, compared to a median of 11 for category selects. All but three users selected more categories than executed queries, and 8 users did not enter a text query at all.

A similar situation exists when the time querying vs. browsing categories is estimated (Figure 3). Such times were estimated by starting a timer when a query or category was selected, and taking all activity between this point and the next query or category select as the user either “querying” or “browsing categories”. As might be expected, the trend is similar to that of Figure 2, with users spending more time browsing categories when compared to executing queries. All but five participants spent more time browsing using the categories than spent querying.

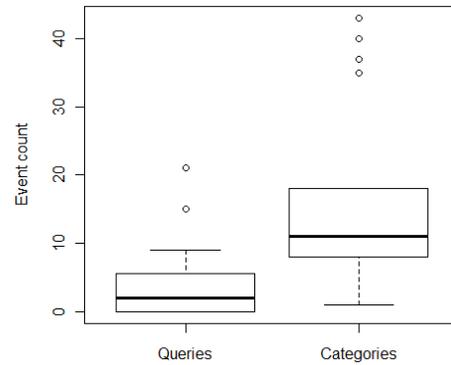


Figure 2: Comparison of query and category select counts

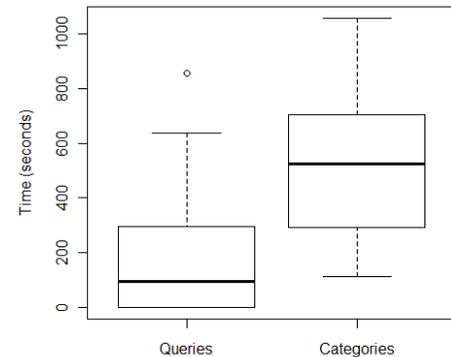


Figure 3: Estimated time querying vs. browsing by category

5.3 “How did you start?”

In addition to the quantitative data above, in the post-session interview two questions were asked of users: “how did you start?” and “Why did you choose to start with a [category/search query]?” It was intended to alter this latter question depending on how the user initiated their exploration. While some users started by examining the results, all users chose the category browser over the search box to initiate searches.

The responses to the first question “how did you start?” mentioned the category browser explicitly in 8 of the 12 answers. In most of these cases this was linked to exploring the interface. For example, participant P3 stated:

“I was drawn to the middle then decided to look around at the interface. I decided to look at categories first, picked politics”

Similarly, participant P10 stated:

“I just looked round to see what I could use to explore things. The category browser looked like the most likely candidates because it had descriptions of stuff.”

As well as being influenced by the interface, responses from some users suggest that prior interests also played a part. For example,:

“I just look at the layout of the website and then found that I had a category browser so I went to what I study actually, and I study languages and I try to find something interesting.” [P8]

“There is no particular task and so I started from browse to see which information is more interesting to me.” [P1]

The design of the interface, with a relatively small search box, appears to also have had an effect on the choices of at least two of

the user, indicated by responses to the second question. Participants P2 and P4 stated:

“Because I only saw that [category]. I didn’t see the search until a bit later on.” [P2]

“I didn’t really see this one at first [the search box] it was a bit obscure.” [P4]

For many users, however, the fact that the category browser allowed easy exploration appeared to be the key, with some users making connections to physical museums. For example:

“If I was going to a museum I would look at the categories [museum sections] that are of most interest to me: arts, old stuff and so this is why I was looking for Mona Lisa.” [P5]

The lack of an explicit task was mentioned by some, and search was explicitly commented on by two users. E.g., P7 stated *“When I wanted to find something specific I went to the search box.”*

6. DISCUSSION

RQ1 asks how participants initiate their exploration of the collection. From Table 1 it can be seen that all 20 participants started their exploration using the category browser, rather than a text search. Indeed, the first action for the majority of users (75%) was to select a category. Quantitative data from Section 5.3 backs this up, with 8 out of 12 of the participants for which text transcripts are available explicitly mentioning the category browser as a way of starting their exploration. Looking at Table 2, it can be seen that there is typically a short delay until participants started their browsing (median 38 seconds, interquartile range of 59). This delay is consistent with participant’s comments which suggested that many first spent some time orienting themselves to the interface before starting (e.g. P10 from Section 5.3).

Moving to RQ2 and RQ3, which asked whether participants have a preference for browse or search and why, it is clear from Figure 2 and Figure 3 that there is a general preference for browsing, e.g. from Figure 3 the median estimated time spent browsing using the categories was 524 seconds (IQR 399), compared to 77 seconds (IQR 394) for text queries. Looking at the participant comments, the lack of any explicit task would appear to have played a part in this preference (e.g. P1 and P5 quotes from Section 5.3). In addition to this the design of the interface, with a relatively small text search box at the top, appeared to also play a part, with some users pointing out that they did not see the search box until later in their session (e.g. P2 and P4).

7. CONCLUSIONS AND FUTURE WORK

The preliminary results reported here would suggest that providing browse functionality to cultural heritage collections is important for users arriving without a specific information need, as may be typical in casual search. For the majority of users, this preference for category browsing continues to hold for the session as a whole, with all but 5 users spending more time browsing than keyword searching. Initial analysis of quantitative interface data backs up the qualitative results, with more of the currently analysed user transcripts explicitly mentioning the category browser. The results presented here are preliminary. Future work will expand on the analysis presented here, both the qualitative and quantitative results. However, these initial results provide evidence of the importance of providing browse functionality to cultural heritage collections, and Europeana in particular.

Acknowledgements: This work was supported by the EU projects PROMISE (no. 258191) and PATHS (no. 270082).

8. REFERENCES

- [1] Gäde, M., Ferro, N., and Lestari Paramita, M. 2011. ChiC 2011 – Cultural Heritage in CLEF: From Use Cases to Evaluation in Practice for Multilingual Information Access to Cultural Heritage. In Petras, V., Former, P., and Clough, P., editors, *CLEF 2011 Labs and Workshops*, Italy.
- [2] Hall, M. and Toms E. 2013. Building a Common Framework for IIR Evaluation. Information Access Evaluation meets Multilinguality, Multimodality, and Visualization, 4th International Conference of the CLEF Initiative.
- [3] Marty, P. F., Rayward, W. B. and Twidale, M. B. 2003. Museum informatics. *Ann. Rev. Info. Sci. Tech.*, 37, 259–294.
- [4] Booth, B. 1998. Understanding the Information Needs of Visitors to Museums, In *Museum Management and Curatorship*, 17(2).
- [5] White, L., Gilliland-Swetland A., and Chandler R. 2004. We're Building It, Will They Use It? The MOAC II Evaluation Project. In *Museums and the Web (MW2004)*, <http://www.museumsandtheweb.com/mw2004/papers/g-swetland/g-swetland.html>
- [6] Amin, A., van Ossenbruggen, J., Hardman, L. and van Nispen, A. 2008. Understanding cultural heritage experts' information seeking needs. In *Proceedings of the 8th ACM/IEEE-CS joint conference on Digital libraries (JCDL '08)*. ACM, New York, NY, USA, 39-47.
- [7] Skov, M. and Ingwersen, P. 2008. Exploring information seeking behaviour in a digital museum context. In *Proceedings of the second international symposium on Information interaction in context (IiX '08)*, ACM, New York, NY, USA, 110-115.
- [8] Black, G. 2005. *The engaging museum*. London: Routledge.
- [9] Treinen, H. 1993. What does the visitor want from a museum? Mass media aspects of museology. In S. Bicknell and G. Farmelo (Eds.), *Museum visitor studies in the 90s*, London, Science Museum, 86-93.
- [10] Wilson, M. L. and Elswiler, D. 2010. Casual-leisure Searching: the Exploratory Search scenarios that break our current models. In: *4th International Workshop on Human-Computer Interaction and Information Retrieval*, Aug 22 2010, New Brunswick, NJ, 28-31.
- [11] Darby, P. and Clough, P. 2013 Investigating the information-seeking behaviour of genealogists and family historians. *Journal of Information Science* February, 39, 73-84.
- [12] Richard Butterworth and Veronica Davis Perkins. 2006. Using the information seeking and retrieval framework to analyse non-professional information use. In *Proceedings of the 1st international conference on Information interaction in context (IiX)*. ACM, New York, NY, USA, 162-168.
- [13] Ingwersen, P. and Järvelin, K. 2005. *The turn: integration of information seeking and retrieval in context*. Dordrecht, The Netherlands: Springer.
- [14] Fernando, S., Hall, M.M., Agirre, E., Soroa, A., Clough, P. & Stevenson, M. (2012) Comparing taxonomies for organizing collections of documents, *Proceedings of COLING 2012: Technical Papers*, 879-894.