

Presenting and visualizing results on an image retrieval user interface

Ayşe Göker¹, Richard Butterworth², Andy MacFarlane² and Simone Stumpf²

¹Ambisense Ltd

²Centre for HCI Design, School of Mathematics, Computer Science and Engineering, City, University of London

Abstract. The volume of images available online has increased significantly but the choices offered by image retrieval systems have not kept pace. We describe the design and evaluation of a ‘high density’ image search interface focusing on the results pages. We quantitatively and qualitatively compared image presentation on a high density interface with a traditional image search interface. Our results show that users had major problems with high-density interfaces for images due to information overload. We point to further work which could improve the user experience.

Keywords: *Image retrieval, user interface, results presentation, high density UIs*

1. INTRODUCTION

The volume of images available online has grown exponentially in recent years. Users face a non-trivial challenge to find an appropriate image from the enormous number available. There a number of freely available image search engines e.g. Google Image Search. These tools provide a limited set of images to the user, and do not give them flexibility to choose the number of images they wish to view. Rodden and Wood (2003) provide some evidence that presenting a large number of thumbnails is what users would generally prefer. We propose an alternative approach that shows the image searcher many images at once i.e. a ‘high density’ image retrieval interface. We hypothesize that this high density interface will be positively received by image searchers.

The paper is structured as follows: Section 2 describes related work in image retrieval search interfaces. Section 3 outlines the design for a high density image search interface. Section 4 then describes an experiment where a prototype high density interface was tested against a more traditional interface with the results of the test reported in section 5. We reflect on the design in section 6 and include recommendations for further work in section 7.

2. RELATED WORK IN IMAGE SEARCH INTERFACES

Many image retrieval and search interfaces are essentially based on a ranked list of search results. Search results are often shown as a list of reduced size images or "thumbnails", usually presented in a linear vertical fashion, ordering the images by relevance, time etc (Data et al, 2008). There are a number of methods for presenting thumbnail images including cropping algorithms (Suh et al, 2003) and representative photos (Huynh et al, 2005).

Plant and Schaefer (2011) provide a survey of the visualisation and browsing of image databases and provide an overview of layout techniques which can be used, based on features such as colour, texture, shape, and text, and possibly a visual metaphor. The visualisation techniques they identify include mapping based, clustering, Graph based and Virtual Reality based visualisations. Browsing techniques identified were horizontal and vertical browsing, panning, zooming, magnification and scaling. Examples of thumbnail display include the use of ‘PileBars’ (Brivio et al, 2012), clustering images (Nguyen and Worring, 2006), Treemaps (Wang et al, 2013) and oval bubbles (Smith et al. (2006). We use this literature to inform the design of our interface in section 3.

An important question in visualization of images is exactly how much should one display to the user at once, and what forms of image interaction might

provide better user experiences. Recent research into displaying a large number of images to a user on a screen (Plant and Schaefer, 2011; Wang et al 2010) has shown conflicting results as to the viability and usability of this approach. For example, it has been found that more than 30 images per screen 'confuses' the user whilst presenting a visual dictionary for English nouns which consist of tens of thousands of tiny 'tiles' that map on to millions of images worked well. The purpose of this work is to investigate the issue by comparing a high density user interface (Butterworth and Goker, 2011) with a google style control interface.

3. DESIGNING THE HIGH DENSITY IMAGE INTERFACES

Our design goal was to display image search results as thumbnails, and display a high number of results onto a single screen using an appropriate layout or organisation of thumbnails. The layout we used was a 'sunflower pattern' (figure1), based on Vogel (1979) and Rose (1999). This is an efficient method of packing the thumbnails onto the screen

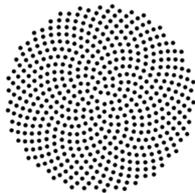


Figure 1. The sunflower spiral

The interface (figure 2) displayed cropped images as very small thumbnails (60 pixels) in the sunflower pattern which would then expand to a more normal thumbnail size (e.g. 150 pixels) on some gesture from the user (e.g. a 'mouse over' event). The search results showed the most relevant images in the centre and less relevant results further out in concentric circles around the centre point of the screen. This is one of the visual structures utilized by Torres et al (2003) however they used visual similarity on images rather than relevancy. Relevancy here refers to the relevance ordering created by the Xapian information retrieval system (Xapian, N.D.) through the keywords associated with the image.

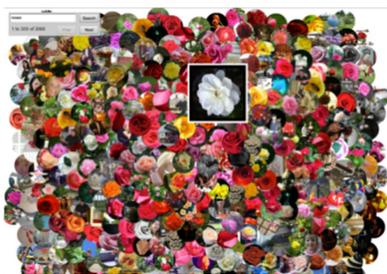


Figure 2. The high density image interface

4. EXPERIMENT DESIGN

We conducted a within-subject experiment comparing our interface design with a more traditional image search interface as a control, capturing both quantitative and qualitative data to investigate how well users perform on search tasks and perceived usability.

To test the two interfaces an online testing system was created, where users were presented with the control and experimental image retrieval interfaces, given some tasks to perform with both interfaces, and then invited to rate their experiences using multiple choice questions, Likert scale ratings, and more open ended questions.

Participants from the experiment were all students recruited from three postgraduate courses at City, University of London. The students came from a wide variety of backgrounds in the arts and humanities, and many had at least one year's work experience in the information professions. From this cohort, 58 students were invited to carry out the experiment. A total of 35 participated, with 18 both completing the experiment and answering the post-test questionnaire fully – an effective participation rate of 31%. There were sufficient numbers to carry out statistical significance tests in order to show any differences between user views on the interfaces presented.

Participants in the study were asked to find images using three different types of task: a targeted search, category search and an exploratory search (Cox et al, 2000). Each participant completed 6 tasks, 2 per condition, randomly assigned from each task type.

The 'Google like' control interface presented users with a linear list of 30 thumbnail images (150 pixels each) per page, with the user having the ability to page forwards and backwards through the list. The test set of images used for this experiment was a subset of the CoPhIR test set of images (Bolettieri, 2009). These images are taken from Flickr, and their title, tags and comments were indexed using the Xapian information retrieval system using BM25 (Robertson et al, 1994) without modification. Both interfaces used the same image dataset with the same information retrieval system; results were the same but presented in a different way.

During the experiment, participants completed the first three tasks with the control interface. When the user had found an image which they felt matched the task they were asked to click on it to confirm that it was the target image, and they then moved on to the next task. If they felt they could not find an image that matched the task there was also a link which took them to the next task, registering the fact that they could not find an appropriate image. After completing the first sets of tasks on the con-

control interface, they were then asked to perform three different randomly assigned tasks from the same list using the high density interface using the same process as the control set of tasks.

After the users had carried out the tasks using both interfaces, a post-experiment questionnaire was presented. There were three preliminary questions to gather background information, including frequency of using search engines such as Google (multiple choice); how often they use the search engines to look for images and photos (multiple choice); and prior knowledge of commercial image search engines (yes/no). Of the participants only two students had used such systems prior to the experiment.

Within the main body of questions, we asked both quantitative and qualitative questions. In questions 1 to 8 we focused on the perceived ease of completion for the search tasks, especially relating to the size of the thumbnails, perceived utility of the search interfaces in retrieving appropriate images, and the perceived user experience with the interfaces. We also asked participants which interfaces they preferred - the high density or control interface. All these questions used a Likert scale of 1 to 5, where 1 represented disagreement ('strongly disagree') and 5 represented agreement ('strongly agree'). We also provided a 'not applicable' option, which meant that the question was not scored. In questions 9 to 12 we asked participants to provide open feedback on the use of the interfaces to provide qualitative data.

To analyse the data, we used descriptive statistics and Wilcoxon tests to compare the ratings between the assessments of both interfaces. To analyse the qualitative data from open comments, we looked for common themes mentioned by test participants.

5. RESULTS

Figure 3 provides an overview of the quantitative results from the experiment. The results showed that users preferred using the control interface. Just over 80% disagreed that the higher density interface was easier to use and only around 20% found that using the high density interface was more useful to complete the set search tasks. Around 75% were negative about the ability of this interface to assist searching. Around 80% found the thumbnails too small to use on the high density interface. Comparing the interfaces, users clearly thought that using the control interface led to the best available image being found – just over 60% of users preferred the control interface compared with around 25% of users for the high density interface. The difference was statistically significant (Wilcoxon, $p < 0.05$). In terms of how happy the users were with the interfaces the results were more in favour

of the control interface with just over 80% expressing satisfaction compared with around 25% for the high density interface – the significance between the results was very significant (Wilcoxon, $p < 0.01$). Most users expressed the view that they would not prefer to use the high density interface over the control interface with just under 75% disagreeing with the statement expressed in Q8.

Question	Mode	Median	Max	Min
1. HD UI easier to use than control UI	1	1.5	4	1
2. Easier to complete the search tasks with HD UI	1	1.5	5	1
3. Thumbnail images on HD UI too small	5	5	5	2
4. Found the best available image using control UI	5	4	5	2
5. Found the best available image using HD UI	1	2.5	5	1
6. Liked using the control UI	5	5	5	3
7. Liked using HD UI	1	2	5	1
8. Prepared HD UI over the control UI	2	2	5	1

Figure 3. Quantitative statistics from questionnaire (1=strong disagree, 5=strongly agree)

From the comments gathered in questions 9 to 12, four clear problems with the high density interface emerge: mousing over images, thumbnail size, distance between images, and information overload. Users found that when they moused over images in the high density interface, other images were obscured – they found this frustrating: “it was very difficult to isolate images and look at them”; “I ended up enlarging thumbnails I did not intend to look at”. Users were not happy with the thumbnail size on the high density interface: “the images were too small to see without expanding them”; “the small photos on the high density interface hurt my eyes”; “The high density interface was TOO high”. These views match the quantitative data collected. Users also expressed dissatisfaction with the inability to go back and click on images previously viewed. They felt that images were not delineated clearly enough and that there was too little space between the thumbnails – leading to a feeling of the

high density interface being “too cluttered”: “I didn’t like the overlapping of pictures on the high density interface”; “space out the high density interface so that each individual image can be seen better”. This is consistent with the work by Rodden et al (1999) who found in their study that most users were negative about overlapping on highly packed thumbnail displays where some objects are partially or wholly obscured. The final problem relates to the ability of the user to absorb the information presented in the interface, which encompasses the points made above, but also had to do with the number of images presented. For the participants of this experiment information overload was something of a problem: “Seeing so many images yet not REALLY seeing all of them (unless I moused over them), makes me feel like I’m missing things in the way the normal interface does not.”; “less is more”.

It should be noted, however, that there was some contradictory evidence as to the usefulness of the high density interface compared with the control interface. Some participants who stated that they did not like the normal interface (‘disagree’ on the Likert scale) wrote comments that they actually liked using the high density interface.

6. A REFLECTION ON THE DESIGN

Here we reflect on the issues identified in the results for further investigation i.e. mousing over images, thumbnail size, distance between images, and information overload. The mousing over problem occurred because the thumbnails began to expand immediately when the mouse cursor passed over them. Therefore, if the user spotted an image that they maybe were interested in and tried to move the mouse cursor over it to expand it they were likely to expand the images close to the target image, on the way there – therefore obscuring the target image. This problem can be fixed by placing a small delay (0.3 seconds) on the expansion of the thumbnail when the mouse cursor was over the image for desktop screens. The problem of the distance between images can be solved by spreading the images out more (Rodden, 2002) but this clearly reduces the overall number of images available on the display. Spreading the images out so that they did not overlap there were not enough images on screen to establish the sunflower pattern. This issue can be addressed together with the problem of thumbnail size and information overload, by allowing the users to ‘zoom’ in and out on the interface e.g. by a slider which allows the users to expand the images from 30 pixels up to 150 pixels (in effect full size thumbnails). A minimum size of 30 pixels may seem a surprising choice as in the previous experiment users expressed dissatisfaction even with 60 pixels, but firstly at 30 pixels the sunflower pattern is visible and around six hundred

images could be displayed per screen giving users an ‘overview’ of their search results – allowing users the flexibility to pan in or focus on a specific part of the image results set.

7. SUMMARY AND CONCLUSION

The results of this experiment were largely negative; the subjects clearly did not prefer the high density interface. This is the main contribution of our work. However engagement was encouraging and there is some evidence from our work that users like high density interfaces. This is our secondary contribution. For the general user this may be less task based needs and more serendipitous exploration and/or affective needs. With a redesign of the display of the user interface and the interaction mechanism described in section 6, it is possible that more users would find the interface a useful way to engage with image sets. Currently our hypothesis that the high density interface is positively received by image searchers is not confirmed by the results here, but we do not rule out the possibility of the design working for specific user types.

There are two areas of concern which could be investigated further. Professional users who search for images and have the cognitive ability to process large numbers of images may find such a design useful. Such users may find that the high density interface is more suitable, providing the problems outlined in section 5 are tackled by the proposed redesign in section 6. The issue of information overload may not be such an issue with professional users, but some flexibility would be advisable e.g. the ability to pan and zoom images as suggested in section 6. We report this work in the same volume (Göker et al, 2017). A further area of fruitful research would be to examine cultural differences for example many Asian languages are pictorial in character (e.g. Chinese), which implies different design criteria for image interfaces for that user cohort (Santos, 2016). It would be worthwhile examining these differences and the impact on results presentation and user interaction in image interfaces.

8. Acknowledgements

This work is supported by the Piclet project, funded by the TSB reference TP14/DIC/6/1/BL209H.

9. References

- Bolettieri, P., Esuli, A., Falchi, F., Lucchese, C., Perego, R., Piccioli, T., Rabitti, F. (2009) CoPhIR: a Test Collection for Content-Based Image Retrieval. CoRR abs/0905.4627.
- Brivio, P., Tarini, M., Ponchio, F., Cignoni, P., & Scopigno, R. (2012). PileBars: Scalable Dynamic Thumbnail Bars. In VAST (pp. 49-56).
- Butterworth, R., Göker. (2011). The potential of a 'high density' image retrieval interface. Proceedings of HCIR 2011. Available on [https://sites.google.com/site/hcirworkshop/hcir-2011/posters – visited: 8th May 2017].
- Cox, I. J., Miller, M. L., Minka, T. P., Papathomas, T. V., and Yianilos, P. N. (2000) The Bayesian Image Retrieval System, PicHunter: Theory, Implementation, and Psychophysical Experiments. IEEE Transactions on Image Processing. 9, 20-37.
- Datta, R., Joshi, D., Li, J., and Wang, J.Z. (2008) Image Retrieval: Ideas, Influences, and Trends of the New Age. ACM Comp. Sur. 40, 5:1-60.
- Göker, A., Butterworth, R., MacFarlane, A., Ahmed, T and Stumpf, S. (2017). Presenting and visualizing image results for professional image searchers: A field evaluation. Proceedings of British HCI (in this volume).
- Huynh, D. F., Drucker, S. M., Baudisch, P., & Wong, C. (2005) Time quilt: scaling up zoomable photo browsers for large, unstructured photo collections. Proc CHI'05 extended abstracts on Human factors in computing systems, pp. 1937-1940. ACM Press, New York.
- Nguyen, G. P., and Worring, M. (2006) Interactive access to large image collections using similarity-based visualization. J. Visual Lang.Comp. 19, 203-224.
- Plant, W., and Schaefer, G. (2011) Visualisation and browsing of Image Databases. In Multimedia Analysis, Processing and Communications, Studies in Computational Intelligence 346, Lin, W. Tao, D., Kacprzyk, J., Izquierdo, E. and Wang, H. (eds). Springer-Verlag, pp 3-57.
- Robertson, S., Walker, S., Jones, S. Hancock-Beaulieu, M., and Gatford, M. (1994) Okapi at TREC-3. Proc. TREC-3. pp 109-126. NIST Gaithersburg.
- Rodden, K. (2002) Evaluating similarity-based visualisations as interfaces for image browsing. University of Cambridge, Computer Laboratory. Technical Report No 543. UCAM-CL-TR-543.
- Rodden, K., Basalaj, W., Sinclair, D., and Wood, K. (1999) Evaluating a Visualisation of Image Similarity as Tool for Image Browsing. IEEE Symposium on Information Visualisation, 36-43, IEEE Press, Hoboken, N.
- Rodden, K., and Wood, K.R., (2003) How do people manage their digital photographs, In: Proc. CHI 2003, pp. 409-416, ACM Press, New York.
- Rose, S.J. (1999) The Sunflower visual metaphor, A new paradigm for dimensional compression. In: IEEE Symposium on Information Visualisation, 128-131, IEEE Press, Hoboken, N.J.
- Santos, M. (2016). The Unberable Homogeneity of Design. Available on: <http://tinyurl.com/kpkqy6d> (Visited: 29th March 2017)
- Smith, G., Czerwinski, M., Meyers, B., Robbins, D., Robertson, G., and Tan, D. S. (2006) Facetmap: A scalable search and browse visualization. IEEE Trans. Vis.Comp. Graphics. 12, 797-803.
- Suh, B., Ling, H., Bederson, B.B. and Jacobs, D.W. (2003) Automatic thumbnail cropping and its effectiveness. Proc.16th annual ACM symposium on User interface software and technology, pp. 95-104. ACM Press, New York.
- Torres, R.S., Silva, C.G., Medeiros, C.B. and Rocha, H.V. (2003) Visual Structures for Image Browsing, In: Proc. CIKM 2004, pp 49-55. ACM Press, New York.
- Vogel, H (1979) A better way to construct the sunflower head. Math. Biosci. 44, 179–189.
- Wang, R., McKenna, S. J., Han, J., Ward, A. A., (2010) Visualizing Image Collections using High-Entropy Layout Distributions. IEEE Trans. Multimedia. 12, 803-813.
- Wang, C., Reese, J. P., Zhang, H., Tao, J., & Nemiroff, R. J. (2013). iMap: a stable layout for navigating large image collections with embedded search. In IS&T/SPIE Electronic Imaging (pp. 86540K-86540K). International Society for Optics and Photonics.
- Xapian Project. <https://xapian.org/> (Last accessed: 9th May 2017).