

MERGING TECHNOLOGY AND USERS: APPLYING IMAGE BROWSING TO THE FASHION INDUSTRY FOR DESIGN INSPIRATION

Annette A Ward¹, Stephen J McKenna¹, Anna Buruma², Peter Taylor², Junwei Han¹

¹School of Computing, University of Dundee, Dundee, Scotland, UK; ²Liberty Fabric, London, UK
fabric@computing.dundee.ac.uk

ABSTRACT

Archives are increasingly important to creative industries, especially the fashion sector, which often relies on historical collections for design inspiration, product-development, and market strategizing. As these collections are digitized and catalogued their accessibility becomes more complex and their usefulness restricted to those skilled in navigating through the metadata. Content-based image retrieval (CBIR) offers only a partial solution to this dilemma. Requirements of the fashion sector necessitate a tool to enhance the browsing experience for creative applications beyond retrieval. This paper highlights limitations of CBIR for applications in creative industries such as fashion. It briefly reviews relevant literature, highlights some research questions, and outlines the approach we are adopting to integrate the requirements of users with the appropriate technology. Examples from Fashion and Apparel Browsing for Inspirational Content (FABRIC), a collaborative 3-year £1.4M Technology Strategy Board project, sponsored by the Department for Innovation, Universities and Skills (DIUS), provide a real-world application.

1. INTRODUCTION

The relationship between art images and creative industries is indisputable. Throughout the textile and clothing sector, which is diverse (horizontally- and vertically-integrated), and dynamic (based on obsolescence and change), images are a constant source of creative inspiration to industry personnel. Images showing continuous and discontinuous pattern repetition (e.g., print, floral, plaid, stripe, and lace fabrics; tiles; mosaics; grillwork; wallpaper; embroidery; stained glass; carpets; and rugs) are excellent sources of inspiration. Unfortunately, their complex and abstract colouration and design make them hard to describe when cataloguing, and as a result, very difficult to access using text. Arcane database retrieval languages add to the problem. As a consequence, searching image databases may be limited to those proficient with the associated

metadata and/or host languages, thus precluding access by those less skilled with searching protocol.

Growing image collections require innovative technology to overcome monumental challenges of data management and to increase their commercial value. Museum, library, archive, and other art image collections, rich in history and culture, also provide invaluable stimuli for creative design applications. Increasingly, these collections are also revenue-producing for the content holders.

Evolving technologies may assist in increasing intrinsic value of these collections by offering an innovative solution to help make images easier to access. Further, integrating user *technical* requirements (i.e., specific attributes that may impact accuracy of retrieval) along with user preferences will help ensure algorithms that drive browsers are developed to meet the needs of users and the businesses and organizations they serve.

This paper highlights limitations of current content-based image retrieval for applications in creative industries such as fashion. It briefly reviews relevant literature, highlights some research questions, and outlines the approach we are adopting to integrate the requirements of users with the appropriate technological applications. Examples from a particular project (FABRIC: Fashion and Apparel Browsing for Inspirational Content) are used to illustrate this discussion. FABRIC aims to provide a browser that is appropriate for creatively-driven functions and more comparable to real-world visual experience while overcoming some of the challenges and deficiencies of retrieving, managing, storing, and presenting image data for desktop and mobile IT applications.

2. CONTENT-BASED IMAGE BROWSING

Humans are uniquely designed to process a tremendous amount of visual information. However, they may have difficulty precisely defining their visual world. For example, it is estimated the brain can distinguish about 10,000 nuances in colour yet individuals can name only a small fraction of colour terms (approximately 12) [1]. As a consequence, accessing images with intricate colour attributes may be challenging using metadata alone.

Content-based Image Retrieval (CBIR), the computer-derived technique for retrieving images based on elements such as colour, texture, and shape, may provide a partial solution. The term “content-based image retrieval” seems to have first appeared in literature in 1992 [2]. Emphasis in most early CBIR systems was on automated retrieval of relevant images based on notions of similarity to a query image [3]. CBIR can use features of a selected painting, photograph, print, drawing, or other object to find visually similar images and locate matches in a collection even if they do not share metadata with the original image.

However, similarity of retrieved images may provide visually confusing results due to inherent characteristics of the images that comprise the database. A collection filled with heterogeneous images will likely provide disappointing results for a similarity search. Such collections, which are not unlike many art image archives, do not lend themselves readily to structured browsing for design inspiration when using CBIR.

Eakins & Graham [4] highlighted the need for prioritizing research into new approaches to semantic image retrieval as well as improved methods for interface design and evaluation of system effectiveness. Similarly, Forsyth [5] also discussed user needs and strategies for evaluating image retrieval systems. Unfortunately, existing technology has failed to provide industry with more practical solutions to image retrieval chiefly because users are far removed from technology development. It should be well understood that computational methods that learn to organise images in an unsupervised way without human supervision or feedback will be deficient.

For our purposes, navigating the image set in a meaningful way is more important than finding semantically-close matches to a query (although query-by-example might be used to initiate starting points for browsing). Mapping image collections into a 2D or 3D space that preserves meaningful inter-image similarities will enable visual browsing directly on a display.

FABRIC is working towards developing an innovative visual user interface and interaction methodologies for accessible browsing, searching, retrieving, and managing digital art image collections in ways similar to humans flicking through a collection and focusing on what “catches their eye”. Humans are part of the system to be modeled (“human in the loop”). We believe taking advantage of human perceptual/cognitive abilities to guide interaction and provide data for offline learning and adaptation will resonate with users’ intuitions and semantic notions. By involving users throughout the process more salient data for modelling, machine learning, performance evaluation, and interface evaluation for iterative prototype development will be derived.

In contrast to commercial content-based image retrieval (CBIR) software, navigation and browsing through collections can be highly user-directed, leveraging human perception and design knowledge rather than relying on automated retrieval of ranked similar matches as in the ubiquitous query-by-example paradigm. For highly competitive and creative industries, navigating the image set in a meaningful way is more important than finding semantically close matches to a single query image.

CBIR has been applied to some digital art databases, but comprehensive user evaluations have been largely lacking. Ward et al. [6] tested CBIR on three major image sites: the British Broadcasting Corporation (ELVIS II), The Corporation of London Guildhall Library and Art Gallery (Collage), and The British Library (Images Online). The goal of the Ward et al. study was to assess the effectiveness of CBIR through a systematic evaluation while providing users with accessibility to images independent of applying indexing terminology beyond conducting the initial search. Results of that study indicated respondents were generally satisfied with the experience, even when similar matches were not returned. However, most respondents visited the sites with the purpose of finding a specific image rather than just browsing.

Users of CBIR at the BBC generally agreed the software had appeal beyond finding similar images. However, lack of practicality of using it as a retrieval system did surface. Only 15% “agreed” or “strongly agreed” the visual search was better than a word search. Yet, over 75% of the users “agreed” or “strongly agreed” that the search was fun to use and nearly 85% concurred that results were interesting despite users’ lack of satisfaction with success of CBIR to retrieve acceptable matches [7]. These results supported the idea that application of a similar technology may be useful for other applications in creative industries.

3. VISUALISING IMAGE COLLECTIONS

When exploring image collections, users’ intentions might be very vague. They expect the system to be able to provide a variety of cues and options to guide their navigation. Visualization should reflect the distribution of images, enabling users to gain a global view of the collection and also focus in quickly on parts of the database that are of interest. As users explore and browse the data, perhaps in search of inspiration, the interface should respond in real-time, forming a closed interaction loop between the user and the system.

The great majority of systems present images to users by displaying a grid of image thumbnails. One-dimensional visualization in which images are typically presented in a rank-ordered relevance list is most common. Relevance often needs to be based on high-level semantics but is more

often computed based on low-level visual similarity or other relationships among images such as temporal ordering.

Lee et al. [8] proposed a user-centred image browsing/navigation model. Others cluster images use visual features or creation time and thus organize the collection as a hierarchical tree or pyramid structure [9, 10, 11, 12]. Each node in the tree is associated with an image category. Users are able to move around the tree to roughly look through the whole database or zoom into an image category to view the details. Torres et al. [13] placed a query image at the centre of the screen and a set of similar images returned by the retrieval system surrounded the query image along a spiral or concentric ring with distances and sizes proportional to their similarities to the query. Although 1D visualization can be easy to implement, it fails to reflect the mutual relationships among images in the database. It is more suitable for the retrieval task that generally assumes the user has a mental image of what they hope to find. However, such an assumption is not always true, such as when the user's intention is ambiguous or the user relies on browsing the database to obtain inspiration. For these situations, the user might prefer to have a global view of a set of images in a way that reflects mutual relations among all images in the collection.

The main purpose of 2D visualization is to offer users a more global view that embeds similar images in 2D display. The yielded 2D distribution of images may augment the understanding of users of the database and suggest a subsequent action. Many image browsers use zooming and panning as the primary interaction mechanism [11, 12, 14, 15, 16]. The user can smoothly zoom into a picture or a group of pictures to see more detail, or zoom out and pan to get the overview, or dynamically scale images according to preference. Rubner et al. [17, 14] introduced 2D visualization to image retrieval/browsing. They used earth mover's distance to measure similarity between images' colour and texture feature distributions. Images were projected into a 2D Euclidean space using multidimensional scaling (MDS). Many related 2D visualization systems follow this idea but vary in the methods used for measuring similarity and performing dimensionality reduction. If the images are parameterized using appropriate continuous feature variables, they will lie on or near a low-dimensional manifold in the high-dimensional feature space. Inter-image similarity can then be computed based on (geodesic) distance on the manifold rather than Euclidean distances in the original feature space. However, finding appropriate mappings onto low-dimensional manifolds is challenging and a suitable fixed mapping is unlikely to exist. Rather, different mappings will be required depending on the user's needs. Furthermore, users might want to

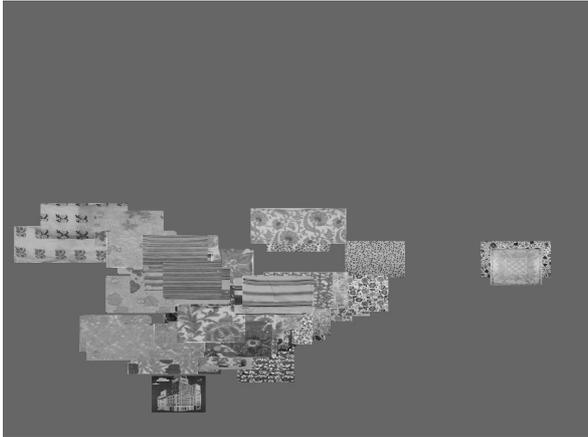
dynamically alter the mapping during an interaction, thus changing the meaning of the display axes.

ANVIL [18] adopted annotation information rather than low-level features to measure the similarity between two images. This can capture semantic information better and thus achieves superior performance. Walter et al. [15] embed images in a 2D hyperbolic space, using a hyperbolic-multidimensional-scaling algorithm. Images near the centre of the display have high resolutions while images located nearer the periphery appear as context with gradually lower resolutions. The user can move focus and indicate interest in an image; an interaction that seems very natural.

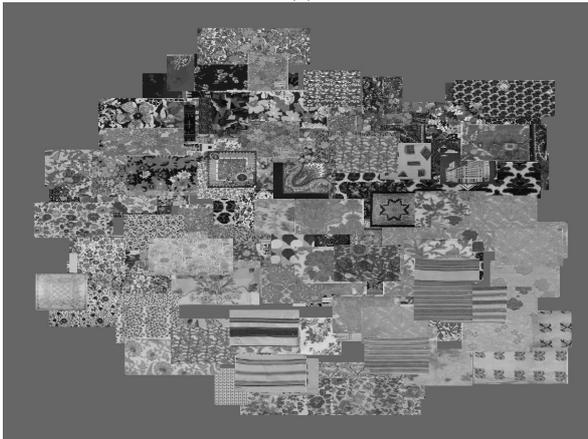
As pointed out by Pecenovc et al. [19], experiments have shown that users can prefer to view results in 2D maps rather than 1D lists since 2D visualization is able to provide them with context indications. Two-dimensional visualization is more appropriate to image browsing and navigation. Nevertheless, the state of the art in 2D visualization is still far from satisfactory. If the image distribution is modelled as a continuous manifold, the only points within the manifold that will be visualised are those at which images exist and these points will not be equally spaced in the manifold. Therefore, care must be taken to display images drawn from local regions on the manifold in such a way that images are not neglected (i.e., never displayed because others hide them from view and display real-estate is utilized efficiently. Most methods for automatically arranging images in 2D suffer from the problem of image overlap. Figure 1(a) shows an example in which part of an image collection is displayed for interaction using Isomap [20], to model the images using a 2D manifold. Such methods treat images as points and do not account for the size of each image. The result is severe overlapping and poor use of display space that can cause difficulties for users [21]. The images can be subsequently moved to optimise a cost function that penalizes overlap [22, 23]. However, a better solution is to directly trade off overlap and content similarity during manifold estimation [24]. Figure 1(b) shows the same image set as in Figure 1(a) but displayed using an approach in which image overlap is measured and used to adjust the spacing of images on the manifold.

The useful *intrinsic* dimensionality of an image distribution will usually be higher than two. It therefore becomes necessary to devise methods for organizing and navigating image collections in more than two dimensions. A 3D visualisation (in which images appear suspended in a 3D world) has been explored by several groups. For example, Nakazato et al. [25] used three dimensions reflecting colour, texture, and structure while Rubner [26] mapped images to 3D using multidimensional scaling. While it can be argued that special 3D display equipment is not necessary for 3D rendering to be useful, 3D visualizations can be computationally expensive. In any case, methods are required to enable users to browse image

collections with intrinsic dimensionality higher than three. For example, an 8D space could be browsed by constraining moves to be along 8 orthogonal directions on the manifold. Further possibilities include learning hierarchical representations and stochastic displays in which higher dimensions are sampled stochastically. The latter approach may prove particularly appropriate for browsing on small screens such as on mobile devices.



(a)



(b)

Fig. 1. Images arranged using manifold learning and colour histogram similarity. (a) A layout obtained using Isomap. (b) A layout obtained by trading off content dissimilarity and display overlap.

4. IMAGE MODELLING

The question of which features to use for image representation has been extensively studied in the CBIR literature [27]. Features usually describe colour [28, 29], texture [30], local geometry, and structure [31]. However, the use of such low-level image features results in the so-called semantic gap: better models that incorporate prior knowledge (context) are needed if similarity measures computed among images are to map

well onto user's semantic understanding of image content similarity. There has been widespread interest in the computer vision community in object and scene class recognition systems that deal with large image databases [31, 32]. However, our focus here is on collections of what may be described as non-representational, repeating image patterns which are characteristics of many fabrics.

Effective browsing models can be derived by considering the processes of producing printed textiles. First, the fabric designers develop a brief (a design framework) from which evolves designs that are created by hand and/or using computer-assisted design. Sometimes designs are resurrected from the digital or physical archive and reintroduced with changes in detail or colouring. Considerations are made regarding the design repeat and the numbers of colours that will be used. Frequent design repeat generally means that fewer meters are required for garment or home furnishings applications. Larger repeats require more fabric to construct a garment so that the pattern may be matched at seams (where the garment pieces are sewn together). More colours added to a run means more expensive production. Budget constraints usually limit the number of colours that are used in a design. Marketability of designs with very large repeats is carefully considered.

Designers may draft an original design by sketching their ideas on paper and doing the finish work by hand or they may use computer-aided design. Once the design is sent for the initial production run and the strike-off (the production sample) is checked for errors, the fabric is printed. Depending on the printing process used, different colours in the fabric are separated, either by hand tracing or computer reduction and separation. Figure 2 shows an attempt to automatically recover this colour separation from a digital image.

The method is based on an assumption that different dyes can be modelled as giving rise to Gaussian colour distributions in the image. Assigning pixels to regions using a Gaussian mixture model (with automatic model order selection) results in reasonable but somewhat noisy segmentation masks (Figure 2(d)). Incorporation of a spatial prior that encourages pixels to be labelled similarly to their neighbours results in better masks as shown in Figures 2(b) and (e). This is achieved using a Markov random field model with a graph-cuts energy-minimisation algorithm [33]. These segmentations enable images to be represented using features that reflect the design process that gave their rise. Users could also potentially isolate and modify aspects of the design (e.g., the shape of a particular colour mask) for subsequent organization and browsing of the image collection. Future work will perform analysis of the shapes of these colour masks, and could lead to analysis in terms of design elements, motifs, and even style. There are also obvious applications for computer-aided design (e.g., recolouring and modifying existing shapes).

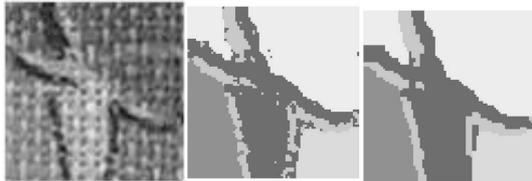


© Liberty Fabric

(a)



(b)



(c)

(d)

(e)

Fig. 2. Design is inferred from the colouration of a textile image: (a) The original image, (b) Colour segmentation using Markov random field model, (c) An extract from the image in (a), (d) Segmentation obtained using colour only (Gaussian mixture model), and (e) Segmentation obtained using colour and a spatial prior (Markov random field model).

Figure 3 shows automatic recovery of the design repeat from a digital archive image. The method hypothesises different repeats and compares them using an approximate Bayesian model comparison. The automatically-extracted repeat can be used to derive image features. Furthermore, knowledge of the repeat enables the image to be appropriately cropped for use as a thumbnail. This contrasts with the more usual practice of creating a thumbnail by reducing image



(a) © Liberty Fabric



(b) © Liberty Fabric

Fig. 3. Discovering repeats can be accomplished using Bayesian model comparison. (a) A detected ‘texton’. (b) An automatically extracted thumbnail image showing approximately four repeats.

resolution, an approach that often obscures important details of the design by rendering the image at an inappropriate scale. These ‘smart’ thumbnails are useful for image browsing because they enable more images to be effectively visualised simultaneously.

5. LEARNING FROM USER INTERACTION

Users may change their focus and their path while browsing through an image collection. When this happens, real-time communication between the user and the system is required. User interaction plays a key role in building such communication. An effective visualization mechanism provides users an easy way to communicate their preferences to the system. On the other hand, it should also enable the system to learn the users’ actions adaptively then rapidly respond to their preferences. Many efforts have been made to achieve this sort of interactive visualization including relevance feedback and layout reconfiguration.

Rui et al. [34] introduced the relevance feedback concept to the CBIR community. It allows users to evaluate search results and label relevant and irrelevant examples. Once the system receives the feedback information, it refines the results to more closely match the user's subjective interest using machine learning techniques [35, 36].

Layout reconfiguration approaches learn users' current focus by asking them to rearrange images [22, 37, 38]. Given a visualization reflecting similarity of image, the user moves images around to indicate preference. Based on the new layout, the system refines its similarity measure and updates the results.

Of additional importance is investigating incorporating semantics and learning when computing the display mappings (over and above semantics implicit in the feature (sub)sets). We will investigate the utility of making explicit to the user notions such as texture directionality, multi-colouredness, dominant spatial frequency, intricacy. For example, a user may wish to browse strong horizontal patterns or complex patterns such as Jacquard. Therefore, it is important to obtain manifolds that capture notions of similarity that are of use to the target users. This requires learning from browsing behaviour and images that are partially labelled or scored by users. For example, a subset of images could be manually scored on a 1D semantic space (e.g., perceived complexity). A mapping could then be learned between the manifold coordinates and this semantic space (using function approximation) resulting in a 2D space for display and browsing. Alternatively, spaces that combine semantic axes with axes determined in terms of representing remaining image variation could be used. Matusik et al. [39] have pursued a related approach for a completely different application in graphics in which users supplied binary labels for examples and support vector machines were used to compute a hyperplane separating the two classes. The direction orthogonal to this hyperplane then defined a "trend".

6. USER-CENTRED DEVELOPMENT

An iterative, user-centred approach for developing an innovative browser requires ongoing requirements gathering and user evaluations throughout the process. In our application with FABRIC, the browsing software will be tested with the people who stand to gain the most from the innovation, Liberty Fabric designers. Liberty Fabric has over 10,000 homogenous textiles images. Additionally, they have highly-experienced design and marketing personnel with over 50 years of collective design expertise who will add their knowledge of the design process to assist in developing the browser and

user interface. It is anticipated that their success in taking textile design to market means they are excellently placed to provide meaningful evaluation. Subsequent versions of the browser will include application to images that comprise the Victoria & Albert Museum (V&A) image collection (approximately 40,000 images). At that stage, a panel of experts from the UK textile community will assist in validating the browser with the V&A collection. When conducting a research project with commercial application, one of the challenges is to get sufficient user evaluation within a relatively small number of users ($n < 30$). Confidentiality is critical with the Liberty Fabric collection since it is commercial and restricted to Liberty staff.

Results from initial data gathering have been used to develop the preliminary browsing software. Multiple methods for evaluation add to the richness of the data collection and greatly assist the software developers in knowing and understanding the requirements of their users. Laboratory and on-line user evaluations will be conducted with Liberty's exclusive images and those of the V&A. V&A staff will also provide user feedback throughout the duration of prototype testing but with V&A images. User evaluation will be designed to collect demographic data of respondents, assess efficacy of the system as a source of inspiration and creative development, and test the effectiveness of the user interfaces. In addition to qualitative user evaluations, quantitative investigation of browsing behaviour will be performed based on analysis of manifold trajectories.

User interface development will involve investigating how best to lay out and update image subsets in 2D in response to real-time user interaction. For example, how can granularity be "zoomed" so that a user moves from browsing over quite dissimilar images to images that are very similar to one another? How can feedback and labelling be provided seamlessly and intuitively by the user for input to machine learning?

7. CONCLUSIONS

Content-based image browsing has great potential for creative industries. Although there is merit to conducting similarity search and retrieval, especially when one is looking for a similar image that was successful in the past (e.g., a best seller), similarity opens up an old design to new possibilities. With a little variation, last years' best seller may be reworked to become this year's successful design. There will always be times when searching by text is probably the best approach; for example, when one is looking for a specific item name.

As public access is widened with art image databases, newer ways of accessing images may be appropriate and open up new potential for creative acquisition of digital resources. Research regarding enhanced search capabilities is important to the future of information retrieval and the

development of searching software. Growing collections necessitate technology that expands browsing and search capability beyond traditional text indexing.

FABRIC moves beyond similarity matching to provide designers an array of images to spark and fuel design inspiration. Objectives include real-time browsing on desktop and mobile platforms, in-depth performance evaluation, and an intuitive user interface for visualisation and navigation. Visualization of images is based on manifold learning to make effective use of limited display size while visual content description applies a novel Bayesian method for regular texture analysis and Markov random field-based segmentation.

It is important that browsing software accommodates individual styles of “flicking” through images, allowing users to tailor their experience for when they have an idea of what they are looking for and when they do not. In-depth performance evaluation throughout a project ensures users’ perspective will be integral to the application of technology.

8. ACKNOWLEDGEMENTS

The authors are grateful to Ian Ricketts and Ruixuan Wang (University of Dundee) and James Stevenson (V&A Museum) for helpful discussions. Ruixuan Wang, and Wei Jia provided result images shown in this paper. This research was supported by the UK Technology Strategy Board grant “FABRIC: Fashion and Apparel Browsing for Inspirational Content” in collaboration with Liberty Fabric Ltd., System Simulation Ltd., and Calico Jack Ltd. The Technology Strategy Board is a business-led executive non-departmental public body, established by the government. Its mission is to promote and support research into, and development and exploitation of, technology and innovation for the benefit of UK business, in order to increase economic growth and improve the quality of life. It is sponsored by the Department for Innovation, Universities and Skills (DIUS). Please visit www.innovateuk.org for further information.

REFERENCES

- [1] Bornstein, M. H. “Perceptual Categories in Vision and Audition,” *Categorical Perception: The Groundwork of Cognition*, ed. S. Harnad. New York: Cambridge University Press, 1987.
- [2] Kato, T., “Database architecture for content-based image retrieval,” *Proceedings of the International Society of Optical Engineering, USA, 3846*, pp. 112-123, 1992.
- [3] Smeulders, A. W. M., Worring, M., Santini, S., Gupta, A., Jain, R., “Content-Based Image Retrieval at the End of the Early Years”, *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22(12), pp. 1349-1380, December 2000.
- [4] Eakins, J. P. and Graham, M. E., “Content-based image retrieval”, *Report to JISC Technology Applications Programme*, January 1999.
- [5] Forsyth, D. A. “Benchmarks for storage and retrieval in multimedia databases,” *Proceedings of SPIE - the International Society for Optical Engineering*, vol. 4676, pp. 240-7. 2001.
- [6] Ward, A. A., Graham, M. E., and Riley, K. J. “Final Report: Evaluation of Content-based Image Retrieval in Operational Settings: The Corporation of London Guildhall Library and Art Gallery, The British Library, and The British Broadcasting Corporation.” *Submitted to Resource: The Council for Museums, Archives and Libraries (CMAL/RE/103)*, UK, 2003.
- [7] Ward, A. A. and Murray, L. W. “Application of Content-Based Image Retrieval at the BBC and The London Guildhall Library and Art Gallery: A Comparison from the User's Perspective,” *Unpublished paper, poster session, CBMI Riga*, 21-22 June 2005.
- [8] Lee, J., Jeong, S. Y., Han, K. S., Chun, B. T., Bae, Y. J., “Image navigation: a massively interactive model for similarity retrieval of images,” *International Journal of Computer Vision*, vol. 56 (1/2), pp. 131-145, 2004.
- [9] Zhang, H. J. and Zhong, D., “A scheme for visual feature based image indexing,” *SPIE/IS&T Conference on Storage Retrieval Image Video Database*, pp. 36-46, 1995.
- [10] Chen, J. Y., Bouman, C. A., and Dalton, J. C., “Hierarchical browsing and search of large image databases,” *IEEE Trans. on Image Processing*, vol. 9(3): pp. 442-455, 2000.
- [11] Bederson, B. B., “Quantum treemaps and bubblemaps for a zoomable image browser,” *ACM Symposium on User Interface Software and Technology*, pp. 71-80, 2001.
- [12] Platt, J. C., Czerwinski, M., and Field, B. A., “PhotoTOC: automatic clustering for browsing personal photographs,” *Fourth IEEE Pacific Rim Conference on Multimedia*, pp. 15-18, 2003.
- [13] Torres, R. S., Silva, C. G., Medeiros, C. B., and Rocha, H. V., “Visual structures for image browsing,” *International Conference on Information and Knowledge Management*, pp. 49-55, 2003.
- [14] Rubner, Y., Tomasi, C., and Guibas, L. J., “A metric for distributions with applications to image databases,” *IEEE International Conference on Computer Vision*, pp. 59-66, 1998.
- [15] Walter, J. A., Webling, D., Essig, K., and Ritter, H., “Interactive hyperbolic image browsing-towards an integrated multimedia navigator,” *ACM Conference on*

- Knowledge Discovery and Data Mining*, pp. 111-118, 2006.
- [16] Huynh, D. F., Drucker, S. M., Baudisch, P., and Wong, C., "Time quilt: scaling up zoomable photo browsers for large, unstructured photo collections," *International Conference on Human Factors in Computing Systems*, pp. 1937-1940, 2005.
- [17] Rubner, Y., Guibas, L. J., and Tomasi, C., "The earth mover's distance, multi-dimensional scaling, and colorbased image retrieval," *The ARPA Image Understanding Workshop*, pp. 661-668, 1997.
- [18] Rose, T., Elworthy, D., Kotcheff, A., Clare, A., and Tsonis, P., "ANVIL: A system for the retrieval of captioned images using NLP techniques", *Third UK Conference on Image Retrieval*, 2000.
- [19] Pecenovic, Z., Do, M. N., Vetterli, M., and Pu, P., "Integrated browsing and searching of large image collections," *Fourth International Conference on Advances in Visual Information Systems*, pp.279-289, 2000.
- [20] Tenenbaum, J. B., de Silva, V., Langford, J. C., "A global geometric framework for nonlinear dimensionality reduction," *Science* 290, pp. 2319-2323, 2000.
- [21] Rodden, K., "Evaluating similarity-based visualisations as interfaces for image browsing," *Ph.D Thesis, University of Cambridge*, 2002.
- [22] Moghaddam, B., Tian, Q., Lesh, N., Shen, C., and Huang, T. S., "Visualization and user-modeling for browsing personal photo libraries," *International Journal of Computer Vision*, vol. 56(1/2) pp. 109-130, 2004.
- [23] Nguyen, G. P. and Worring, M., "Interactive access to large image collections using similarity-based visualization," *Journal of Visual Languages and Computing*, 2006.
- [24] Wang, R., Han, J. and McKenna, S. J., "Visualization of Image Collections based on Display Overlap and Content Similarity", *Submitted to International Conference Pattern Recognition*, 2008
- [25] Nakazato, M. and Huang, T. S., "3D MARS: Immersive virtual reality for content-based image retrieval," *IEEE International Conference on Multimedia and Expo*, pp. 22-25, 2001.
- [26] Rubner, Y., "Perceptual metrics for image database navigation," *Ph.D Thesis, Stanford University*, 1999.
- [27] Long, F., Zhang, H. J. and Feng, D. D., "Fundamentals of Content-Based Image Retrieval, Multimedia Information Retrieval and Management: Technological Fundamentals and Applications," D. Feng, W. C. Siu, and H J Zhang (Eds.), *Springer*, 2003.
- [28] Gevers, T., "Color in Image Search Engines," *Principles of Visual Information Retrieval*, M. S. Lew (Ed.), Springer-Verlag, 2001.
- [29] Schettini, R., Ciocca, G., and Zuffi, S., "A Survey of Methods for Color Image Indexing and Retrieval in Image Databases," *Color Imaging Science: Exploiting Digital Media*, R. Luo and L. MacDonald (Eds.), John Wiley & Sons, 2001.
- [30] Sebe, N. and Lew, M. S., "Texture Features for Content-based Retrieval," *Principles of Visual Information Retrieval*, M. S. Lew (Ed.), Springer-Verlag, 2001.
- [31] Carson, C., Belongie, S., Greenspan, H. and Malik, J., "Blobworld: Image Segmentation Using Expectation-Maximization and Its Application to Image Querying," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 24(8), pp. 1026-1038, August 2002.
- [32] Li, Y., "Object and concept recognition for content-based image retrieval," *Ph.D. Thesis, University of Washington*, 2005.
- [33] Boykov, Y., Veksler, O. and Zabih, R., "Fast approximate energy minimization via graph cuts," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 23(11), 1222-1239, 2001
- [34] Rui, Y., Huang, T. S., Ortega, M., and Mehrotra, S., "Relevance feedback: A power tool in interactive content based image retrieval," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 8(5): pp. 644-655, 1998.
- [35] Datta, R., Joshi, D., Li, J., and Wang, J. Z., "Image retrieval: Ideas, influences, and trends of the new age," *ACM Trans. on Computing Surveys*, vol. 29 pp. 1-65, 2007.
- [36] Liu, Y., Zhang, D., Lu, G., and Ma, W. Y., "A survey of content-based image retrieval with high-level semantics," *Pattern Recognition*, vol. 40 pp. 262-282, 2007.
- [37] Santini, S., Gupta, A., and Jain, R., "Emergent semantics through interaction in image databases," *IEEE Trans. on Knowledge and Data Engineering*, vol. 13(3) pp. 337-351, 2001.
- [38] Heesch, D. and Ruger, S., "Three interfaces for content-based access to image collections," *ACM International Conference on Image and Video Retrieval*, pp. 491-499, 2004.
- [39] Matusik, W., Pfister, H., Brand, M., McMillan, L., "A data-driven reflectance model," *ACM Transactions on Graphics*, vol. 22(3), pp. 759-769, 2003.