# Relevance Feedback in Content-Based Image Retrieval: Promising Directions

Bart Thomee Michael S. Lew
LIACS Media Lab
Leiden University
{bthomee, mlew}@liacs.nl

**Keywords**: relevance feedback, content-based image retrieval, image search

## ABSTRACT

In this paper we describe the current state of the art in relevance feedback as seen from a content-based image retrieval perspective. Recent technological advancements have spurred new research directions and have contributed new insights in the problem of visual information retrieval. We survey the main scientific problems and approaches in the last 4 years and give some recommendations for the future.

## 1 INTRODUCTION

In this digital age, millions of items are stored in huge databases and on the internet, and browsing through these collections for items of interest requires a new approach. This is especially true, considering that most of these items are not ours and we thus have no previous knowledge that might help us where to search to find what we're looking for. If we manually search these databases, we will eventually find the items of interest; while our capabilities in recognizing the content of media are unmatched, one of the greatest challenges is our lack of speed in searching libraries consisting of thousands or millions of items. By taking advantage of the processing power of modern day computers and having them perform the search instead of ourselves, the retrieval process for digital media is improved considerably.

The field of content-based image retrieval has been working for years on methods that allow a computer to have some understanding of images, through creating digital representations of image content and groupings of these representations based on their similarities. Yet, it is important to realize that a single search query is unlikely to return satisfactory results, because complexity and variations in visual images makes finding those that are sought for as hard as finding a needle in a haystack.

The concept of relevance feedback was initially developed to improve document retrieval [Rocchio1966]. In general, a document retrieval system presents a ranked set of documents relevant to the user's initial query and from thereon iteratively solicits the user for feedback on the quality of the documents and uses the feedback to compose an improved set of documents. Since keywords in a query can be matched one-on-one with text in a document, it is relatively easy for the retrieval system to determine which documents are relevant and which not.

This same technique can be applied to image retrieval – or multimedia retrieval in general – and hence it was also introduced into these fields. But, unlike with document retrieval, a user's query cannot be directly mapped onto items in the database. Relevance feedback is nonetheless still very well suited for this task as, with help of the user, the system can learn which (features, parts of) images the user is interested in; the feedback resolves many of the uncertainties that arise as the system tries to learn what the user is looking for.

In this paper, we focus mainly on the usage of relevance feedback in content-based image retrieval. First, we reflect back on the current state of the art, as has been published in the past few years, and then focus on the areas that look promising for the near future.

## 2 A CLOSER LOOK AT RELEVANCE FEEDBACK IN IMAGE RETRIEVAL

An image retrieval system using relevance feedback generally consists of four components, namely *user interface*, *offline learning*, *data storage* and *online learning*.

| | |
|---|---|
| **user query and feedback models** | a good user interface can help reduce the amount of interaction needed by making it easy for the user to submit a query and to give feedback |
| **semantic concept detection** | automatically detect the high-level concept of an image |
| **kernel tricks** | develop optimized or new kernel functions for clustering relevant images in a lower- or higher-dimensional space. |
| **weight adjustment and query point movement** | in the original feature space: employ weighted networks/lists or issue multiple simultaneous queries and merge the results |
| **using unlabeled images** | exploit unlabeled images by asking the user to label them |
| **long-term learning** | use information gathered during previous retrieval sessions to improve the retrieval results in future sessions |
| **sub-images and regions** | concentrate on matching query image segments with database image segments, instead of matching the image as a whole |
| **combining multiple knowledge sources** | visual features and annotations can complement each other to provide more accurate results |

Table 1. Current research directions

- The user interface allows the user to perform the initial query, mostly either by typing keywords, through sketching or by providing an example image. It additionally allows the user to indicate the relevance of returned images, often in the form of irrelevant-neutral-relevant.
- In the offline learning phase, image descriptors are extracted from all images in the database. Since users think in terms of high-level semantic concepts and not in low-level image features as available to the system, it is very important to select the most useful image descriptors, to help narrow this so-called *semantic gap*.
- The data storage component acts as a virtual file system and is responsible for storing and loading of necessary data. It is essential to use suitable multi-dimensional indexing techniques for an efficient search in high-dimensional feature space, especially considering that performance quickly suffers with an increase in dimensionality [Böhm2001].
- The online learning component performs the actual relevance feedback; through interaction with the user it discovers which image features the user deems relevant and which not and consequently ranks the images according to how well they conform to the relevant features.

As put forward in [Zhou2003], relevance feedback using both positive and negative training examples exhibits amongst others the following characteristics: the *small sample problem*, which is the issue that the user will only label a tiny amount of images while the amount of dimensions in feature space is huge, making it very challenging to discover the user's interest; the necessity of *asymmetric treatment*, as positive images are assumed to belong to a particular class of images, while negative images can belong to any class; and that users require *real-time interaction* and thus the retrieval system should be sufficiently fast.

## 3 RECENT WORK

During the last 4 years there has been tremendous interest in relevance feedback approaches toward content-based image retrieval. The last focused survey on relevance feedback was in 2003 [Zhou2003], so this work serves to give an overview of the recent advances in the area; Table 1 lists a short summary of the current research directions that we discuss in this section. Other more general introductions to the area of multimedia information retrieval can be found in [Smeulders2000] and [Lew2006].

*User query and feedback models from a user interface perspective*

Requiring a minimal amount of effort of a user is key in relevance feedback; if too much involvement is demanded, the user will be reluctant to use the system. Usually a user is searching for a particular category of items (*category search*) but the user can also be looking for something very specific (*target search*). A user searching for the latter will be harder to satisfy and the search process can then take many feedback rounds. A good user interface can help reduce the amount of interaction needed by making it easy for the user to submit a query and to give feedback.

The method used for submitting a query should depend on the target audience. For example, real estate agents might want to be able to submit a picture of a particular house in order to search for similarly looking ones, while graphic designers could prefer to sketch an outline of a particular piece of stock imagery they need. Commonly, users are able to begin a search by entering keywords, by sketching, by providing an example image, by selecting images from a random collection, or through a combination of the aforementioned.

After the initial search a ranked set of images needs to be shown to the user. The presentation choices are essentially limitless: the amount of images to display, the size of the images, the ability to scroll or page-flip through the list of images, how to indicate (ir)relevance. Depending on the learning algorithm, there are variations in how the user can indicate preference. Current work sees the following variations: positive feedback only [Jin2005], positive-negative feedback [Zhang2005] [Hoi2006], positive-neutral-negative feedback [Gosselin2004] [Zhou2006] and feedback with a degree of (ir)relevance [Haas2004]. [Nguyen2005] takes on an uncommon approach in showing the images to the user, which they call *similarity-based visualization*, where images are not displayed in a grid, but are displayed in a partially overlapping pile with related images placed near each other, see Figure 1 for an impression; one user action suffices to select several related images, greatly reducing the effort the user needs to invest.

Figure 1. Similarity-based visualization

*Semantic concept detection*

During the offline learning phase, many different types of image descriptors are extracted from the images. Low-level features such as color, texture and shape are commonly selected. However, humans instantly grasp the set of concepts involved when seeing an image and do not necessarily look at the low-level features as such. Even more so, low-level features are not really powerful enough to completely represent the complete semantic content of an image [Zhang2003].

Ideally, all images in the database are preprocessed by humans to yield a set of concepts for each image, but this is of course very labor intensive for large databases. Moreover, humans can have different subjective interpretations of the same concept. To alleviate these issues, current work is directed towards automatically detecting semantic concepts. These concepts are represented by keywords, reducing the content-based similarity search to finding images whose keywords match those of the images selected as relevant by the user.

[Fan2004] applies automatic image segmentation to decompose images into salient objects; the goal is to not only identify particular image objects (e.g. sky, sand, water) in the images, but also to combine all or a subset of those objects into higher-level concepts (e.g. beach), see Figure 2 for an example. This work is continued in [Gao2005] where a hierarchical feature subset algorithm is developed which selects suitable features for different semantic image concepts. In [Carneiro2005] the authors assume a database centric view of the retrieval problem, which is defined as finding the database class to which the query belongs; in this view a direct mapping is established between semantic classes and images in the database.


Figure 2. Detecting the concept ‘beach’

*Kernel tricks*

Similar images are not always nicely clustered together in feature space, in fact they are usually scattered throughout the space. A linear classifier is thus not very useful in this situation. Kernel methods are used to map the original image data into a higher-dimensional non-linear space, and with a good kernel function similar images can be mapped on the same side of a hyperplane in the new space: linear classification in the new space is equivalent to non-linear classification in the original space. Recent research has mostly focused on the application of kernel functions in Support Vector Machines (SVMs), often using enhanced versions of Tong’s original $SVM_{active}$ [Tong2001].

In [Xie2004] a unique kernel is generated for each user by using the probabilistic information that is contained in the user feedback; this approach is more flexible than using a fixed kernel as it can choose the optimal kernel parameters for each query. Kernel training is also applied in [Wang2005], where their scheme can quickly test image samples by correcting the results of the user's selections during the previous feedback round; this quickness is especially useful in large image databases, where traditional SVMs show low performance. The authors of [Ferecatu2004] argue that by using kernel functions that are insensitive to the scale of data the performance of SVMs can remain high; they also suggest that the learner should strive towards reducing redundancy between the images that it presents to the user, so as to better explore the frontier between the class of interest and the other images. [Goh2004] also introduces diversity into the presented images; they employ an adaptive image selection scheme which improves the ease of learning a query concept based on the complexity of the presumed concept. A new selection criterion is developed in [He2004a] that for any specific learning task can maximally shrink the version space (the set of hyperplanes that separate the labeled examples) and requires only labeling of a single image at a time by the user.

Kernel usage is not solely restricted to SVMs. In [Zhou2002] biased discriminant analysis is proposed to address the small sample problem and the asymmetric clustering of positive and negative images.

*Weight adjustment and query point movement*

Many approaches still operate on the original feature space. Recent work here is mainly concerned with maintaining a weighted network or list and using the feedback to adjust the weights along the edges. In [Xin2004] a Bayesian network is used where at each stage of iteration the belief values of all relevant image nodes are re-evaluated, resulting in a probabilistic ranking of the relevant images. [Zhou2006a] maintains a linked list for each image pointing at the most similar images and performs a multipoint query after which the similarity distances along the links are updated based on the user feedback; the images having small similarity distances are likely to be relevant. Multiple simultaneous queries are also issued in the approach by [Jin2005]. The results are merged in two ways: those containing images in the same visual cluster are combined to filter irrelevant images and those containing images in different clusters are combined to introduce visual variety.

*Using unlabeled images*

In typical supervised learning the machine learning process only involves labeled data to train the learner. The user helps to train the system by labeling images and together with the original query the learner is retrained, in order to produce a new rank that puts more relevant images at a higher position than in previous attempts. However, a user generally cannot be troubled with labeling a lot of images.

As image databases usually contain unlabeled images, these can be exploited to help supervised learning by asking the user to label them. The goal is naturally to minimize asking help from the user. It isn't always optimal to ask the user for feedback on the *most-positive* images; a higher information gain can be obtained by asking for feedback on the *most-informative* images, which are those images that the learning algorithm is most uncertain about.

[Lin2005] works from the assumption that images spread as a manifold and by learning its structure a similarity measure can be easily computed. [He2004b] inspects the relationships of all the data points in the feature space using a manifold ranking algorithm and constructs a weighted graph that contains all images; the ranking scores of labeled examples are iteratively propagated to nearby labeled and unlabeled images. Two generative models, one for relevant images and one for irrelevant images, is constructed in [He2005]; their goal is to estimate the likelihood functions that given an image class a certain data point is generated: from these functions the probability of the relevancy of each data point to the query concept is computed, and the unlabeled images are ranked accordingly.

[Zhou2006b] employs two learners which are trained from the labeled data provided by the user after every round of relevance feedback and consequently each learner then independently labels some unlabeled images and trains the other learner with it; their image classifications are merged and positive images having a high-confidence degree are presented to the user.

*Long-term learning*

Because each user is unique and will have a different view on relationships between items and thus of their relevancy, each relevance feedback session will follow a different path. However, since people do share certain characteristics, it is likely that if user A likes image I, user B probably will also like image I. Long-term learning is an approach that utilizes this kind of information gathered from previous retrieval sessions to improve the retrieval results in future sessions. [Hoi2004] constructs a relevance matrix to describe the relevance relationships between images and proposes a soft label support vector machine that can deal with training data having varying degrees of confidence. [Shah2004] creates a semantic space of the image database by clustering images into memberships through the use of transaction matrices that contain all user feedback scores; new queries are compared to previous search results and the images that received high scores in similar queries are recommended to the user. In [Yoshizawa2004] semantic groups are formed through inferring relationships between images if they have been jointly selected during previous search sessions.

*Sub-images and regions*

Often, the user is only interested in a specific object, which can mostly appear anywhere in an image. Most image retrieval techniques focus on images as a whole and consequently are not able to accurately capture the user's interest. Even most region-based image retrieval methods, which segment the image into (semantic) regions, generally attempt to globally match the regions found in the query images the user selected to the other database images. Current work concentrates on locally matching query image segments with database image segments, working from the assumption that, despite the fact that images contain many segments, only a few segments of the query images will match those of the database images. Additionally, the important segments should appear more often in positive images than less important segments and the important segments should appear less often in negative images than unimportant segments.

There are various ways to segment an image, e.g. segmenting an image into similarly looking (based on color values) regions, segmenting it into meaningful semantic regions or by simply by chopping the image up into small adjacent and possibly overlapping tiles. In [Luo2004] images are split into adjacent tiles and modeled in a hierarchical tree – the root tile is the whole image, the bottom layer consists of the smallest tiles – and in order to match a query image to the database images, their tree structures are slid over each other and distances between tiles are assessed; each tile will incur a certain penalty during each iteration, with more similar tiles having lower penalties and thus more similar images having lower total penalties.

Similar to semantic concept detection, semantic regions are formed in [Zhang2005]; images are segmented using a wavelet-based technique and the resulting regions are clustered by similarity through applying a genetic algorithm; the assumption is that a user is generally interested only in one particular semantic region present in the query image(s) and after the initial query the primary region of interest is distilled and used as the main hypothesis during further iterative learning. [Rhamani2005] uses identical techniques, but the authors pursue a set of hypotheses instead of only one and combine them when ranking images.

*Combining multiple knowledge sources*

To enhance the retrieval process, it is worthwhile to combine image features with other sources of knowledge. Low-level visual features alone do not completely convey the content of an image, and text annotations by themselves are generally directly related to the high-level semantics of an images. Together, visual features and annotations can complement each other to provide more accurate results.

[Zhang2003], [Yang2005] and [Ferecatu2005] incorporate WordNet [Fellbaum1998] to link annotations of images with generalized concepts. Finding similar images to the query image is facilitated by using WordNet to find semantically similar images. [Haas2004] takes this a step further by using WordNet as a bridge between various annotated databases to improve the accuracy of searching in an unannotated database; a genetic algorithm is used to evolve the image sets between iterations.

## 4 FUTURE DIRECTIONS AND MAJOR CHALLENGES

Based on the recent literature, the major challenges are the following: (i) *bridging the semantic gap* through improved concept detection techniques, (ii) *overcoming the curse of dimensionality* by selecting only the most optimal features, (iii) *reducing the amount of iterations* before the user is satisfied.

The field of sub-image retrieval shows potential for the future as indeed most user queries are directed towards a central object or concept. For instance, a user looking for a red sports car will not start caring about the other details in the image until a lot of images containing a red sports car are actually found. On the other hand, if the user is interested in a general concept, e.g. a beach, then a localized approach is less useful as this concept occupies the whole image. An approach assuming a global view and at the same time a local view might bring the benefits of both worlds together.

The common beliefs and interpretations of people are exploited by long-term learning algorithms, allowing a new query to converge quicker because steps that previous users had to go through can now be skipped if the search paths show similarities. Reducing the amount of effort required by users is an absolute necessity and this technique looks very promising to reaching that goal.

One of the fundamental problems with relevance feedback approaches from the past is that systems have been limited to the relatively small space of database images. Even if a database contains a million images, it would be insignificant compared to the typical high dimensional feature space which are often in the range of 1000 dimensions. Thus, when we want to illuminate whether an area of feature space is relevant to the user there frequently will not be a database image to present to the user for clarification. When we are learning new visual concepts, our own imagination often constructs synthetic examples based on real objects or scenes to help understand or clarify the primary features which are associated with the concept. By giving the learning algorithm this ability, it can synthesize images that will help it to clarify uncertain regions in the relevance space. This synthesized example would presumably target a particular feature which could be important to the query; it would be able to target this feature more closely than an existing image from the database and aid the learner in finding the relevant images quicker.

Another fundamental problem is the lack of high quality test databases. In the research literature the Corel dataset is most frequently used, however, it has two rather challenging faults in that it is both copyrighted and it is no longer published. Other image sets such as Benchathlon typically have minor or negligible annotation nor any rigorous testing procedure so that results can not be compared scientifically between different algorithms. This is to say that a good test set should contain both the media and a credible method of determining the accuracy of the system. Large piles of images are simply not useful to the research community. We need to have high quality international image test sets, which can be used for both measuring a system's performance and be used to improve a system by illuminating the weaknesses.

## REFERENCES

[Böhm2001] C. Böhm, S. Berchtold, "Searching in high-dimensional spaces: index structures for improving the performance of multimedia databases", in *ACM Computing Surveys v.33 n.3*, pp. 322-373, 2001.

[Carneiro2005] G. Carneiro, N. Vasconcelos, "A database centric view of semantic image annotation and retrieval", in *Proc. of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 559-566, 2005.

[Fan2004] J. Fan, Y. Gao, H. Luo, G. Xu, "Automatic image annotation by using concept-sensitive salient objects for image content representation", in *Proc. of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 361-368, 2004.

[Fellbaum1998] C. Fellbaum, G. Miller, editors. "WordNet: an electronic lexical database", *The MIT Press*, 1998.

[Ferecatu2004] M. Ferecatu, M. Crucianu, N. Boujemaa, "Retrieval of difficult image classes using SVM-based relevance feedback", in *Proc. of the 6th ACM SIGMM international workshop on Multimedia information retrieval*, pp. 23-30, 2004.

[Ferecatu2005] M. Ferecatu, N. Boujemaa, M. Crucianu, "Hybrid visual and conceptual image representation within active relevance feedback context", in *Proc. of the 7th ACM SIGMM international workshop on Multimedia information retrieval*, pp. 209-216, 2005.

[Gao2005] Y. Gao, J. Fan, "Semantic image classification with hierarchical feature subset selection", in *Proc. of the 7th ACM SIGMM international workshop on Multimedia information retrieval*, pp. 135-142, 2005.

[Goh2004] K.-S. Goh, E.Y. Chang, W.-C. Lai, "Multimodal concept-dependent active learning for image retrieval", in *Proc. of the 12th annual ACM international conference on Multimedia*, pp. 564-571, 2004.

[Gosselin2004] P.H. Gosselin, M. Cord, "A comparison of active classification methods for content-based image retrieval", in *Proc. of the 1st international workshop on Computer vision meets databases*, pp. 51-58, 2004.

[Haas2004] M. Haas, J. Rijsdam, B. Thomee, M.S. Lew, "Relevance feedback: perceptual learning and retrieval in bio-computing, photos, and video", in *Proc of the 6th ACM SIGMM international workshop on Multimedia information retrieval*, pp. 151-156, 2004.

[He2004a] J. He, M. Li, H.-J. Z. H. Tong, C. Zhang, "Mean Version Space: a new active learning method for content-based image retrieval", in *Proc. of the 6th ACM SIGMM international workshop on Multimedia information retrieval*, pp. 15-22, 2004.

[He2004b] J. He, M. Li, H-J. Zhang, H. Tong, C. Zhang, "Manifold-ranking based image retrieval", in *Proc. of the 12th annual ACM international conference on Multimedia*, pp. 9-16, 2004.

[He2005] J. He, H. Tong, M. Li, W.-Y. Ma, C. Zhang, "Multiple random walk and its application in content-based image retrieval", in *Proc. of the 7th ACM SIGMM international workshop on*

*Multimedia information retrieval*, pp. 151-158, 2005.
[Hoi2004] C.-H. Hoi, M.R. Lyu, "A novel log-based relevance feedback technique in content-based image retrieval", in *Proc. of the 12th annual ACM international conference on Multimedia*, pp. 24-31, 2004.
[Hoi2006] C.H. Hoi, M.R. Lyu, R. Jin, "A unified log-based relevance feedback scheme for image retrieval", in *IEEE Transactions on knowledge and data engineering v.18 n.4*, pp. 509-524, 2006.
[Jin2005] X. Jin, J.C. French, "Improving image retrieval effectiveness via multiple queries", in *Multimedia Tools and Applications*, pp. 221-245, 2005.
[Lew2006] M.S. Lew, N. Sebe, C. Djeraba, R. Jain, "Content-based multimedia information retrieval: state of the art and challenges", in *ACM Transactions on Multimedia computing, communications, and applications v.2 n.1*, pp. 1-19, 2006.
[Lin2005] Y.-Y. Lin, T.-L. Liu, H.-T. Chen, "Semantic manifold learning for image retrieval", in *Proc. of the 13th annual ACM international conference on Multimedia*, pp. 249-258, 2005.
[Luo2004] J. Luo, M.A. Nascimento, "Content-based sub-image retrieval using relevance feedback", in *Proc. of the 2nd ACM international workshop on Multimedia databases*, pp. 2-9, 2004.
[Nguyen2005] G.P. Nguyen, M. Worring, "Scenario optimization for interactive category search", in *Proc. of the 7th ACM SIGMM international workshop on Multimedia information retrieval*, pp. 159-166, 2005.
[Rhamani2005] R. Rhamani, S.A. Goldman, H. Zhang, J. Krettek, J.E. Fritts, "Localized content based image retrieval", in *Proc. of the 7th ACM SIGMM international workshop on Multimedia information retrieval*, pp. 227-236, 2005.
[Rocchio1966] J.J. Rocchio. "Document retrieval systems – optimization and evaluation", *PhD thesis, Harvard Computational Laboratory, Cambridge, MA*, 1966.
[Shah2004] A. Shah-Hosseini, G.M. Knapp, "Learning image semantics from users relevance feedback", in *Proc. of the 12th annual ACM international conference on Multimedia*, pp. 452-455, 2004.
[Smeulders2000] A. Smeulders, M. Worring, S. Santini, A. Gupta, R. Jain, "Content-based retrieval at the end of the early years", in *IEEE Transactions on Pattern Analysis and Machine Intelligence v.22 n.12*, pp.1349-1380, 2000.
[Tong2001] S. Tong, E. Chang, "Support Vector Machine active learning for image retrieval", in *Proc. of ACM international conference on Multimedia*, pp. 107–118, 2001.
[Wang2005] L. Wang, X. Li, P. Xue, K.L. Chan, "A novel framework for SVM-based image retrieval on large databases", in *Proc. of the 13th annual ACM international conference on Multimedia*, pp. 487-490, 2005.
[Xie2004] H. Xie, A. Ortega, "An user preference information based kernel for SVM active learning in content-based image retrieval", in *Proc. of the 6th ACM SIGMM international workshop on Multimedia information retrieval*, pp. 1-6, 2004.
[Xin2004] J. Xin, J.S. Jin, "Relevance feedback for content-based image retrieval using Bayesian network", in *Proc. of the Pan-Sydney area workshop on Visual information processing*, pp. 91-94, 2004.
[Yang2005] C. Yang, M. Dong, F.Fotouhi, "Semantic feedback for interactive image retrieval", in *Proc. of the 13th annual ACM international conference on Multimedia*, pp. 415-418, 2005.
[Yoshizawa2004] T. Yoshizawa, H. Schweitzer, "Long-term learning of semantic grouping from relevance-feedback", in *Proc of the 6th ACM SIGMM international workshop on Multimedia information retrieval*, pp. 165-172, 2004.
[Zhang2003] H. Zhang, Z. Chen, M. Li, Z. Su, "Relevance feedback and learning in content-based image search", in *World Wide Web: internet and web information systems*, pp. 131-155, 2003.
[Zhang2005] C. Zhang, X. Chen, "OCRS: an interactive object-based image clustering and retrieval system", in *Proc. of the 6th international workshop on Multimedia data mining: mining integrated media and complex data*, pp. 71-78, 2005.
[Zhou2002] X. Zhou, T. Huang, "Relevance feedback in content-based image retrieval: some recent advances", in *Information Sciences v.148 n.1*, pp. 129-137(9), 2002.
[Zhou2003] X. Zhou, T. Huang, "Relevance feedback in image retrieval: a comprehensive review", in *Multimedia Systems v.8 n.6*, pp. 536-544, 2003.
[Zhou2006a] X. Zhou, C. Ang, T. Ling, "Dynamic interactive spatial similarity retrieval in iconic image databases using enhanced digraph", in *Proc. of the 2006 ACM symposium on Applied computing*, pp. 46-51, 2006.
[Zhou2006b] Z.-H. Zhou, K.-J. Chen, H.-B. Dai, "Enhancing relevance feedback in image retrieval using unlabeled data", in *ACM Transactions on Information Systems v.24 n.2*, pp. 219-244, 2006.