Utilizing Context in Ranking Results from Distributed CBIR[1]

Christian Hartvedt
Department of Information Science and Media Studies
University of Bergen

Abstract

Selection and ranking of relevant images from image collections remains a problem in content-based image retrieval. This problem becomes even more visible and acute when attempting to merge and rank multiple result sets retrieved from a distributed database environment. This paper presents findings from a project that investigated if combining text and image retrieval algorithms with the use of image context can help reduce the problem of merging and ranking distributed results [1]. The evaluation of our approach, implemented in a system called CAIRANK (Context-Aware Image Ranking), shows that it returns significantly better results than a more traditional ranking approach based on using DBMS-normalized image similarity scores alone.

1 Introduction

The increase of computational power in hardware and software combined with improvements in bandwidth, has greatly improved organisations possibilities to store, manipulate and display images. This has lead to the rapid deployment and use of digital image collections on the Internet. Museums and other cultural organisations can capitalize on this development, and may now with relative ease present their material on the web through presentations and virtual exhibitions either alone or in a joint effort as a cooperating consortium.

There has been much research done on image retrieval resulting in various approaches and methods used to retrieve images. Currently, the two main approaches to image retrieval are low-level image retrieval commonly known as content based image retrieval (CBIR), and high-level image retrieval based on text retrieval of images using image annotations.

Low-level image retrieval is based on the use of an integrated feature-extraction/object-recognition subsystem that automates the process of feature-extraction and object-recognition. This process is based on analysis of the low-level image features colour, texture, location and/or shape to index images and later retrieve images based on similarity [2].

By contrast, high-level retrieval is most often based on either using the semantic content of the images as modelled by a set of manually assigned attributes managed within the framework of a conventional database management system (DBMS), or by annotating images using free text describing the image content. Traditional information retrieval techniques are then employed to carry out the text-based image retrieval (ibid).

Current CBIR algorithms are not optimal. They are most often based on retrieving images with a feature structure *similar* to the feature structure of an example image, and the better the system, the more similar images it will retrieve. The actual relevance an image has to a user though, may reside on other factors than similarity alone. Different users may also have different information needs varying from query to query, but they will perhaps still want to use the same example image. Thus, a system executing queries using low-level features alone also returns images based solely on the extracted image features available without considering semantic information. This normally results in a

[1] *This paper was presented at the NIK-2007 conference. For more information, see //www.nik.no/.*

somewhat haphazard result with regards to relevance, and this is one example of the problem called the semantic gap in image retrieval.

An important part of the semantic gap problem, is the fact that visual similarity does not necessarily correspond to a semantic relationship. Therefore, the images retrieved on the basis of the low-level features of an example image are not necessarily related to it on a semantic level [3]. The semantic gap thus refers to the discrepancy that exists between the information currently possible to extract from visual data and the interpretation the same data has for a user in a given situation [4].

As an information need does not appear in a vacuum, context also plays a vital role both concerning what a user is interested in, and perhaps more importantly with regard to the image context of a particular image. Image context can be regarded as all information not directly derived from the visual properties, or low-level features [3]. Alternatively image context can be considered as any information that can be used to characterize the situation of an entity [5].

1.1 Challenges in distributed image retrieval

With distributed image retrieval, the difficulties associated with traditional CBIR are also present, but the distributed image retrieval process is further complicated by such factors as the context of the individual image collections and the similarity algorithms used in the image retrieval systems.

One of the main challenges in distributed image retrieval is how to take advantage of the potential that lies in having several databases available for image retrieval. Distributed collections can help users gain access to a richer and better source of information, and the challenge is how to provide a user with the most relevant result set possible. One goal here is to have one result list regardless of the number of databases queried, and that the result list has the same quality as would be experienced retrieving images from a single source [6].

A major obstacle associated with distributed systems consisting of database systems from different vendors is that different DBMS' use different algorithms for both query processing and result ranking [7]. These database system (DBS) specific algorithms are most often kept secret in commercial systems and thus are not available to the developers of systems to be used for distributed image retrieval. The consequence is that result sets returned from different sources may not be ranked in a compatible way, resulting in a situation that complicates the process of merging distributed result sets.

Prior work demonstrates that significant improvement is often seen when combining CBIR algorithms with text retrieval algorithms on an arbitrary collection of result sets from different information retrieval systems [8].

This paper summarizes results from a study aimed at investigating if a new way of combining both image and text retrieval algorithms with the utilization of image context in distributed image retrieval is a viable approach to improving the ranked results with regards to relevance.

The main motivation behind the study was to evaluate if the utilization of existing DBMS functionality in a new way could contribute in narrowing the Semantic Gap.

The experimental approach was evaluated through the development of a prototype called Context Aware Image Ranking – CAIRANK [1].

2 Previous work

There has been much research done on both low-level and high-level image retrieval, and some proposals for combining the two approaches have also been put forth. There has also been some research done on using various forms of context in image retrieval, and some of the contributions provided are briefly discussed in the following sections.

2.1 Combining Low- and High-level image retrieval

There are both strengths and weaknesses associated with the low- and high-level approaches to image retrieval. Current low-level image retrieval algorithms cannot capture the high-level abstractions contained in images, while most of the current high-level image retrieval methods depend on annotations that may be partial or incomplete and subjective [2].

A possible approach could be to integrate the two approaches to reduce their weaknesses while capitalizing on their strengths [2, 3, 8, 9].

2.2 Combining Image Queries with Text-based Queries

There have been several suggestions for methods combining image queries with text-based queries. Some of these approaches utilize keywords derived from the top results in a result set produced by a content-based query. Keywords extracted from these results then act as query terms in the subsequent queries which retrieves images associated with documents containing these keywords [8, 10, 11]. Queries consisting of example images submitted to systems such as these, have most often been set up as either *query by example* (QBE), *query by sketch* (QBS).

Other approaches use a text-based query as a starting point and retrieve images in near proximity of words similar to those given as query criteria [3, 9, 12].

2.3 Retrieving Information Items Utilizing Context

User-evaluation of the relevance of an image is often dependent on context, and this implies that the performance of a raking function may also be very context dependent [13].

A key component of a retrieval system operating in such an environment is context awareness in the query process. An early approach to developing methods utilizing context in information retrieval was so-called context-aware retrieval applications. These applications could be interactive, where the user directly issued a request to retrieve relevant data items, or proactive, where documents were presented to the user automatically [14].

To help define the field of context-aware applications, Dey et al. [5] presented a list of context-aware features that these applications should support:

1. *presentation* of information and services to a user dependent on the perceived user context.
2. automatic *execution* of a service triggered by the context of the user, and
3. *tagging* of context data to information elements for later retrieval

The *Fast Search & Transfer* (FAST) search engine [15] is one example of a system conducting context-aware computing called contextual insight. The system supports text-based information retrieval from distributed databases by collecting, processing and storing the processed data it in a central database.

None of the approaches discussed above combine content-based queries with text-based queries while also utilizing context in a distributed setting, and none of the systems supports the use of a combination of content-based queries with text-based queries on full-text document collections. Instead, most of them rely on manually added annotations given to all images stored in single image collections. The FAST system [15] is an exception and supports information retrieval from full-text documents, but does not support content-based image retrieval in the current version of the system. In addition, the FAST approach makes use of a centralized database solution where the contents of all participating databases are copied into a central repository for further processing. The other systems proposed for combining content-based image retrieval and text retrieval support only keyword searches against recorded annotations in combination with the content-based image retrieval, and these systems are not developed for distributed settings. Moreover, most of these systems do not allow for users to choose both the seed image and specify the query terms.

The proposed CAIRANK system, discussed next, represents an approach that combines content-based image retrieval algorithms with text-retrieval algorithms developed for full-text queries with the use of context for improved quality in the ranking of result sets in a distributed setting. The CAIRANK approach also involves the user more in the process of formulating queries in that the user specify both the seed image and the query terms to be used in the query.

3 The current study

One purpose of the project was to investigate if the approach represented by the CAIRANK system could result in more relevant images in merged and ranked result sets from distributed image queries.

The research question underlying the project was if better quality in the search results could be achieved by combining text and image retrieval algorithms with the use of full-text context descriptions. This opposed to most current solutions, which utilize only a DBMS-normalized similarity score from image comparison when ranking the results.

The primary novelty of the approach presented here lies in the way in which some methods and approaches from different research fields are adapted and customized for use in the CAIRANK prototype and thus combined into a new approach for retrieving and ranking of images from distributed database systems.

What sets this approach apart from most existing approaches that combines high-level image retrieval with low-level image retrieval is primarily the usage of full-text document searches as opposed to the use of keyword searches on recorded annotations implemented in most existing systems. Moreover, contrary to other systems that combines high- and low-level in their image retrieval solutions, the CAIRANK approach also differ in that image context is used actively to help support users retrieve relevant images.

In addition, the CAIRANK system also combines the scores returned from the text queries and image queries into a new global score where the image score and text score have equal weighting. This means that results consisting of results scoring high on *both* image and text will surpass results scoring high with regards to image score *or* text score alone, thus preventing that the obvious superiority of text-based information retrieval influences the results too heavily. This should make sure that the results recorded are a result of the combination of scores from image and text retrieval, and not a result of text-based retrieval being superior to content-based image retrieval.

3.1 Framework for specifying context

The framework used for defining image context descriptions in this project drew on the context categories presented in Dey et al. [5]: *Activity*, *Identity*, *Location* and *Time*.

Activity, answers a fundamental question of what is occurring in the situation depicted, while *Identity* characteristics provides us with the possibility of acquiring many pieces of information about the image. *Location* can be vague or concrete and may aid in determining what other objects or people are near the entity and perhaps activities occurring near the entity. *Time*, like different hours of the day, is also helpful to use in specifying image context, e.g. in searching for day or night images [5].

3.2 The CAIRANK approach

The CAIRANK approach as suggested by Hartvedt [1], is built on meta-search engine technology and focuses on image retrieval from distributed databases using data items consisting of images accompanied by full-text context descriptions.

This approach makes use of each participating DBS' implemented solutions for content-based image retrieval from the image collection and text-based queries against stored context information in order to create improved grounds for relevance-based ranking in a combined-score approach. This solution should return images similar to the syntactical low-level features of a seed image as well as being relevant from a context standpoint. To help ensure that the final ranking of the distributed query results are correct regarding their relative relevance, the similarity scores in sub-results from the participating DBS' were normalized locally in each DBS by user defined functions using a *Min-Max Normalization* approach: For each query, the minimum score was subtracted from each score in the result, and this sum was divided by the product of subtracting the minimum score from the maximum score for the given query. This produced a new normalized score with a range going from "0" towards "1" (ibid).

The merging and ranking approach taken here, hence was to calculate a new score by combining the DBS-normalized similarity scores returned from content-based and text-based queries in each of the participating DBS' and multiply them with the weight assigned to each DBS, before merging and ranking the results using an adaptation and customization of the *weighted merging* method described in [16-18].

The weights assigned to each participating DBS were determined using a series of test queries to assess the effectiveness of each participating DBS by using the formula depicted in equation 1. This formula is presented in Yu and Meng [19] where the authors describe the ProFusion formula originally introduced by Gauch et al. [16].

$$c * \frac{\sum_{i=1}^{10} N_i}{10} * \frac{R}{10}$$

Equation 1 - Formula for calibrating the effectiveness of a database

In the equation above, c is a constant applied to all participating databases. N_i is $1/i$ if the i-th ranked document is useful and 0 otherwise. R is the number of relevant documents in the first 10 retrieved documents for a given query. Using this approach in the CAIRANK environment opened for distinguishing between the influence of results retrieved from different DBS' thus possibly alleviating some of the potential problems associated with using normalized sub-scores from DBS' using different retrieval and ranking algorithms in the image retrieval process. This differentiation was to be achieved by factoring in the relative weight for each participating DBS into the global score.

3.3 Prototype development

As illustrated in figure 1, the CAIRANK processing system consists of two main components: a search engine for meta-searching and an application for merging and ranking of data items according to global scores. Two databases with text and image collections provided the test environment.

The development platforms for the database systems used by the processing system were Oracle 9*i* and IBM UDB DB2 v8 using their respective procedural languages. Microsoft Excel was used for merging and ranking of the results.

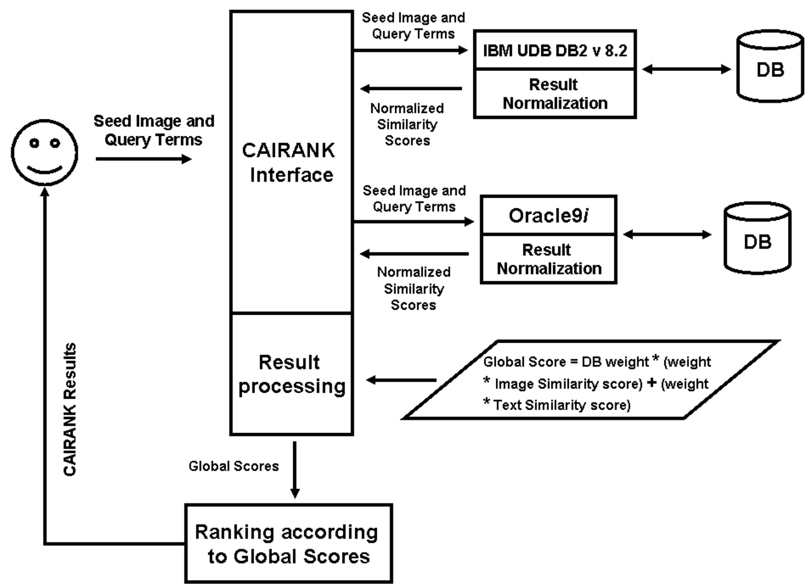

Figure 1 - CAIRANK Ranking Process

4 Evaluation of the CAIRANK approach

In order to investigate if any differences existed between the combined-score approach and the single-score approach concerning abilities to rank multiple result sets in a distributed system environment, the CAIRANK prototype was equipped with an extended version of a Raw Score Merging method used for comparison. The traditional Raw Score merging method used only DBS-normalized similarity scores from CBIR, while the version implemented in the CAIRANK prototype combined DBS-normalized similarity scores from text and image queries multiplied with a pre-determined DB weight assigned to each participating DB.

Two hypotheses were put forth. The first was that this new combined-score approach would improve local result sets. The second hypothesis was that this improvement would hold when combing result sets from multiple DBS'.

The first hypothesis was aimed at investigating differences in a single-score approach versus a combined-score approach, i.e. using a normalized similarity score from CBIR alone versus normalized similarity scores from both CBIR *and* text-based image retrieval from the stored context descriptions. The second hypothesis was aimed at investigating differences between the combined-score approach implemented in the CAIRANK prototype and a traditional Raw Score Merging approach, based on a single score approach.

4.1 Experimental setup

An experiment was set up to compare the two different ranking approaches in two different settings using the IBM and Oracle DBS'. In the first setting queries were submitted to each DBS. The experimental results using the single-score approach were compared to results from the combined-score approach. In the second setting, queries were submitted simultaneously to the distributed environment consisting of both DBS', i.e. results obtained using the Raw Score Merging method (single-score approach) were compared to results from the extended version implemented in the CAIRANK prototype (combined-score approach) that used two similarity scores and database weights to calculate the combined score score.

The experiment set out to record the effect of introducing the usage of two similarity scores in the process of ranking the query results. This effect was measured by recording two different Precision measures from the search results obtained in the two different settings.

Twelve test queries, each consisting of an example image accompanied by text query terms representing one of the context categories described above, were submitted to the DBS'. The query in figure 2 is an example of a query to be translated to SQL and used in the experiment.

Find images resembling seed image 1. context of interest: Bridges being struck by or being near lightning

Figure 2 - Query consisting of image with query text representing the activity category

The SQL statement formulated from the query shown in figure 2 contained only the word "lightning" to be used together with the query image. Reducing the number of words used was done in an effort to reduce the possibility of researcher bias in the queries since the author created both the context descriptions and the queries. A single-word query was thought as having less chance of being unintentionally adapted to the context descriptions which was not actually written by the author, but comprised from relevant articles published in the free encyclopaedia wikipedia. These articles were stored unchanged.

Based on the query submitted, images similar to the seed image were identified, given a similarity score and retrieved in both settings by the DBMS'. In addition, the CAIRANK processing system also gave a global score to each image retrieved in the distributed setting.

According the query setting, Distance or Precision was used to measure the quality of the result sets. Distance was defined as the number of places between the placement of an image in the ideal result set for a query and the actual placement of the same image in the result set for that query. Precision is the fraction of the retrieved documents which is relevant.

For the queries submitted separately to each of the participating DBS', performance was evaluated using a set of ideal images and Distance to measure displacement. Displacement refers to the actual number of places an image in the result set for a query has been misplaced compared to the placement in the ideal set for that query. For the queries submitted in the distributed setting, Precision was used to measure the quality of the results provided by CAIRANK. This created two different measures that was used to evaluate the performance of the proposed combined-score approach compared to the single-score approach.

4.2 Test collection

The test collection consisted of 800 images of 84 different bridges, collected from various sources on the Internet. The collection was divided into two sub sets consisting of roughly the same number of images of the same bridges. Each sub set was then assigned to one of the two participating DBS' resulting in an evenly distributed image collection.

The reasons for constructing an image collection exclusively for this particular project were twofold. Firstly, there was a lack of suitable image collections available when starting up this project. Secondly, by developing a specialized collection, maintaining control and having a clear overview over the images and ideal result sets was possible.

Every image was given a general full-text context description of various aspects of the image content using the different context categories discussed above. In the project, *Activity* referred to what was going on in an image, e.g. fog, lightning or construction. *Identity* referred to identifying characteristics such as names or nicknames, e.g. Big Apple and Big Easy for New York and New Orleans respectively. *Location* referred to placement and could be vague or concrete. For instance, America is less concrete than Boise, Idaho, which is less concrete than 1600 Pennsylvania Avenue, Washington DC, USA. *Time* referred to what time related information was present in an image, e.g. day versus night, or high versus low tide.

The author chose both the image and context description collections, developed the text-based queries used to accompany the seed images in the experiment, as well as selecting the seed images used in the queries and determining the ideal result sets.

4.3 Results

The result data used in the ranking process were normalized image similarity scores alone for the Raw Score method, and a new calculated score consisting of image- and text similarity scores multiplied with the database weights for the CAIRANK approach.

Precision was calculated for each query by using the top 20 results with intervals of 5 documents as the inspection point for calculating the Precision in 20 % intervals for the results. Each inspection point of the distributed queries thus resulted in 4 inspection points where the average Precision value was recorded for the 5 images constituting each interval. Based on the average Precision results for each query, an average Precision value for all queries was calculated for the two ranking methods.

The average results Precision are displayed in figure 3. As illustrated by the graphs in the figure, the results recorded using the CAIRANK prototype are much better than results achieved by using the Raw Score approach indicating that the combined-score approach by far outperforms the single score approach.

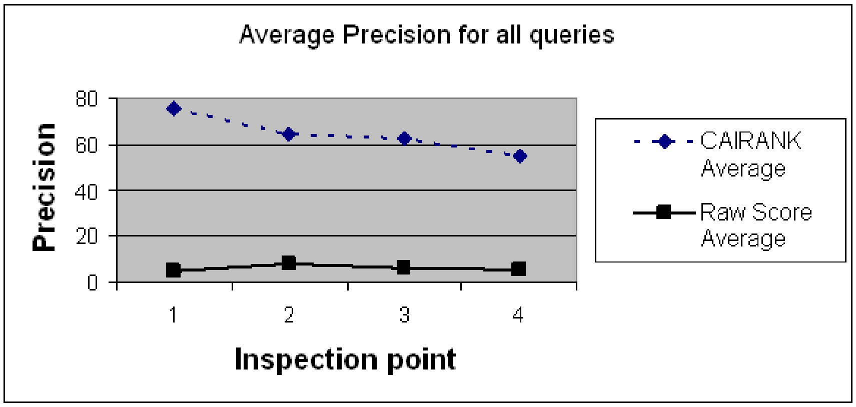

Figure 3 - Average Precision values

Distance was calculated in order to assess the actual degree of displacement of images by the two ranking approaches for each query submitted the participating databases by using all images in each result set. By using the ideal relevance set developed for each query, the distance between the placement of images in the ideal set and their placement in the result set was counted.

Figure 4 displays the combined average Distance between images in the ideal set for each query from the actual placement of those images in the result set for each query in the two participating DBS'. Again, results from the combined-score approach are far better than results recorded using the single score approach on almost every query.

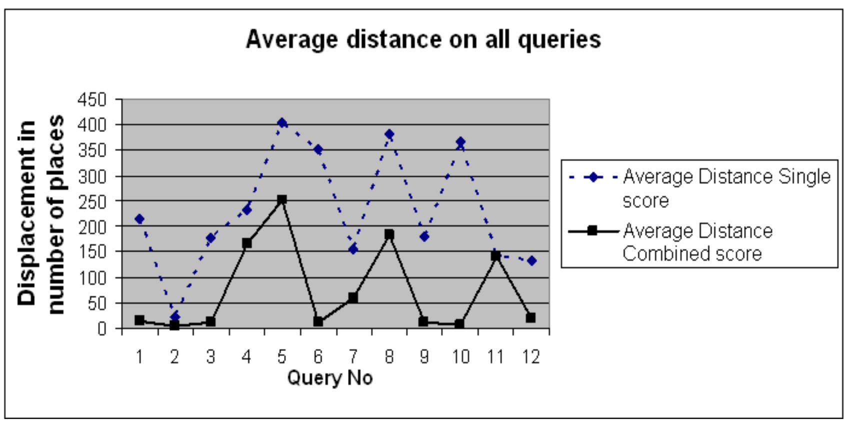

Figure 4 - Average Distance values

Approximately five ideal images were chosen in each DBS for each query. Some queries had more than five images while other had less. A Distance value showing the average displacement from the ideal images for each query was also calculated as illustrated in figure 5 and figure 6. The figures show the average number of positions between the images placed from 1-5 in the ideal set and the actual placement of all images placed as number one, number two, number three etc. on all queries in each of the participating DBS'. The combined-score approach resulted in far less displacement from the ideal image placement in both DBS' than was the case with the single score approach.

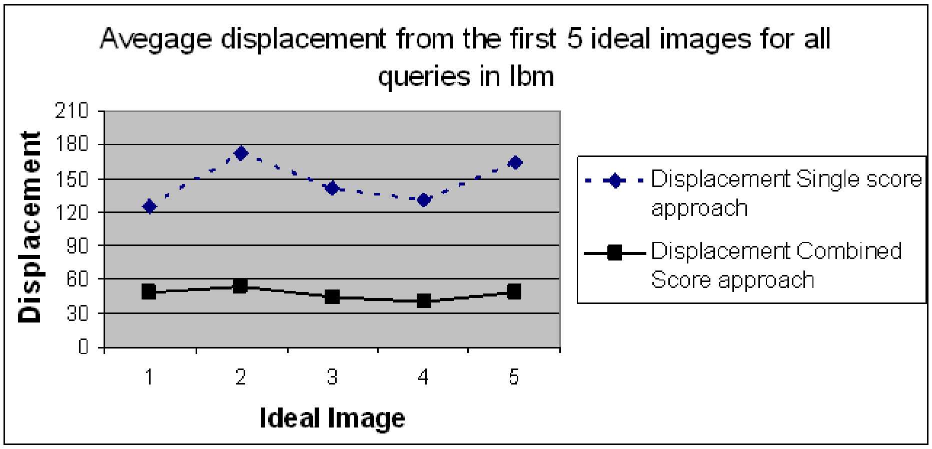

Figure 5 - Average displacement Ibm results

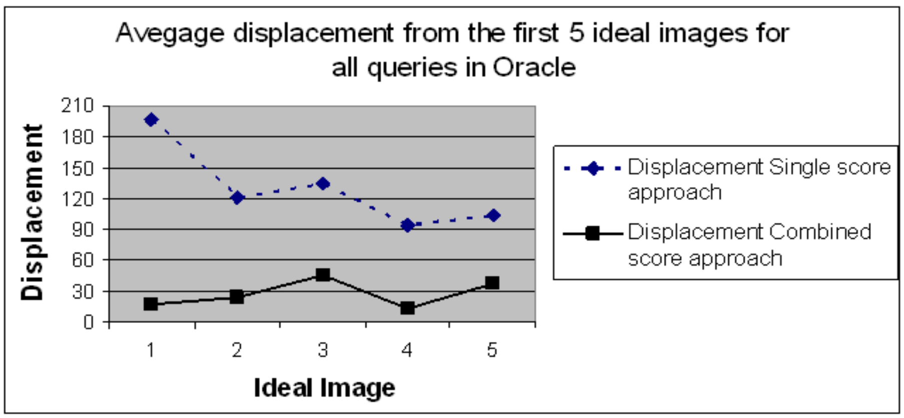

Figure 6 - Average displacement Oracle results

As both hypotheses put forth in the project was directional, the average results from both ranking approaches in both the local settings and the distributed query settings were subjected to a paired, one-tailed Student-t significance test to determine if the results show a significant difference between the approaches. The alpha level was set to 0.05 in all tests. The significance level for the average Precision measures was well below the alpha level: ($t = 12{,}7$; $p<0{,}001$; Df 3). This was also the case with the average Distance measures where results from the significance tests was ($t = 4{,}3$; $p< 0{,}001$; Df 11) and ($t = 4{,}0$; $p=0{,}001$; Df 11) in Ibm and Oracle respectively. This indicates that all results recorded in the experiment were statistically significant.

5 Conclusions and future research

The combined-score approach represented by the CAIRANK prototype has shown that an implementation of a method that can merge and rank distributed results using both content-based and text-based scores combined with the use of image context is feasible.

The results recorded in the experiment indicate that ranking of result sets using a combined-score approach clearly outperforms result ranking using a single-score approach with regards to both Precision and Distance measurements. Precision is a great deal higher on all inspection points when ranking distributed results using the CAIRANK prototype compared to the traditional Raw Score approach. In the results recorded in each of the participating DBS', Distance was generally much lower with the combined-score approach than with the single-score approach.

Results from the significance tests indicate that there are significant differences in favour of the combined-score approach both for the datasets produced in each participating DBS and when using the CAIRANK prototype in a distributed setting.

The experimental results generated in this project support the hypotheses; a combined-score approach is capable of achieving significantly better Precision and Distance values in both a local and distributed setting than a system based on use of normalized similarity scores from CBIR alone.

However, both the experiment and the foundation from which the results are generated were created specifically for this project, and were also based on several assumptions. Firstly, it was assumed that describing content in full-text using general descriptions in natural language from an external source could function better than annotations since the same context descriptions could be associated with all images displaying the same object in similar circumstances. This could potentially lessen the burden of manually annotating all images with keywords, and also possibly remove some of the negative drawback associated with subjectivity in the descriptions. Secondly, it was assumed that having the text stored as full-text documents written in natural language was better that using annotations as this allowed for usage of a wider range of query functionality in the DBMS, e.g. fuzzy queries, queries about words in near proximity, or queries about certain themes.

The project also made several assumptions about the nature of the underlying data used in the experiment. It was assumed that the samples were independently and randomly drawn from the source population, that the scale of measurement for both samples had the properties of an equal interval scale, and that the source populations could be reasonably supposed to have a normal distribution. These assumptions must be addressed and investigated before any final conclusions may be drawn regarding answering the research question. Still, there are clear indications that it may be well worth the effort to conduct further studies of a combined-score approach to image retrieval.

More thorough testing than what was possible within the timeframe of a master's degree is needed in order to be able to draw any well founded conclusions about the usefulness and real world performance of the CAIRANK prototype. In order to conduct a more extensive evaluation of the framework presented here, the CAIRANK prototype should therefore be tested on a larger scale, having more DBS' participating, adding several new image collections from different domains and involving participants drawn randomly from an external population.

One fairly clear indication given by the experimental results is the need for better CBIR systems than the two systems used in this experiment. It would be very interesting to see how the CAIRANK prototype would perform if used with a CBIR system that provides good support for object recognition in images, but unfortunately no such system exists at this time.

The actual impact of using only full-text documents in the form used by CAIRANK for retrieving images would also be interesting to investigate further in order to determine how much the text influence the results, and to see if this influence is just positive or both positive and negative.

Acknowledgements

The work referenced here has been done as a master's thesis within the framework of NFR project #176858, CAIM - Context Aware Image Management. I would like to thank my advisor, associate professor Joan Nordbotten, for her guidance and valuable discussions and insights.

References

1. Hartvedt, C., *Utilizing context in ranking results from distributed image retrieval. Master's Thesis*, in *Dept. of Information Science and Media Studies*. 2007, University of Bergen: Bergen.
2. Lu, G., *Multimedia database management systems*. Artech House computing library. 1999, Boston: Artech House. 373.
3. Westerveld, T. *Image Retrieval: Content versus Context*. in *RIAO 2000 Conference Proceedings*. 2000. Paris.
4. Smeulders, A.W.M., et al., *Content-Based Image Retrieval at the End of the Early Years* IEEE Transactions on Pattern Analysis and Machine Intelligence 2000. **22**(12): p. 1349 - 1380
5. Dey, A.K., et al., *Towards a better understanding of context and context-awareness.* Handheld and Ubiquitous Computing, Proceedings, 1999. **1707**: p. 304-307.
6. Voorhees, E.M., N.K. Gupta, and B. Johnson- Laird. *The Collection Fusion Problem*. in *TREC-3*. 1994.
7. Steidinger, A. *Comparison of different Collection Fusion Models in Distributed Information Retrieval*. in *Proceedings of the First DELOS Network of Excellence Workshop on "Information Seeking, Searching and Querying in Digital Libraries"*. 2000.
8. Müller, H., P. Ruch, and A. Geissbuhler, *Enriching content-based image retrieval with multi-lingual search terms.* Swiss Medical Informatics, 2005. **54**: p. 6-11.
9. Zhou, X.S. and T.S. Huang, *Unifying keywords and visual contents in image retrieval.* Ieee Multimedia, 2002. **9**(2): p. 23-32.
10. Jones, G.J.F., et al. *Dublin City University at CLEF 2004: Experiments with the ImageCLEF St Andrew's Collection*. in *Proceedings of the CLEF 2004: Workshop on Cross-Language Information Retrieval and Evaluation*. 2004. Bath, UK.
11. Orphanoudakis, S.C., C.E. Chronaki, and D. Vamvaka, *I Cnet: Content--based similarity search in geographically distributed repositories of medical images.* Computerized Medical Imaging and Graphics, 1996. **20**(4): p. 193 - 207.
12. Sclaroff, S., M.L. Cascia, and S. Sethi, *Unifying textual and visual cues for content-based image retrieval on the world wide web.* CVIU, 1999. **75**(1-2): p. 86 - 98.
13. Wang, L., et al., *Ranking Function Discovery by Genetic Programming for Robust Retrieval.* , in *The Twelfth Text Retrieval Conference*. 2003, NIST: Gaithersburg. p. 828-836.
14. Brown, P.J. and G.J.F. Jones, *Context-aware Retrieval: Exploring a New Environment for Information Retrieval and Information Filtering.* In press, 2001.
15. Øhrn, A. *Contextual insight in search: enabling technologies and applications* in *Proceedings of the 31st international conference on Very large databases* 2005. Trondheim, Norway: VLDB Endowment
16. Gauch, S., G. Wang, and M. Gomez, *Profusion: Intelligent fusion from multiple, distributed search engines.* The Journal of Universal Computer Science, 1996. **2**(9): p. 637-649.
17. Yu, C. and W. Meng. *Web Search Technology* The Internet Encyclopedia 2003 [cited 2006 November 4th]; Available from: www.cs.binghamton.edu/~meng/pub.d/yumeng2.doc
18. Callan, J.P., Z. Lu, and W.B. Croft. *Searching distributed collections with inference networks*. in *ACM-SIGIR*. 1995. Seattle (WA).
19. Yu, C. and W. Meng, *Web Search Technology* in *The Internet Encyclopedia*, H. Bidgoli, Editor. 2003 Wiley Publishers. p. 738-753.