BTU DBIS' Personal Photo Retrieval Runs at ImageCLEF 2013

Thomas Böttcher, David Zellhöfer, and Ingo Schmitt

Brandenburg Technical University, Database and Information Systems Group,
Walther-Pauer-Str. 2, 03046 Cottbus
tboettcher@tu-cottbus.de, david.zellhoefer@tu-cottbus.de,
schmitt@tu-cottbus.de

Abstract. This paper summarizes the results of the BTU DBIS research group's participation in the Personal Photo Retrieval subtask of ImageCLEF 2013. In order to solve the subtask, a self-developed multimodal multimedia retrieval system, PythiaSearch, is used. The discussed retrieval approaches focus on two different strategies. First, two automatic approaches that combine visual features and meta data are examined. Second, a manually assisted relevance feedback approach is presented. All approaches are based on a special query language, CQQL, which supports the logical combination of different features.
Considering only automatic runs without relevance feedback that have been submitted to the subtask, DBIS reached the best overall results, while the relevance feedback-assisted approach is placed second amongst all participants of the subtask.

Keywords: Content-Based Image Retrieval, Preference Based Learning, Relevance Feedback, Polyrepresentation, Experiments

1 Introduction

This paper summarizes the results of the BTU DBIS research group's participation in the Personal Photo Retrieval subtask of ImageCLEF 2013 [3].

As in DBIS' participations in various ImageCLEF tasks between 2011 and 2012, the discussed approaches rely on the commuting quantum query language (CQQL) [16]. CQQL is capable of combining similarity predicates as found in information retrieval (IR) as well as relational predicates common in databases (DB) and has been on of the main research fields of the database and information systems work group at the Brandenburg Technical University (BTU).

CQQL is an extension of the relational domain calculus, i.e., it can be directly executed within a relational DB system [11]. To combine both data access paradigms, CQQL relies on the mathematical foundations of quantum mechanics and logic. For the sake of brevity, the theoretical background of the query language is omitted. For further details, please refer to the central CQQL publication [16]. Additional information, e.g., the relation of CQQL to fuzzy logic can be found in [17]. Its relation to probabilistic IR models is discussed in [26].

In the scope of this paper, CQQL is used for the matching within the used multi-modal multimedia retrieval system PythiaSearch [24, 23], which has been developed by DBIS. The system consists of an extraction module for both visual features and meta data that supports various image formats and PDF, a matching component relying on CQQL, and a full-featured GUI supporting graded relevance feedback. In order to carry out the matching between query documents and a document collection, CQQL combines various features with the help of logical connectors.

1.1 Personal Photo Retrieval Subtask

The personal photo retrieval subtask 2013 is an extension of 2012's pilot task. The current subtask uses 5,555 image documents that have been sampled from personal photo collections. In contrast to the pilot phase of the task, 2013's focus lies on the evaluation of retrieval algorithms using different search strategies and user groups. One objective of the task is to assess whether a retrieval algorithm's effectiveness is stable for different user groups [28, 10]. To test the effectiveness for different users, multiple ground truths are provided reflecting relevance assessments of CBIR/MIR experts, laypersons or the like.

The subtask does not provide any training data. Hence, it has to be solved ad-hoc. The participants are given multiple query-by-example (QBE) documents and/or browsed documents and are asked to find the best matching documents illustrating an event or depicting a visual concept. In total, 74 topics are available. In contrast to the last year, the topics are no longer separated into visual concepts or events [10]. Furthermore, the information need (IN) for each topic is not explicitly given. Instead, the IN is concealed inside the query or/and browsed documents. To infer an IN, the participants get 0-1 QBE document and up to 3 browsed documents. For some topics, there are no QBE documents in order to model the following usage behavior: a user browsed a personal photo collection and toggled an action to show more similar images without stating an explicit preference [10]. The provided browsed documents can sometimes be irrelevant or have only a low degree of relevance. A more detailed description of the subtask's experimental setup, objective, and participation is available in [28].

2 PythiaSearch - an Interactive and Multi-modal Multimedia Retrieval System

The interactive retrieval system PythiaSearch [24, 23] forms the core for both the interactive and non-interactive retrieval experiments that are described in this paper. In order to express their IN, users can input images (following the QBE paradigm that is used in the subtask), (multilingual) texts, or PDF documents. Additionally, it supports a relevance feedback (RF) process that can be used to personalize the query results based on the user's interaction with the system. The interactive parts rely on a common code base for feature extraction and similarity calculation with the baseline system [28] that has been provided by the organizer of the subtask. Figure 1 shows the GUI of the system. A full description of the GUI and its conceptual model has been published before [23].

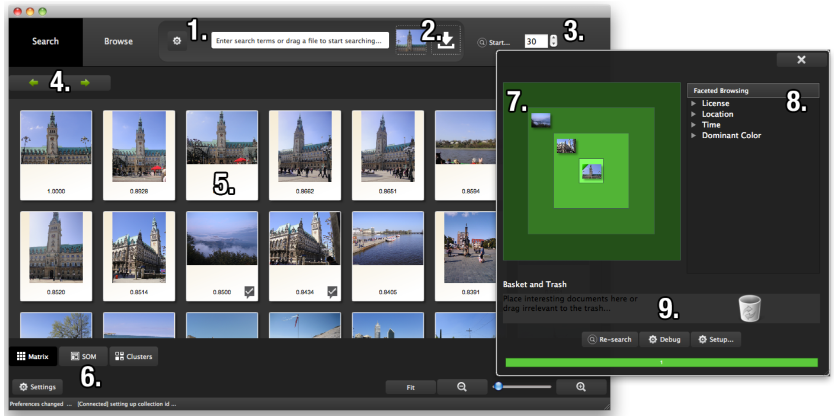

Fig. 1. PythiaSearch - graphical user interface

2.1 Evaluation of CQQL

As said before, we will neglect the theoretical foundations of CQQL to facilitate the understanding of this paper. Mathematically interested readers are recommended to refer to [16] and [13]. The arithmetic evaluation of a CQQL statement which consists of multiple conditions that are connected by logical connectors is directly derived from the mathematical framework of quantum mechanics and logic. In this section, we will sketch the arithmetic evaluation of CQQL as far as it is necessary for the understanding of this paper.

Let $f_\varphi(d)$ be the evaluation of a document d w.r.t. a CQQL query. To construct a CQQL query, various conditions φ can be linked in an arbitrary manner using the conjunction (1), disjunction (2), or negation (3). If φ is atomic, $f_\varphi(d)$ can be directly evaluated yielding a value out of the interval $[0; 1]$. For the scope of this paper, an atomic condition is the result of a similarity measure, e.g., the similarity of the QBE document's color histogram and d's color histogram; or a Boolean evaluation calculated by a DB system or the like.
After a necessary syntactical normalization step [27], the evaluation of a CQQL query is performed by recursively applying the succeeding formulas until the atomic base case is reached:

$$f_{\varphi_1 \wedge \varphi_2}(d) = f_{\varphi_1}(d) * f_{\varphi_2}(d) \tag{1}$$

$$f_{\varphi_1 \vee \varphi_2}(d) = f_{\varphi_1}(d) + f_{\varphi_2}(d) - (f_{\varphi_1}(d) \wedge f_{\varphi_2}(d)) \tag{2}$$

$$f_{\neg\varphi}(d) = 1 - f_\varphi(d) \tag{3}$$

An example of the arithmetic evaluation of the query that is used in this paper is given in Section 3. In accordance with the Copenhagen interpretation of quantum mechanics, the result of an evaluation of a document d yields the probability of relevance of d w.r.t. the query. This probability value is then used for the ranking of the result list of documents.

Weighting in CQQL In order to reflect the need for the personalization of a query and as a necessary step for the support of relevance feedback (RF), CQQL has been extended with a weighting scheme [15]. The weights in CQQL can be used to steer the impact of a condition on the overall evaluation result. Weighting is a crucial part of the machine-based learning supported RF mechanism that is used in Section 3 and discussed in more detail in [27] and [23].

The weighting in CQQL is fully embedded into the logical query. That is, a query maintains its logical properties while weights are used. To illustrate, Equation 4 denotes a weighted conjunction, whereas Equation 5 states a weighted disjunction. A weight θ_i is directly associated with a logical connector and steers the influence of a condition φ_i on the evaluation. To evaluate a weighted CQQL query, the weights are syntactically replaced by constant values according to the following rules:

$$\varphi_1 \wedge_{\theta_1,\theta_2} \varphi_2 \rightsquigarrow (\varphi_1 \vee \neg\theta_1) \wedge (\varphi_2 \vee \neg\theta_2) \tag{4}$$

$$\varphi_1 \vee_{\theta_1,\theta_2} \varphi_2 \rightsquigarrow (\varphi_1 \wedge \theta_1) \vee (\varphi_2 \wedge \theta_2) \tag{5}$$

2.2 Result Personalization and Relevance Feedback

As implied before, the relevance judgement of a query's results is very subjective with respect to the user's IN. To refine a subjective IN, PythiaSearch supports a gradual relevance feedback on the basis of partially ordered sets (posets) [27]. Users can input a poset of documents which contains an arbitrary amount of documents at various relevance levels. For instance, a poset can define a preference expressing that a document D_i is better than a document D_j. This form of user input requires no background information of the underlying features and is based on the subjective qualitative perception of the user alone. Figure 2 illustrates the mechanism as it is implemented in PythiaSearch's GUI. In this example, the second ring contains documents considered more relevant than those on the third etc., while the center contains the current QBE document.

Internally, a machine-based learning algorithm (a downhill simplex variant) is used to find appropriate weight values for a given CQQL query fulfilling the input preferences. The actual algorithm and its properties is described separately in [27].

3 Experimental Setup and Results

Motivated by CQQL's support for formulating multi-modal queries, DBIS participated in the 2011 Wikipedia Retrieval task at ImageCLEF [25] combining textual and visual features. This year's participation in the Personal Photo Retrieval subtask focuses on a CQQL-based combination of visual features and the accompanying meta data. This poses a new challenge for the working group because the studies carried out before were not relying very much on meta data.

Our experiments for the Personal Photo Retrieval subtask can be subdivided into two types of runs. First, fully automatic runs which demonstrate the effectiveness of a CQQL-based logical combination of features with different origin, i.e., visual and meta data features. Second, the performance of the aforementioned RF mechanism is investigated (see Section 3.4).

3.1 Used Features

Over the last years, the DBIS working group conducted a lot of experiments on various image collections, ranging from the Caltech collections [7, 8] to MSRA-MM [20] in order to assess the retrieval effectiveness of different low-level visual features. This investigation of single features forms the basis for the decision which features to combine with CQQL.

PythiaSearch supports the extraction of low-level global and local visual features, e.g., color, edges and texture features or local features like SIFT and SURF. In total, the extraction component offers more than 30 visual features. Additionally, the system allows the extraction of common image meta data such as Exif, IPTC, or XMP. This meta data, e.g., GPS coordinates, the camera model, or the image orientation extends the variety of features that can be used for the matching of documents. In accordance with the rules of the subtask, IPTC-based data is ignored in the following experiments. Table 1 lists all features that are used in the experiments.

The feature extraction and similarity calculation functionality used by PythiaSearch resembles the baseline system that is provided by the subtask organizer[1]. For a description, see [28]. The main differences between the baseline system and the system used for the described experiments are the CQQL support, the supplementary RF mechanism, and the GUI.

3.2 Examined CQQL Query

Based on preparatory study on the retrieval effectiveness of various visual low-level features and an examination of the subtask's thematic orientation on both visual concepts and events, an appropriate CQQL query had to be defined. The core idea of the examined CQQL query, which is shown in Equation 6, is to use low-level features that did show a good performance over all six test collections

[1] Please note that the organizer of the subtask did not actively participate in the experiments described in this paper nor did he release additional information to the working group that other participants could not obtain. Alas, he carried out many of the pre-studies including the investigation of generally effective CQQL queries. Furthermore, he had a major impact on the development of PythiaSearch and the underlying learning algorithm.

Table 1. Features and type, * denotes features in the scope of MPEG-7 [12]

Name	Type	Origin
Auto Color Correlogram (ACC)	color-related, global	[9]
BIC	color-related, global	[18]
CEDD	texture/color-related, global	[4]
Color Histogram (region based)	color-related, pseudo-local	[1]
Color Histogram	global	256 bin RGB histogram (own implementation)
Color Layout*	color-related, global	[6]
Color Structure*	color-related, global	[6]
Dominant Color*	color-related, global	[6]
Edge Histogram*	edge-related, global	[6]
FCTH	texture/color-related, global	[5]
Scalable Color*	color-related, global	[6]
Tamura	texture-related, global	[19]
Region-based Shape*	global	[6]
Person Detection	facial features	Own implementation (based on OpenCV)
Time of creation	temporal	Exif
GPS coordinate	spatial	Exif
Camera model	metadata	Exif

Table 2. Overview over the examined test collections in the preparatory study

Collection	Collection Size	Collection Type
Caltech 101 [7]	9,197	Object recognition
Caltech 256 [8]	30,607	Object recognition
MSRA-MM [20]	65,443	Web Image Sample
UCID [14]	904	Personal photos
Wang [21]	1,000	Stock photography
Pythia [22]	5,555	Personal photos

that are listed in Table 2 to cope with the visual content of a document. In order to retrieve similar events, we assume that the presence (or absence) of persons in a picture, spatial and temporal proximity as well as a similar camera model are valid indicators. Hence, the core CQQL query is enriched by a person presence condition in form of a Boolean predicate and the aforementioned features derived from Exif meta data.

This concept results in the following CQQL query that uses a weighted conjunction of 18 conditions, whereas all weights are set to 1 initially to express the equal importance of all conditions.

$$\begin{aligned}\bigwedge_{\theta_i}(&ACC_{sim}, BIC_{sim}, CEDD_{sim}, ColorHistBorder_{sim}, ColorHistCenter_{sim},\\ &ColorHist_{sim} ColorLayout_{sim}, ColorStructure_{sim}, DominantColor_{sim},\\ &EdgeHist_{sim}, FCTH_{sim}, Regionshape_{sim}, ScalableColor_{sim}, Tamura_{sim},\\ &GPS_{sim}, model_{sim}, time_{sim}, Person_{sim})\end{aligned} \tag{6}$$

The value of each condition is determined by a distance measure such as the Euclidean distance of the corresponding feature between the QBE document and the retrieved document which is then transformed into a similarity measure in the interval of $[0; 1]$. Boolean conditions are evaluated traditionally. The calculation of the GPS coordinate similarity is carried out as we did for the ImageCLEF 2012 Plant Identification task [2]:

$$GPS_{sim} = 1 - \frac{\sqrt{(71.5 \cdot (long_x - long_y))^2 + (111.3 \cdot (lat_x - lat_y))^2}}{6378.388} \quad (7)$$

whereas *long* stands for longitude and *lat* for latitude.

Following the transformation rules that have been described in Section 2.1, the arithmetic evaluation of the presented CQQL is as follows:

$$(ACC_{sim} + \neg\theta_1 - ACC_{sim} * \neg\theta_1)* \quad (8)$$

$$(BIC_{sim} + \neg\theta_2 - BIC_{sim} * \neg\theta_2)* \quad (9)$$

$$(CEDD_{sim} + \neg\theta_3 - CEDD_{sim} * \neg\theta_3)* \quad (10)$$

$$\dots \quad (11)$$

3.3 Automatic Runs

The 2013 Personal Photo Retrieval subtask provides QBE documents as well as browsed documents. For the approaches without RF, we use the provided data in two ways.

First, we use only the QBE defined documents (*run1*) because of the fact, that the provided browsed documents can contain misleading information, respectively contain images that did not fulfill the users IN. For the topics for which no QBE document is available, we use all browsed documents instead. Whether the browsed documents are really relevant for the user's IN cannot be determined automatically. Thus, this approach might be affected by irrelevant input to the retrieval system.

Second, we assume that all documents (no matter if they are QBE or browsed documents) are equally meaningful regarding the user's IN. Hence, we use all documents as QBE documents with no special ranking for the labelled QBE document. This approach is labeled *run2* in Figure 4.

3.4 Manual Relevance Feedback Run

The main objective for the manually assisted approach (*run 3*) is to remove misleading browsed image documents from the initial query and to improve the retrieval quality using the graded RF approach that is described in Section 2.2. The experiment is carried out interactively with the PythiaSearch GUI (see Section 2). Using the aforementioned preference-based approach, irrelevant, relevant documents and the relationship between them can be expressed as a poset. To simplify the user interaction, the GUI offers 3 levels of relevance and a "garbage can" to collect completely irrelevant documents. Figure 2 shows the three levels where the center (level 1) contains the query document(s). All documents in level 2 are more relevant than documents in level 3 and 4, whereas documents in level 3 are more relevant than documents in level 4. All preferences together create a poset. Documents marked as completely irrelevant are removed from the query results and have a negative impact on the machine-based learning algorithm similar to negative QBE documents.

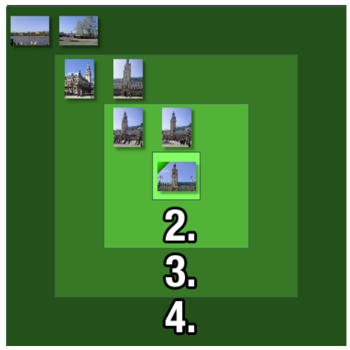

Fig. 2. Graded Relevance Feedback - Preference Levels

In order to keep control over the time consumption of submitting 74 queries manually, we have defined some restrictions for the RF-based experiment.

First, at most one RF iteration is carried out, i.e., we model the behavior of an impatient user.

Second, the assessment of the quality of the results is based on the top-30 results only. This number of documents can be easily inspected without scrolling and requires significantly less time than inspecting the top-50 or top-100. Because of this strategy, it may happen that no RF is carried out all because the top-30 results seem relevant to the interacting user.

Third, to simulate a user that avoids a large amount of interaction with the system, a total of 6 images is used to define the preferences used during RF.

In general, obviously irrelevant images from the given IN specification were removed from the input. Nevertheless, during the submission of all 74 sample queries it was not always possible to identify the IN without background knowledge. In these cases, the RF process is skipped.

4 Results

With reference to the official results (see Figure 4), our best run, i.e., *run 3*, achieves rank 5 in the overall ranking for the average user. Compared to the results of all participants of the subtask, DBIS is ranked second. Focussing on the NCDCG_cut 5 we reach about 97%, on NDCG_cut 100 about 87% and on MAP_cut 100 78% of the best obtained retrieval score.

When only automatic runs are considered, i.e., runs without RF, DBIS achieves the best results. Unfortunately no other runs without RF that use all available modalities (visual and meta data features) and IN information (QBE and browsed documents) were submitted. Due to these circumstances, a resilient interpretation of our results is hardly possible. Anyhow, we assume that the inclusion of meta data helps to distance our from the other approaches. In any case, further information about the techniques used by the other participants is needed.

Generally speaking, the outcome of the presented experiments is fully satisfying. Though, we acknowledge room for improvements for the RF-based run. We assume that an inspection of more than the top-30 results and the inclusion of more preferences might have an impact on the retrieval effectiveness. Furthermore, a specification of the actual IN in textual form will help human assessors during RF because it will enable them to provide RF for every topic. As said before, we could not provide RF for all topics because of the lack of this kind of information. In consequence, we expect an improvement of the RF effectiveness when this information can be used.

One objective of this subtask is to examine the robustness of a retrieval approach with respect to different user groups (e.g. IT experts, non-IT users or gender-specific groups). Figure 3 shows the variance of the MAP_cut 100 scores of our three submitted runs between the different user groups. The differences between all groups is relatively small, e.g., 0.395 vs. 0.425 for *run 3*. The difference in MAP_cut 100 between the best and the worst run is about 7-9%. Interestingly, the results for female users tend to be the best, whereas the average user score tends to be worst. This effect can also be observed in the results of the other groups. Alas, we are not sure why this effect is present amongst all groups.

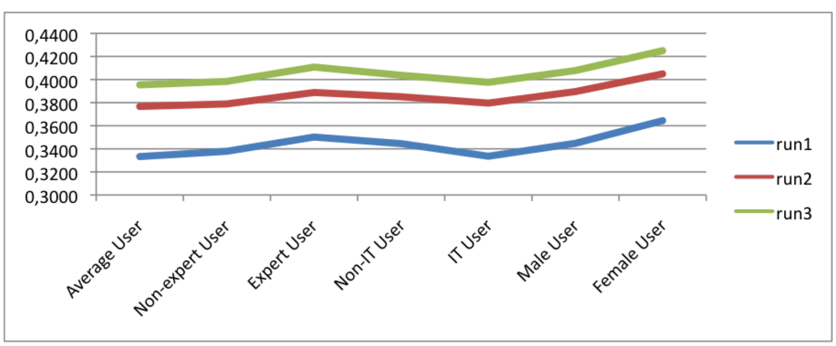

Fig. 3. MAP_cut 100 score comparison on different user types

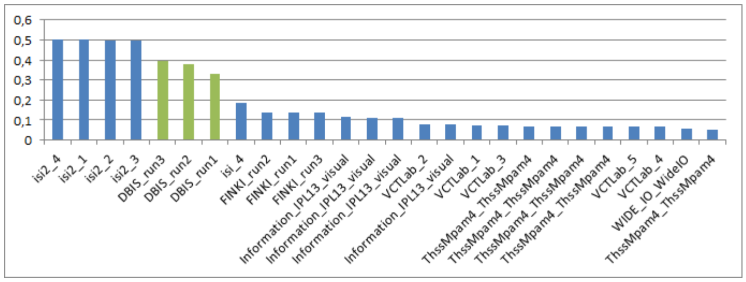

Fig. 4. Overall results considering MAP_cut 100 (average user)

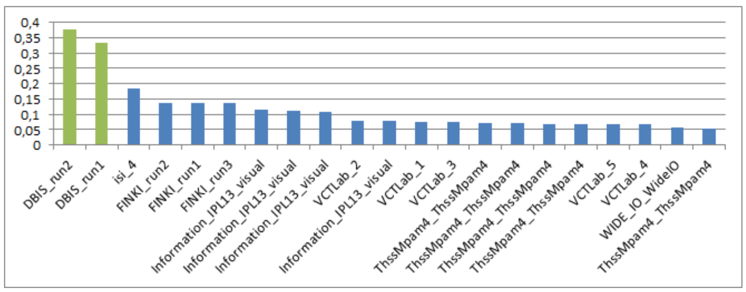

Fig. 5. Results excluding RF runs considering MAP_cut 100 (average user)

5 Conclusions and Future Work

The results of our participation in the ImageCLEF 2013 Personal Photo Retrieval subtask are motivating. Although, DBIS achieved a good effectiveness rank, there are areas that need further research.

First, we plan to analyze single features and meta data in more detail to find out which features or meta data contributes most to the retrieval quality. As said before, for the RF-supported approach there are various optimizations possible. In particular, the restriction to one RF iteration seems to limit the retrieval quality. First informal experiments show that up to three iterations can give a great performance boost. Furthermore, we assume that the inclusion of RF on all topics will lead to a performance improvement. Another interesting research question is the development of the weight values during the RF iterations in order to reveal whether some features do not contribute to the retrieval effectiveness at all.

References

1. Balko, S., Schmitt, I.: Signature Indexing and Self-Refinement in Metric Spaces. Cottbus (2012)
2. Böttcher, T., Schmidt, C., Zellhöfer, D., Schmitt, I.: Btu dbis' plant identification runs at imageclef 2012. In: CLEF (Online Working Notes/Labs/Workshop) (2012)

3. Caputo, B., Mueller, H., Thomee, B., Villegas, M., Paredes, R., Zellhoefer, D., Goeau, H., Joly, A., Bonnet, P., Martinez Gomez, J., Garcia Varea, I., Cazorla, M.: ImageCLEF 2013: the Vision, the Data and the Open Challenges, CLEF 2013 Working Notes, Valencia, Spain (2013)
4. Chatzichristofis, A.S., Boutalis, S.Y.: CEDD: color and edge directivity descriptor: a compact descriptor for image indexing and retrieval. In: Proceedings of the 6th international conference on Computer vision systems. pp. 312–322. ICVS'08, Springer-Verlag (2008), `http://dl.acm.org/citation.cfm?id=1788524.1788559`
5. Chatzichristofis, A.S., Boutalis, S.Y.: FCTH: Fuzzy Color and Texture Histogram - A Low Level Feature for Accurate Image Retrieval. In: Proceedings of the 2008 Ninth International Workshop on Image Analysis for Multimedia Interactive Services. pp. 191–196. WIAMIS '08, IEEE Computer Society (2008), `http://dx.doi.org/10.1109/WIAMIS.2008.24`
6. Cieplinski, L., Jeannin, S., Ohm, J.R., Kim, M., Pickering, M., Yamada, A.: MPEG-7 Visual XM version 8.1. Pisa, Italy (2001)
7. Fei-Fei, L., Fergus, R., Perona, P.: Learning generative visual models from few training examples an incremental Bayesian approach tested on 101 object categories. In: Proceedings of the Workshop on Generative-Model Based Vision (2004)
8. Griffin, G., Holub, A., Perona, P.: Caltech-256 Object Category Dataset (2007), `http://authors.library.caltech.edu/7694`
9. Huang, J., Kumar, R.S., Mitra, M., Zhu, W.J., Zabih, R.: Image Indexing Using Color Correlograms. In: Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR '97). pp. 762–. CVPR '97, IEEE Computer Society (1997), `http://dl.acm.org/citation.cfm?id=794189.794514`
10. ImageClef: Personal Photo Retrieval 2013. `http://www.imageclef.org/2013/photo/retrieval`, 6. June 2013.
11. Lehrack, S., Schmitt, I.: QSQL: Incorporating Logic-Based Retrieval Conditions into SQL. In: Kitagawa, H., Ishikawa, Y., Li, Q., Watanabe, C. (eds.) Database Systems for Advanced Applications, 15th International Conference, DASFAA 2010, Tsukuba, Japan, April 1-4, 2010, Proceedings, Part I, Lecture Notes in Computer Science, vol. 5981, pp. 429–443. Springer (2010)
12. Manjunath, B., Salembier, P., Sikora, T.: Introduction to MPEG-7: Multimedia Content Description Interface. John Wiley & Sons, Inc., New York, NY, USA (2002)
13. van Rijsbergen, C.: The Geometry of Information Retrieval. Cambridge University Press, Cambridge, England (2004)
14. Schaefer, G., Stich, M.: UCID - An Uncompressed Colour Image Database. In: Proc. SPIE, Storage and Retrieval Methods and Applications for Multimedia, pp. 472–480. San Jose, USA (2004)
15. Schmitt, I.: Weighting in CQQL. Cottbus (2007)
16. Schmitt, I.: QQL: A DB&IR Query Language. The VLDB Journal 17(1), 39–56 (2008)
17. Schmitt, I., Zellhöfer, D., Nürnberger, A.: Towards quantum logic based multimedia retrieval. In: IEEE (ed.) Proceedings of the Fuzzy Information Processing Society (NAFIPS). pp. 1–6. IEEE (2008), `10.1109/NAFIPS.2008.4531329`
18. Stehling, O.R., Nascimento, A.M., Falcão, X.A.: A compact and efficient image retrieval approach based on border/interior pixel classification. In: Proceedings of the eleventh international conference on Information and knowledge management. pp. 102–109. CIKM '02, ACM (2002), `http://doi.acm.org/10.1145/584792.584812`

19. Tamura, H., Mori, S., Yamawaki, T.: Texture features corresponding to visual perception. IEEE Transactions on System, Man and Cybernatic 8(6), 460–472 (1978)
20. Wang, M., Yang, L., Hua, X.S.: MSRA-MM: Bridging Research and Industrial Societies for Multimedia Information Retrieval (2009)
21. Wang, Z.J., Li, J., Wiederhold, G.: SIMPLIcity: Semantics-sensitive Integrated Matching for Picture Libraries. In: Proceedings of the 4th International Conference on Advances in Visual Information Systems. pp. 360–371. VISUAL '00, Springer-Verlag (2000), `http://portal.acm.org/citation.cfm?id=647061.714442`
22. Zellhöfer, D.: An Extensible Personal Photograph Collection for Graded Relevance Assessments and User Simulation. In: Proceedings of the ACM International Conference on Multimedia Retrieval. ICMR '12, ACM (2012)
23. Zellhöfer, D.: A permeable expert search strategy approach to multimodal retrieval. In: Proceedings of the 4th Information Interaction in Context Symposium. pp. 62–71. IIIX '12, ACM, New York, NY, USA (2012), `http://doi.acm.org/10.1145/2362724.2362739`
24. Zellhöfer, D., Bertram, M., Böttcher, T., Schmidt, C., Tillmann, C., Schmitt, I.: PythiaSearch – A Multiple Search Strategy-supportive Multimedia Retrieval System. In: Proceedings of the 2nd ACM International Conference on Multimedia Retrieval. p. to appear. ICMR '12, ACM (2012)
25. Zellhöfer, D., Böttcher, T.: BTU DBIS' Multimodal Wikipedia Retrieval Runs at ImageCLEF 2011. In: Vivien Petras, Pamela Forner and Paul D. Clough (eds.) CLEF 2011 Labs and Workshop, Notebook Papers, 19-22 September 2011, Amsterdam, The Netherlands (2011)
26. Zellhöfer, D., Frommholz, I., Schmitt, I., Lalmas, M., van Rijsbergen, K.: Towards Quantum-Based DB+IR Processing Based on the Principle of Polyrepresentation. In: Clough, P., Foley, C., Gurrin, C., Jones, G., Kraaij, W., Lee, H., Murdoch, V. (eds.) Advances in Information Retrieval - 33rd European Conference on IR Research, ECIR 2011, Dublin, Ireland, April 18-21, 2011. Proceedings, Lecture Notes in Computer Science, vol. 6611, pp. 729–732. Springer (2011)
27. Zellhöfer, D., Schmitt, I.: A Preference-based Approach for Interactive Weight Learning: Learning Weights within a Logic-Based Query Language. Distributed and Parallel Databases (2009), `doi:10.1007/s10619-009-7049-4`
28. Zellhöfer, D.: Overview of the ImageCLEF 2013 Personal Photo Retrieval Subtask. CLEF 2013 working notes, Valencia, Spain, 2013 (2013)