

BTU DBIS' Multimodal Wikipedia Retrieval Runs at ImageCLEF 2011

David Zellhöfer and Thomas Böttcher

Brandenburg Technical University, Database and Information Systems Group,
Walther-Pauer-Str. 1, 03046 Cottbus
`david.zellhoefer@tu-cottbus.de`

Abstract. In this work, we summarize the results of our first participation in the Wikipedia Retrieval task. For our experiments, we rely on a cognitively motivated IR model: the principle of polyrepresentation. The principle's core hypothesis is that a document is defined by different representations such as low-level features, or textual content that can be combined in a structured manner reflecting the user's information need. For our first participation, we used mono-lingual English retrieval in combination with global low-level features without further user interaction or query modification techniques.

Our best **NOFB** reached rank 64 or rank 13 of the mono-lingual English runs. This result is promising as we have not used structural information about the documents. Additionally, our findings are indicating the correctness of the polyrepresentative hypothesis for multimodal retrieval.

Keywords: Content-Based Image Retrieval, Information Retrieval, Result Fusion, IR Model, Polyrepresentation, Experiments

1 Introduction

The Database and Information Systems Group of the Brandenburg Technical University (DBIS) is participating for the first time in ImageCLEF's Wikipedia Retrieval task [17]. Hence, our first submitted runs are focussing at gaining experience with the handling of the task and at testing recently published theoretical results.

Over the past years, DBIS has worked on establishing a new query language, the commuting quantum query language (CQQL) [14], which is capable of combining similarity predicates as found in information retrieval (IR) as well as relational predicates common in databases (DB). In addition, CQQL is an extension of the relational domain calculus, i.e., it can be directly executed within a relational DB system [8]. To combine both data access paradigms, CQQL relies on the mathematical foundations of quantum mechanics and logic. For the scope of this paper, we will not emphasize the theoretical implications of the quantum theoretical background. However, the relation of CQQL to fuzzy logic [19] can be found in [15]. Its relation to probabilistic IR models is discussed in [21, 25].

Recently, the correlation between the principle of polyrepresentation [6] and quantum mechanics-inspired IR models has been discussed in the community [21, 5, 24]. Being based on van Rijsbergen’s basal work [11], two approaches are researching a combination of the principle of polyrepresentation with ideas derived from quantum mechanics [10, 23]. Only the latter approach is dealing with multimodal retrieval explicitly. Thus, it will be used as the theoretical IR model for the remainder of this paper.

2 Retrieval Model

The discussed IR model is using CQQL to implement the findings of the cognitively motivated principle of polyrepresentation. The main hypothesis of the principle is that a document is defined by various functionally and cognitively different representations [6]. “*Functionally* different representations are the title of a document, the actual textual content, or an abstract that has been created by a sole actor, e.g., the author. *Cognitively* different representations denote such representations that have been created by another actor such as a retrieval model¹, summaries written by a librarian, or former related INs. Hence, a combination of different IR systems that describe the same document (e.g. examined in [7]) can be regarded polyrepresentative as well.” [25]

If one combines these representations into a so-called “*cognitive overlap*” (CO), documents that are contained within this overlap are likely to be relevant w.r.t. a user’s information need (IN). As a consequence, it is assumed that this combination of representation can compensate the insecurity of relevance assessment of retrieved documents [6]. The utility of the principle of polyrepresentation has been shown for (textual) IR [16, 7]. In addition, both aforementioned works address actual means of construction of a cognitive overlap. As a result, it could be shown that highly structured queries² modeling a CO lead to higher precision in direct comparison to unstructured queries such a bag-of-words.

As said before, our research relies on CQQL to reflect the principle of polyrepresentation in a query language. We believe that the principle provides a sound theoretical framework for dealing with the various representations that are inherent within a multimedia document. As a cognitively motivated model, it does not stop here. Furthermore, it can include additional information that is important during a search task, such as the user’s context. The representations of a document can have different origins such as copyright information stored in a DB, text annotations or low-level features like a color histogram. This data has to be accessed and combined in a consistent manner in order to retrieve multimodal data. CQQL provides means to query both IR and DB systems while maintaining their semantics. For instance, representations from a relational DB can be restricted in Boolean fashion to evaluate only against *true* or *false* (i.e. 1

¹ Ingwersen and Järvelin regard an IR system as a peer actor because it assesses the relevance of a document on basis of different representations. Thus, it participates as well in the search.

² A query combining Boolean connectors to combine parts of a query.

or 0). In contrast, representations originating from an IR system maintain their similarity semantics in an interval of $[0; 1]$ (for details see [14]).

To simplify, CQQL can be considered a query language dealing with probabilities that is consistent with the laws of the Boolean algebra. The probabilities denote how likely it is that a certain document’s representation is relevant to the CO and therefore the user’s IN. In the next section, we will sketch the arithmetic evaluation of CQQL as it is necessary for the understanding of this paper.

2.1 Evaluation of CQQL

Given that $f_\varphi(d)$ is the evaluation of a document d w.r.t. a CO modeled by a CQQL query. To model a CO, various representations 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 $[0; 1]$ As stated before, the actual value of a representation can be calculated by a similarity measure or a Boolean evaluation carried out by a DB system or the like.

After a necessary syntactical normalization step [22], the evaluation of a CQQL query modeling the CO is performed by recursively applying the succeeding formulas until the atomic base case is reached:

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

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

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

An example of an arithmetic evaluation of a sample query is given in Sec. 3.1. The result of an evaluation of a document d yields the probability of relevance of d w.r.t. the given CO. This probability value is then used for the ranking of the result list of documents.

2.2 Weighting in CQQL

In order to steer the influence of certain representation onto the CO, CQQL has been extended with a weighting scheme [13]. This weighting scheme can be used for relevance feedback (RF) during the retrieval process. Weighting is a crucial part of our machine-based learning supported user interaction model discussed in [22, 25]. Although an evaluation of RF for multimodal retrieval is not in the scope of this paper, we will outline how weights are embedded in a CQQL query because they will occur with fixed settings in Sec. 3 as part of the experiments.

Eq. 4 denotes a weighted conjunction, whereas Eq. 5 states a weighted disjunction. A weight θ_i is directly associated with a logical connector and steers the influence of a representation φ_i on the evaluation. To evaluate a weighted

CQQL query, the weights are syntactically replaced by constant values according to the following rules:

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

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

3 Experiments

Recently conducted experiments provided promising results about the utility of the principle of polyrepresentation for the content-based image retrieval (CBIR) domain [25]. The same findings apply to the textual IR domain [16, 7]. As a consequence, our experiments are focussing on the combination of both worlds. Yet, we are not researching representations derived from image metadata or the like. For the scope of this paper, experiments have been conducted only on textual information and low-level features alone. In addition, our first runs do not incorporate any user modification of the query or interaction during the retrieval process such as relevance feedback. For the sake of simplicity, we do not use other languages than English during the experimental retrieval. Although all nodes within the provided XML documents were indexed, structural information such as “title” has not been used explicitly.

In order to obtain the different representations for each document, we used the Indri 5 IR engine from the Lemur project³. All text retrieval has been carried out by this engine. The global low-level features of the images involved in the task were extracted by LIRE 0.8 [9]. Tab. 2 lists all available low-level features that are mostly part of the MPEG-7 standard besides CEDD (also provided during the Wikipedia Retrieval task [17]) and FCTH. Regarding the principle of polyrepresentation, the LIRE features are different *functional* representations of a document while Indri provides a *cognitively* different one. To address the fact that multiple query-by-example (QBE) image documents were available for each query, we have created averaged representations to keep the evaluation simple (see Sec. 4 for details). To obtain averaged representations, we calculated a similarity value for each representation of each QBE image and a given document of the collection. Then, we calculated the arithmetic mean over all QBE image similarities for a given document.

Before our ImageCLEF-related experiments, we have analyzed the performance of different queries that involved LIRE-based representations alone (see Tab. 1). In order to obtain meaningful results, we tested various features and feature combinations with the following well-known collections “Wang” [18], “UCID v2” [12], and “Caltech 101” [4]. Additionally, we have used the “Bali” collection [25] as a representative of personal holiday photographs. Fig. 1 illustrates some of the findings. For the scope of this paper it is notable that the highly structured queries Q_{10} and Q_9 perform almost as well as Q_{AND} or Q_{MAX} while relying on about half of the available features. The features used in Q_{10} and Q_9 have

³ <http://www.lemurproject.org/>

been chosen because of their performance when used on their own and to mirror assumptions about the utility of a combination of structural and color-based features as stated in [3, 2].

To evaluate the principle of polyrepresentation’s utility for multimodal retrieval, we used two different strategies to model the CO and to assess the relevance of a document w.r.t. a given query.

Table 1. Analyzed Sample Queries (Excerpt) from [25]

Name	Query
<i>COLORHIST</i>	$COLORHIST_{sim}$
Q_{MAX}	$\max_{sim}(SCALABLECOLOR_{sim}, COLORLAYOUT_{sim},$ $EDGEHISTOGRAM_{sim}, TAMURA_{sim}, COLORHISTOGRAM_{sim},$ $CEDD_{sim}, GABOR_{sim}, AUTOCOL.CORRELOGRAM_{sim}, FCTH_{sim})$
Q_{AND}	$\bigwedge_{\theta_i}(SCALABLECOLOR_{sim}, COLORLAYOUT_{sim},$ $EDGEHISTOGRAM_{sim}, TAMURA_{sim}, COLORHISTOGRAM_{sim}$ $, CEDD_{sim}, GABOR_{sim}, AUTOCOL.CORRELOGRAM_{sim}, FCTH_{sim})$
Q_9	$(CEDD_{sim} \vee_{\theta_1, \theta_2} FCTH_{sim}) \wedge (COLORLAYOUT_{sim} \vee TAMURA_{sim})$
Q_{10}	$(CEDD_{sim} \vee_{\theta_1, \theta_2} FCTH_{sim}) \wedge (COLORLAYOUT_{sim} \vee (TAMURA_{sim} \wedge$ $EDGEHISTOGRAM_{sim}))$
Q_{13}	$COLORLAYOUT_{sim} \wedge TAMURA_{sim}$
	Weights (θ_i) are initially set to 0.5. XY_{sim} denotes similarity of a representation to the QBE document.

3.1 CQQL-based Fusion

Our first approach used the direct evaluation of CQQL queries. The experiment subdivides into three parts that are related to the used modalities, features, and structural characteristics of the query itself. If weights were present in a query, they were set to 0.5.

Group 1 consists of the queries Eq. 6, Eq. 7, and Eq. 8. All have in common that the LIRE-based part of the CO model is based on Q_{10} (see Sec. 3). For the discussed runs, *COLORLAYOUT* has been removed. This is due to the fact that this feature could not be extracted for a number of image documents. See Tab. 2 for a list of all missing features. This problem has been caused by corrupted ICC profiles within the ImageCLEF image document corpus. These corruptions resulted in a crash of Java’s ImageIO during file access.

The textual part of the CO (denoted as $INDRI_{sim}^x$, whereas x stands for the query type) has been calculated by Indri using different search strategies. All Indri queries used Krovetz stemming and case normalization⁴. The English

⁴ See <http://lemur.sourceforge.net/indri/IndriIndexer.html> for the standard settings.

query text defined by the task [17] served directly as query terms. The differences of the Indri query formulation are as follows.

$$Q_{71-72} : INDRI_{sim}^{ordered\ window} \wedge ((CEDD_{sim} \vee_{\theta_1, \theta_2} FCTH_{sim}) \wedge (TAMURA_{sim} \wedge EDGEHISTOGRAM_{sim})) \quad (6)$$

Here, the English query text has been input directly including the information that the query terms in a document have to appear ordered with at most 4 terms between them using Indri’s ordered window operator. A sample Indri query for the text “close up of trees” would look like this: `#combine(close up of trees #5(close up of trees))`.

In combination with the LIRE-based representations this results in the following arithmetic evaluation given in pseudo code:

$$q_{10} = (1 - ((1 - (cedd * (1 - (1 - \theta_1)))) * (1 - (fcth * (1 - (1 - \theta_2))))));$$

$$relevance = indri * q_{10};$$

$$Q_{73} : INDRI_{sim}^{ordered\ window+SW} \wedge ((CEDD_{sim} \vee_{\theta_1, \theta_2} FCTH_{sim}) \wedge (TAMURA_{sim} \wedge EDGEHISTOGRAM_{sim})) \quad (7)$$

This query is equal to 6 but uses stop word elimination.

$$Q_{79} : INDRI_{sim}^{plain\ combine} \wedge ((CEDD_{sim} \vee_{\theta_1, \theta_2} FCTH_{sim}) \wedge (TAMURA_{sim} \wedge EDGEHISTOGRAM_{sim})) \quad (8)$$

This query is equal to 6 but drops the ordered window operator resulting in queries such as `#combine(close up of trees)`.

Group 2 consists of one query that uses the modified version of Q_{10} alone in order to measure the “raw” low-level feature performance.

$$Q_{110} : (CEDD_{sim} \vee_{\theta_1, \theta_2} FCTH_{sim}) \wedge (TAMURA_{sim} \wedge EDGEHISTOGRAM_{sim}) \quad (9)$$

Group 3 contains COs that have not been sent in before. Because of the missing features for some documents (see Tab. 2), we have investigated the influence of such incomplete information on the retrieval performance of our approach.

$$Q_{500} : INDRI_{sim}^{ordered\ window} \wedge \left(\bigwedge_{\theta_i} (All\ LIRE\ features) \right) \quad (10)$$

For the first query, we used the best performing Indri query in a conjunction with a weighted conjunction of all 9 LIRE features.

$$Q_{600} : \bigwedge_{\theta_i} (INDRI_{sim}^{ordered\ window}, All\ LIRE\ features) \quad (11)$$

The second examined CO is modeled by a weighted conjunction of all LIRE features in addition to the representation provided by Indri. The difference to Eq. 10 is that the textual representation does not have more impact on the result. Here, it is a mere peer amongst all other representations.

Table 2. Missing features after LIRE 0.8 extraction

Feature	Absolute Percentage	
COLORHISTOGRAM	28339	11.90%
SCALABLECOLOR	10116	4.26%
AUTOCOLORCORRELOGRAM	10116	4.26%
COLORLAYOUT	10065	4.23%
EDGEHISTOGRAM	1	<0.01%
CEDD	1	<0.01%
FCTH	1	<0.01%
TAMURA	36	<0.01%
GABOR	36	<0.01%

3.2 Pre-Filtering Approach

For our second approach, we have examined a pre-filtering approach that is closely related to [1]. First, we have calculated the relevance of each document according to Indri’s ordered window operator including stop word elimination as described for Eq. 7. Second, we have calculated the average similarity for each image document. That is, we have computed each LIRE feature similarity. Then, all similarity values have been averaged to obtain an averaged similarity value. Furthermore, we averaged the result over all QBE documents for each query. The result list of the image retrieval was then reduced to the top 5000 documents.

To combine both result lists, we regarded the Indri list as a pre-filter. In other words, we only considered documents that are present in this list as potentially relevant. In a next step, we obtained the relevance score of a document d in the Indri list as the document’s base score. If document d was also present in the low-level feature-based result list, we improved its relevance score by 0.15. This reflects the principle of polyrepresentation’s hypothesis that two cognitively different representations “pointing” at one document are more likely to be relevant.

This approach leaves the proposed usage of CQQL and is merely based on heuristics. The choice of a boosting factor of 0.15 is based on experiments carried out with the ImageCLEF 2010 collection and ground-truth alone. A formal justification cannot be given. This approach is listed in Tab. 3 as run #64.

3.3 Strategy Comparison

In compliance with the findings from [1] our pre-filtering approach performed best. A direct comparison to the other contributors is difficult as not enough information about their runs is available at the present moment. Nevertheless, we compare our approaches with the best performing, automatic English-only textual run #45. Tab. 3 gives an overview of our overall performance. We assume that run #45 uses some kind of optimized query formulation strategy while we are using a fairly naive approach towards the textual retrieval part. That is, we basically hand the query text directly to the Indri IR engine and neglect additional structural information, e.g., if a term is present in a document’s title.

Still, our pre-filtering approach #64 reaches 84,39% MAP performance of the best textual run #45, whereas the pure CQQL runs are in the range of 80.58% to 75.39%. As expected, the plain visual run #110 performs very bad (1.31% MAP of #45). This behavior is in accordance with the result trend from former ImageCLEF Wikipedia Retrieval tasks.

An interesting finding from [25] could be duplicated within the Wikipedia Retrieval task. It seems that there is evidence that a structured combination of functionally different representations (here: low-level features) can compensate a large amount of representations or even surpass it in terms of retrieval quality. In our scenario, run #500 makes use of 9 CBIR-based representation while run #64 and #71-79 use only 4, i.e., 44 % of the available features. This is strengthening the polyrepresentative hypothesis as well as it might be a fruitful area of future research because the actual calculation of such representation can be a costly operation – especially in the field of multimedia retrieval. Concluding from the results of [25], we did expect a slightly better performance for queries involving all features (see run #500) in comparison to, e.g., #71. In our opinion, the presence of missing representations lead to the inverse effect. Hence, it cannot be generalized that a structured combination of representations will outperform an unstructured query in every case. This leaves room for further research.

Run #600 (without an emphasis on textual representations) results in a MAP performance of 19.38% clearly showing that the main contribution of the retrieval quality is derived from textual representations.

Regarding fusion strategies, the findings from the textual domain as stated in [16, 7] that highly structured queries *alone* will improve the retrieval quality cannot be justified for the multimodal retrieval domain for now. Our current results indicate that a cleverly chosen fusion strategy itself is contributing to the overall retrieval quality of a multimodal retrieval system.

4 Conclusion and Future Work

Regarding our low experience with the Wikipedia Retrieval task, we are satisfied with our results. Our best run is place at rank 64, i.e., rank 13 of the mono-lingual English runs. If we consider NOFB runs only, we get rank 4 being the only multimodal approach in this mono-lingual ranking. In accordance with

Table 3. DBIS' runs in comparison to the best automatic English-only textual run

Run	Modality	Language	Feedback	MAP	P@10	Percent of Best Run Performance
<i>45</i>	<i>Textual</i>	<i>English</i>	<i>None</i>	<i>0.2601</i>	<i>0.456</i>	<i>Best Eng.-only Text Run (UNED)</i>
64	Mixed	English	None	0.2195	0.418	MAP: 84.39%; P@10: 91.67%
71	Mixed	English	None	0.2096	0.408	MAP: 80.58%; P@10: 89.47%
72	Mixed	English	None	0.2096	0.408	
73	Mixed	English	None	0.2096	0.408	
79	Mixed	English	None	0.1961	0.382	MAP: 75.39%; P@10: 83.77%
500	Mixed	English	None	0.2061	0.400	MAP: 79.24%; P@10: 87.72%
600	Mixed	English	None	0.0504	0.1660	MAP: 19.38%; P@10: 36.40%
110	Visual	N/A	None	0.0034	0.020	MAP: 1.31%; P@10: 4.39%

our findings about the contribution of textual representations, we assume that the poor performance in comparison to the text-only runs is due to our naive approach towards textual IR. Hence, we would really appreciate the cooperation with other working groups having a strong textual IR experience. In future, we will incorporate the usage of structured information within the XML documents as it is already supported by the CQQL-based query model.

Another criticism of our approach is the over-simplification of the principle of polyrepresentation for the CBIR part of our work. The averaging of representations in order to deal with multiple QBE documents within one query removed important semantics from the query. To model the CO in a better way, each QBE document has to be used as an additional cognitively different representation of the query. We assume that this might improve the retrieval quality because it is likely to reflect the user's subjective IN more complete. In combination with our RF approach [22, 20], this might improve the overlap between the user's cognitive structures – e.g. expressed by different images with different contributions to the user's IN – and the actual query.

Another field for further research is the dealing with missing representations. This is an important issue as it cannot be assumed that all queried representations are always present. For instance, this problem is likely to re-appear when distributed IR systems are used or cross-media retrieval is concerned.

Although it is somewhat out of the scope of system-centric initiatives such as ImageCLEF, we will continue our research on integrating contextual information about the user.

References

1. Clinchant, S., Csurka, G., Ah-Pine, J., Jacquet, G., Perronnin, F., Sánchez, J., Minoukadeh, K.: XRCE's Participation in Wikipedia Retrieval, Medical Image Modality Classification and Ad-hoc Retrieval Tasks of ImageCLEF 2010. In: Braschler, M., Harman, D., Pianta, E. (eds.) CLEF 2010 LABs and Workshops, Notebook Papers, 22-23 September 2010, Padua, Italy (2010)

2. Deselaers, T., Keysers, D., Ney, H.: Features for Image Retrieval: An Experimental Comparison. *Information Retrieval* 11, 77–107 (2008)
3. Eidenberger, H.: How good are the visual MPEG-7 features? In: *SPIE & IEEE Visual Communications and Image Processing Conference*. pp. 476–488 (2003)
4. 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)
5. Frommholz, I., van Rijsbergen, C.: Towards a Geometrical Model for Polyrepresentation of Information Objects. In: *Proc. of the "Information Retrieval 2009" Workshop at LWA 2009* (2009)
6. Ingwersen, P., Järvelin, K.: *The Turn: Integration of Information Seeking and Retrieval in Context*. Springer-11645 [Dig. Serial], Springer, Dordrecht (2005), <http://dx.doi.org/10.1007/1-4020-3851-8> / <http://www.zentralblatt-math.org/zmath/en/search/?an=1079.68532>
7. Larsen, B., Ingwersen, P., Kekäläinen, J.: The polyrepresentation continuum in IR. In: *IIX: Proceedings of the 1st international conference on Information interaction in context*. pp. 88–96. ACM (2006)
8. 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)
9. Lux, M., Chatzichristofis, A.S.: Lire: Lucene Image Retrieval: An Extensible Java CBIR Library. In: *MM '08: Proceeding of the 16th ACM International Conference on Multimedia*. pp. 1085–1088. ACM (2008)
10. Piwowarski, B., Frommholz, I., Lalmas, M., van Rijsbergen, K.: What can quantum theory bring to information retrieval. In: *Proceedings of the 19th ACM international conference on Information and knowledge management*. pp. 59–68. *CIKM '10*, ACM (2010), <http://doi.acm.org/10.1145/1871437.1871450>
11. van Rijsbergen, C.: *The Geometry of Information Retrieval*. Cambridge University Press, Cambridge, England (2004)
12. 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)
13. Schmitt, I.: *Weighting in CQQL*. Cottbus (2007)
14. Schmitt, I.: QQL: A DB&IR Query Language. *The VLDB Journal* 17(1), 39–56 (2008)
15. 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
16. Skov, M., Pedersen, H., Larsen, B., Ingwersen, P.: Testing the Principle of Polyrepresentation. In: Ingwersen, P., van Rijsbergen, C., Belkin, N. (eds.) *Proceedings of ACM SIGIR 2004 Workshop on "Information Retrieval in Context"*. pp. 47–49 (2004)
17. Tsirikika, T., Popescu, A., Kludas, J.: Overview of the wikipedia image retrieval task at ImageCLEF 2011. In: *CLEF (Notebook Papers/LABs/Workshops) : CLEF 2011 working notes* (2011)
18. 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>

19. Zadeh, A.L.: Fuzzy Logic. *IEEE Computer* 21(4), 83–93 (1988)
20. Zellhöfer, D.: Inductive User Preference Manipulation for Multimedia Retrieval. In: Böszörményi, L., Burdescu, D., Davies, P., Newell, D. (eds.) *Proc. of the Second International Conference on Advances in Multimedia* (2010)
21. 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)
22. 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
23. Zellhöfer, D., Schmitt, I.: Approaching Multimedia Retrieval from a Polyrepresentative Perspective: (to appear). In: Detyniecki, M., Knees, P., Nürnberger, A., Schedl, M., Stober, S. (eds.) *8th International Workshop on Adaptive Multimedia Retrieval* (2010)
24. Zellhöfer, D., Schmitt, I.: Ein Polyrepräsentatives Anfrageverfahren für das Multimedia Retrieval. In: Atzmüller, M., Benz, D., Hotho, A., Stumme, G. (eds.) *Proceedings of LWA2010 - Workshop-Woche: Lernen, Wissen & Adaptivitaet* (2010), <http://www.kde.cs.uni-kassel.de/conf/lwa10/papers/ir5.pdf>
25. Zellhöfer, D., Schmitt, I.: A User Interaction Model based on the Principle of Polyrepresentation. In: Nica, A., Suchanek, M.F. (eds.) *Proceedings of the 4th workshop on Ph.D. students in information and knowledge management. PIKM '11, ACM* (2011) (to appear)

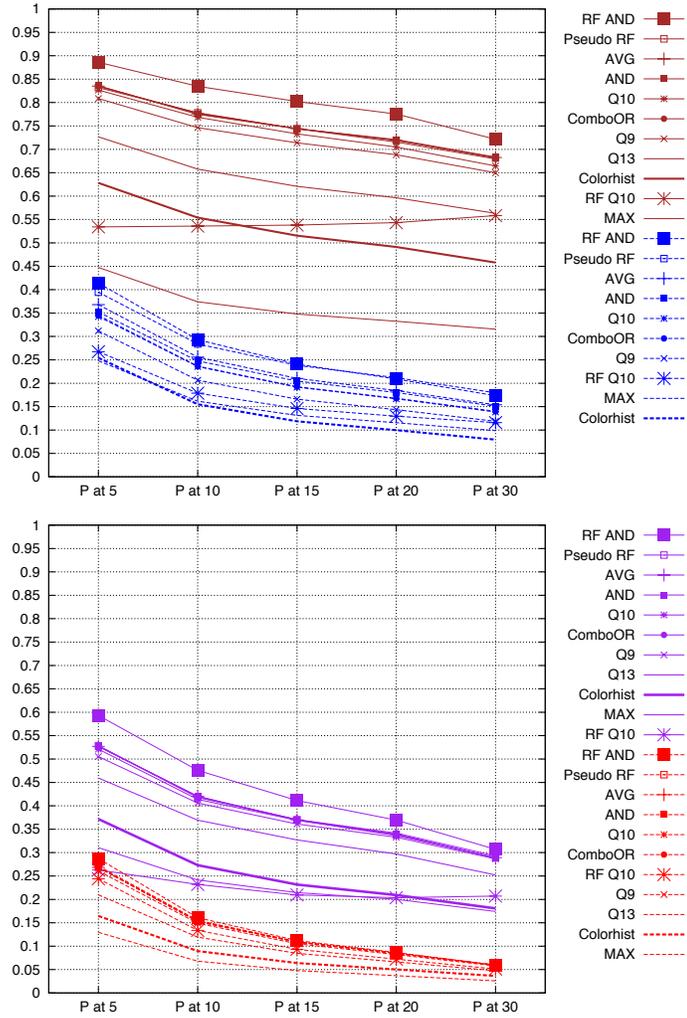


Fig. 1. Comparison of P@X w.r.t. Different Queries, Relevance Feedback Strategies and Document Collections (left: Wang [brown, solid], Caltech 101 [blue, dashed]; right: Bali [purple, solid], UCID [red, dashed]) [25]