In: Eibl, Maximilian; Wolff, Christian; Womser-Hacker, Christa (Hg.): Designing Information Systems. Festschrift für Jürgen Krause. Konstanz: UVK Verlagsgesellschaft mbH, 2005. S. 85 – 99.

An Information Retrieval Prototype for Research and Teaching

Christa Womser-Hacker

Abstract

This paper reports on MIMOR[1] which was modelled as an open meta information retrieval system at the handling of diverse information retrieval objects. It is used to gain further insights concerning the behaviour of information retrieval systems as well as for teaching in the field of Information Science. The main issue of MIMOR is the exploitation of users' relevance feedback in order to optimise the fusion of several retrieval engines or resources. As yet, MIMOR has been applied to a domain specific collection, to a multilingual corpus, and to music data. Experiments within the CLEF evaluation framework and additional ones are discussed here.

Key words

information retrieval, fusion, relevance feedback, evaluation, meta system, multilingual information retrieval

1 Introduction

In Information Science, information retrieval (IR) is seen as the storage and representation of knowledge and the retrieval of information relevant for a special user problem. The information seeker formulates a query which is compared to document representations extracted during the indexing phase. The main issue is to create added value during the transfer process from knowledge to information. According to the basic principle of Information Science which focuses on the user and his environment, the query formulation process is influenced by the users' state of knowledge, their context as well as the quality of the chosen interaction mode. The representations of documents and queries are typically matched by a similarity function such as the Cosine or the Dice coefficient. The most similar documents are presented to the users who then evaluate the relevance with respect to their information need.

Information Retrieval techniques which go beyond the Boolean model must be able to deal with vagueness and uncertainty (Baeza-Yates & Ribeiro-Neto

[1] Multiple Indexing for dynamic Method-Object-Relations

1999). Usually, documents and queries contain natural language or multimedia objects such as graphics, pictures, photos, video sequences or pieces of music. It is a challenging task for the new generation of IR systems to analyse the content of these heterogeneous objects and to make access simple for users. At present, it is not obvious which kind of representation or query mode works best with various kinds of objects. As an example for this heterogeneity MIMOR has been applied to music information retrieval (Mandl & Womser-Hacker 2003) where content-based and symbol-based approaches have been developed to describe musical objects (Fingerhut 2002).

Due to large-scale evaluation studies like TREC[2] (Voorhees & Harman 2001) and CLEF[3] (Peters et al. 2003) many important insights have been gained during the last decade. Furthermore, these initiatives have led to a higher level of comparability in information retrieval. Critical limitations of traditional IR systems have been worked out.

This paper starts by describing the MIMOR model pointing out its basic characteristics and possible expansions. In section 3 we describe the underlying architecture and certain techniques aiming at a progress in efficiency. In section 4, we report on some evaluation results gained by MIMOR until now. The final section will comment our work in progress.

2 The MIMOR Model

MIMOR is modelled as an open information retrieval system which is able to combine individual approaches of information retrieval within one meta system and which could be expanded at different points over time (cf. Womser-Hacker 1997). In our view, MIMOR can serve as an instrument for research and teaching likewise. On the one hand it acts as a research tool for the exploration of the performances of particular retrieval devices and on the other hand students can advance their understanding of retrieval systems and their practical programming skills by working on improvements and additions in the scope of object-oriented programming courses. This combination seems to be very fruitful for both parties.

MIMOR profits from users' relevance assessments in order to learn which combinations of object representations and information retrieval functionality lead to good performance of the overall system. An internal evaluation procedure, which is realized via a blackboard model, permanently registers which

[2] Text REtrieval Conferences
[3] Cross-Language Evaluation Forum

resource produces good results and which one does not. Well-performing techniques gain high weights, poorly-performing ones are excluded over time.

2.1 Basic Assumptions

The MIMOR model takes advantage of the main outcomes of TREC. One of the most important results of this study is that many IR systems perform similarly well in terms of recall and precision but do not lead to the same sets of documents. This means that the systems find the same percentage of relevant documents, but the overlap between their results often is low. Because of these findings, fusion has been pointed out as a promising strategy. This is confirmed by many application experiments of fusion techniques in IR (Fox & Shaw 1994; Vogt & Cottrell 1998; McCabe et al. 1999, Savoy 2002). It turned out that fusion of various IR techniques can improve the overall quality of a system. Fusion methods delegate a task to several systems and integrate their results into one final result set presented to the user. In information retrieval, fusion is mostly implemented as a combination of several algorithms where different probabilities for the relevance of a document query relation are integrated into one final similarity measure. Another kind of fusion is applied by internet meta search engines. These machines have been developed in order to expand the range of single search engines by combining the results. However, it is not clear whether they really lead to better performance. Some empirical studies have shown no improvement (Wolff 2000). In cross-lingual approaches the different languages are brought together by combining multilingual query terms.

A further promising strategy in information retrieval is relevance feedback. It turned out that by taking user assessment into account the systems reached a better quality (cf. Harman 1992, Voorhees 1998).

2.2 Formalization and Technical Background

From a computational point of view, MIMOR is designed as a linear combination of the results of different retrieval systems. The contribution of each system or algorithm to the fusion result is governed by a weight for that system.

$$RSV_{MIMOR}(doc_i) = \frac{\sum_{system=1}^{N} (\omega_{system} RSV_{system}(doc_i))}{N}$$

A central aspect in MIMOR is learning. The weight of the linear combination of each IR system is adapted according to the success of the system measured by the relevance feedback of the users. A system which assigned a high retrieval status value (RSV) and consequently a high rank to a document which then received positive relevance feedback should be able to contribute to the final result with a higher weight. The following formula enables such a learning process, which is also illustrated in figure 1:

$$\omega_{system} = \varepsilon \; RF_{user}(doc_i) \; RSV_{system}(doc_i)$$

$$\varepsilon \quad learning\ rate$$

However, the optimal combination may depend on the context and especially on the users' individual perspectives as well as the characteristics of the documents. Therefore, MIMOR needs to consider context.

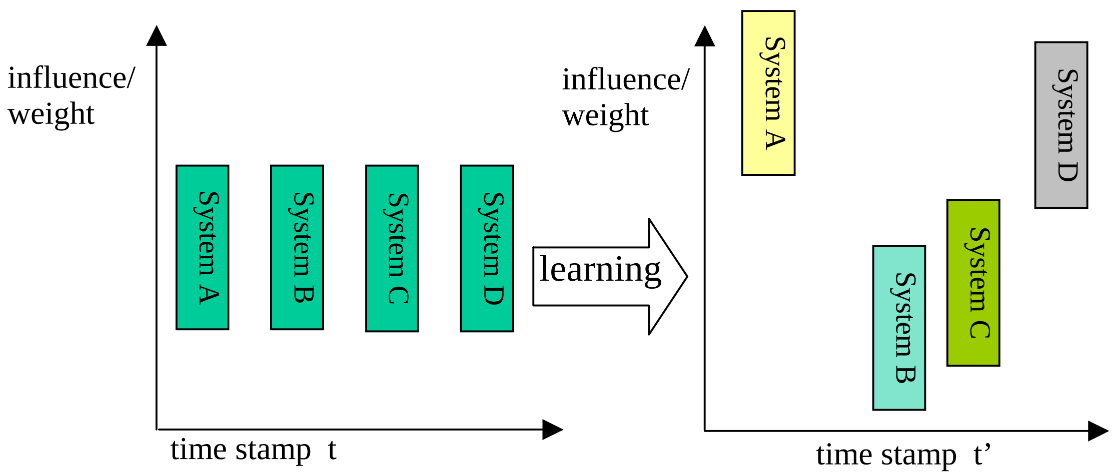

Fig. 1: Learning the optimal linear combination over time

2.2.1 Modeling Context via Clusters

The performance of IR systems differ from domain to domain. TREC found out that particular characteristics of the documents relevant for the indexing procedure may be responsible for this effect. In one experiment, for example, optimal similarity functions especially for short queries could be developed (Kwok & Chan 1998). MIMOR is based upon the idea that formal properties of queries and documents can be exploited in order to improve the overall fusion system. Within fusion, the weight of a system should be high for the type of documents it was optimised for only. Some characteristics of text documents seem to be good candidates for such distinctions. Length, difficulty, syntactic complexity and even layout can be assessed automatically. In MIMOR, these properties are modelled as clusters whereas all documents having a property in common belong to the same cluster. For each cluster an individual MIMOR model is developed with own weights for all participating systems. The clustering process is not restricted to algorithms based on unsupervised learning. Pre-defined classes and even human assignment are compatible with MIMOR.

A theoretical justification for a cluster model can be found in the evaluation strategies for clustering algorithms like minimal description length or category utility (Witten & Frank 2000). Category utility estimates the value of a cluster by checking how well it can be used to predict attribute values of objects. Clusters are good if the probability of an object having a certain value is higher for objects in a specific cluster that for all objects. If good clusters are found and one attribute is an appropriate retrieval system, then the probability is high that a good retrieval system for that specific object is used.

The final result considers only the weight of the cluster the document belongs to. The learning formula has to be modified accordingly. The change in the weight is now applied only to the cluster containing the document.

Clustering documents is a challenging task. In many cases, the exact assignment of a document to only one class is difficult. Therefore, this condition needs to be relaxed. With respect to this approach, fuzzy clustering has been integrated into MIMOR (Mandl & Womser-Hacker 2001).

2.2.2 User Model

Further refinement of MIMOR can be achieved by integrating a user model. Unlike other user models in information retrieval, MIMOR effects an adaptation in the core of an IR system and applies it to the calculation of the RSV.

Similar to the properties of the documents, an additional MIMOR model for each person could be introduced. This would lead to optimal user models. However, the training of a MIMOR model requires a substantial amount of relevance feedback decisions. Therefore, the user is forced to submit many decisions before he can use the system effectively. Another disadvantage is common to all inductive and incremental learning algorithms. The occurrence of some unusual cases in the initial learning phase may lead to an unstable learning curve. This may result in a degradation of the retrieval behaviour.

Both problems are solved by introducing separate private and public models. The private model contains a user specific MIMOR model optimised by all the relevance feedback decisions of that user. The public model is trained with all decisions of all users of the system. The public MIMOR system is optimised but not individualized. Therefore, it can be used for any user beginning to work with MIMOR because an individual model is not available. Over time, such a beginner will collect a significant number of relevance judgements and will eventually reach a fully individualized and saturated model. During this process, the public model will lose its influence while the importance of the private model grows.

The user model in MIMOR differs from many individualization approaches in information retrieval. Often, the individual preferences are stored as a content model. Many systems use interest vectors. MIMOR applies individualization to the algorithmic layer of the system.

3 System Architecture

MIMOR relies on the combination of multiple retrieval components. This fact has led to the demand for a device facilitating the fusion of e.g. different types of databases (relational, object-oriented, XML etc.) or HTTP search engines or technologies like web services. A special REtrieval COmponent INtegrator (RECOIN) takes on the role of a translator between heterogeneous software components within MIMOR and aims at creating a modular and therefore easily scalable application by at the same time reducing programming effort[4]. RECOIN delivers a plug-in mechanism for dealing with different data models, database technologies, protocols, and query languages and works as an integrating component in the middle tier of a multitier application.

According to the principle of an EIS (Enterprise Information Systems) or backend tier where usually a variety of database management systems (DBMS) and legacy applications can be found, RECOIN is located in the middle tier connecting not only the client and the backend tier, but also providing ways to integrate components working in the same or other middle tiers. Figure 2 illustrates a three-tier model. The tiers themselves can be located on different computers collaborating in a distributed effort. RECOIN was primarily designed with the question in mind how it could be used in information retrieval and of what avail it would be especially for the MIMOR model (cf. Kassem 2000).

[4] The development of RECOIN is described in detail in Scheufen 2002, a master thesis at the University of Hildesheim.

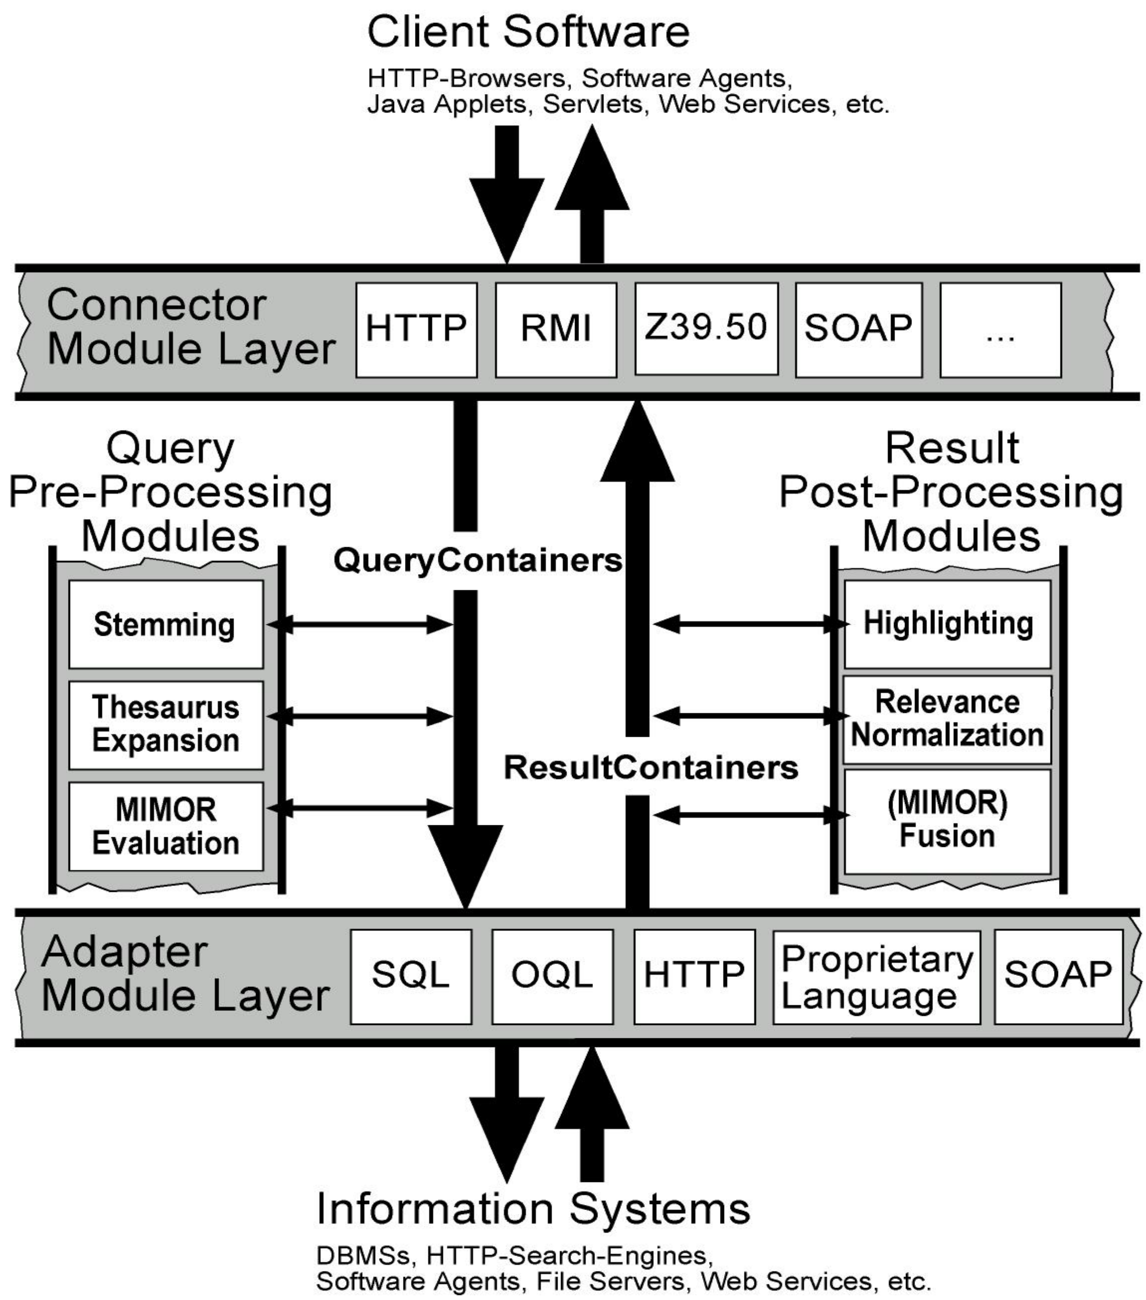

Fig. 2: Three-tier model and inner structure of RECOIN (cf. Kassem 2000:7, Scheufen 2002)

3.1 RECOIN in the Retrieval Process

In order to identify components that are candidates for working with RECOIN, the retrieval process was divided into different stages. These could be associated with groups of components connected to RECOIN:

- *Connector* objects handle the client communication. Receiving information needs and creating query objects represent the first stage of the retrieval process while formatting results and returning them to the clients can be seen as the last stage of the process.

- *Adapter* objects represent interfaces to IR systems. They accept query objects to be processed by the retrieval systems (matching) and generate result lists (rankings).
- Prior to matching queries to information objects inside the IR systems autonomous components can provide basic operations like stop word removal, stemming, relevance feedback, etc.
- Similarly, the generated rankings can be processed in a post-matching stage. Examples for operations conducted at this stage include re-ranking, highlighting, fusion of result lists, etc.

This modular approach permits to create a chain of components implementing basic retrieval services. The components and objects used in the retrieval process need to be described by meta characteristics that can be correlated in the learning component of the model. Generally a collection of meta data of the individual components in RECOIN is considered useful in several ways:

- Metadata can hold information the components need to function properly. This could be, e.g. security and authentication data like passwords and user names or simply the name of a database driver.
- Vital data for RECOIN can also be stored as an attribute. An example is the name of the Java™ class that represents a customized component. Through the class name the component can be loaded at runtime, which forms part of the plug-in mechanism.
- Apart from the afore mentioned possibilities a component can be marked with metadata concerning aspects of information retrieval. Databases, e.g., could be characterized by their content (e.g. text only, facts, images, or mixed), the retrieval model, or the ranking algorithm.

RECOIN has been implemented with special regard to MIMOR. However, it can be applied to other retrieval systems as well as the following description will show. The different components involved in the retrieval process are represented in RECOIN as pluggable modules. Since the Java™ class through which the components' functions are accessed has to be tailored to the individual component, the module object in RECOIN serves merely as a wrapper for the plug-in functions and as an interface to pass data to the specialized class inside the module object. Any data that is passed through RECOIN is encapsulated in containers serving as a common medium of transportation. Thus, RECOIN is not concerned with the format and structure of the data, but guarantees that heterogeneous components are able to send data to one another. How the data is interpreted and processed is up to the components (for details see Scheufen 2002).

RECOIN provides the basic infrastructure together with a number of abstract classes that re-present objects related to the information retrieval process. These can be extended and customized to fulfil individual requirements. The most important of these classes are *Query*, *Result* and *ModuleComponent*. The latter one represents any component to be integrated into the retrieval process. This may be a database, a thesaurus, or any other component.

3.2 RECOIN and MIMOR

RECOIN can be easily incorporated into the MIMOR model. It was already mentioned that the metadata repository can provide the necessary information about all objects that are part of the retrieval process.

It is therefore possible, on the one hand, to integrate MIMOR as a query-pre-processing module that examines the information request, i.e. the query, and selects the components (IR systems, linguistic operations, etc.) to be used in the retrieval process according to its experience accumulated in the learning component. This scenario serves as an example for a dynamically created RetrievalChain object that is generated by the MIMOR module at runtime. Additionally, MIMOR can work together with a relevance feedback component in order to make use of the feedback from a preceding retrieval cycle so as to draw conclusions about the performance of the individual components in that cycle.

On the other hand, a MIMOR module can be used as a result-post-processing module that is responsible for the fusion of individual result lists. This step also involves the object relations which are stored as weights in the learning component of the MIMOR model. The module can deduce from these weights whether the results of a special component, i.e. IR system, should be ranked higher, because the system has proven to return high quality results in the training phase of the model.

RECOIN is currently still in a testing phase and further evaluation in real life scenarios is necessary. The software has been made publicly available[5] under the GPL (General Public License) on the Sourceforge[6] platform.

[5] The Retrieval Component Integrator Project (Recoin) Homepage (2004). http://recoin.sourceforge.net [Access September 2004].
[6] Open Source Development Network, Inc (2004). Sourceforge Homepage. http://www.sourceforge.net [Access September 2004].

4 Evaluation

Information retrieval is a highly experimental discipline in which systems are evaluated with the goal of system optimisation and scientific progress. Following this tradition, MIMOR was evaluated at different development stages. We participated in the CLEF 2002 and 2003 campaigns (Kluck & Gey 2000, Peters et al. 2003) with different tasks. Only the CLEF 2003 experiments are presented in this paper. The results of MIMOR's participation in the domain-specific track are reported in Hackl et al. 2003.

4.1 Cross-Language Information Retrieval

In CLEF 2003, a fully automatic MIMOR system was applied to cross-language Information retrieval with the four languages English, French, German and Spanish using English as source language because most of the web based translation services offer translations to and/or from English. The employed tools included Lucene 1.31[7], MySQL 4.0.122[8] and JAVA TM-based snowball3[9] analyzers.

In a first step after formal pre-processing, customized snowball stemmers were used to stem the data and to eliminate stop words. Then the collection was indexed by Lucene and MySQL and the topics were translated into French, German and Spanish via machine translation tools (FreeTranslation, Reverso and Linguatec5[10]). During an analysis of the various translation tools it became apparent that the quality of the machine translations was not satisfying, but that, at the same time, the individual translation systems did not show the same weaknesses nor make the same mistakes (cf. Plödt 2003). Due to this fact, fusion of various systems was applied as well in the context of multilinguality. According to the respective language the topics were also stemmed and then merged together to form an entire query. The experiments on the training data strongly favoured Lucene as retrieval engine in comparison to MySQL which was reflected in the defaults of weight setting.

---

[7] The Apache Jakarta Project (2004). Jakarta Lucene Homepage. http://jakarta.apache.org/lucene/docs/index.html [Access September 2004].

[8] MySQL AB (2004). MySQL Homepage. http://www.mysql.com/ [Access September 2004].

[9] The Apache Jakarta Project (2002). Snowball Stemmers for Lucene Homepage. http://jakarta.apache.org/lucene/docs/lucene-sandbox/snowball/ [Access September 2004].

[10] SDL International (2004). FreeTranslation Homepage. http://www.freetranslation. com/ [Access September 2004], Reverso: http://www.reverso.net/, Linguatec GmbH (2004). E-Translation Server. http://www.linguatec.net/online/ ptwebtext/index.shtml [Access September 2004].

In order to further improve retrieval quality, blind relevance feedback (BRF) was implemented. Expansion terms were selected by applying the Robertson selection value or the Kullback-Leibler (KL) divergence measure (Carpineto et al. 2001). Thus, the submitted runs used BRF KL from the top five documents adding 20 terms.

In our test runs we were able to show that fusion helped to improve at least the recall, although the official CLEF results for 2003 did not confirm this finding. The Lucene-based runs generally outperformed the fusion runs, except for a marginally better recall in the merged monolingual run:

Run	**Documents retrieved**	**Average precision**
UHImlt4R1	3944 / 6145	0.285
UHImlt4R2	4137	0.306
UHImnenR1	951 / 1006	0.363
UHImnenR2	945	0.380

Table 1: CLEF Results 2003

In order to gain more insights a number of additional experiments was conducted beyond CLEF 2003. On the one hand the isolated IR systems were examined and it could be proved that by applying more intensive optimisation techniques the solo performance of each individual system could be improved (Hackl et al. 2004). On the other hand it turned out that BRF worked generally well but using the original (perfect) translations of the CLEF queries instead of the automatically translated ones for all four languages. BRF as well as the fusion runs had a negative influence on the overall performance. More detailed experiments are needed to ensure these findings.

4.2 Further Experiments on IBM's DB2 Text Extender

In a further empirical step, MIMOR was tested with a commercially available retrieval software, in this case IBM's DB2 text extender. Details of this evaluation are described in Li 2002, a master thesis at the University of Hildesheim. One part of the German CLEF corpus ("*Der Spiegel*") and all 30 CLEF topics from the 2000 campaign were used for these experiments. Text Extender allows many parameter settings mainly based on different linguistic processing modules. Some of these settings were used to establish different systems for the fusion experiment including Boolean retrieval models as well as probabilistic ones. Linguistic pre-processing like stemming and statistical n-gram-methods are comprised.

It turned out that MIMOR's fusion strategy is indeed a fruitful strategy which can increase the quality of retrieval results. The results showed that fusion works well for the Boolean and for the probabilistic model. However, the quality of the Boolean runs is lower overall.

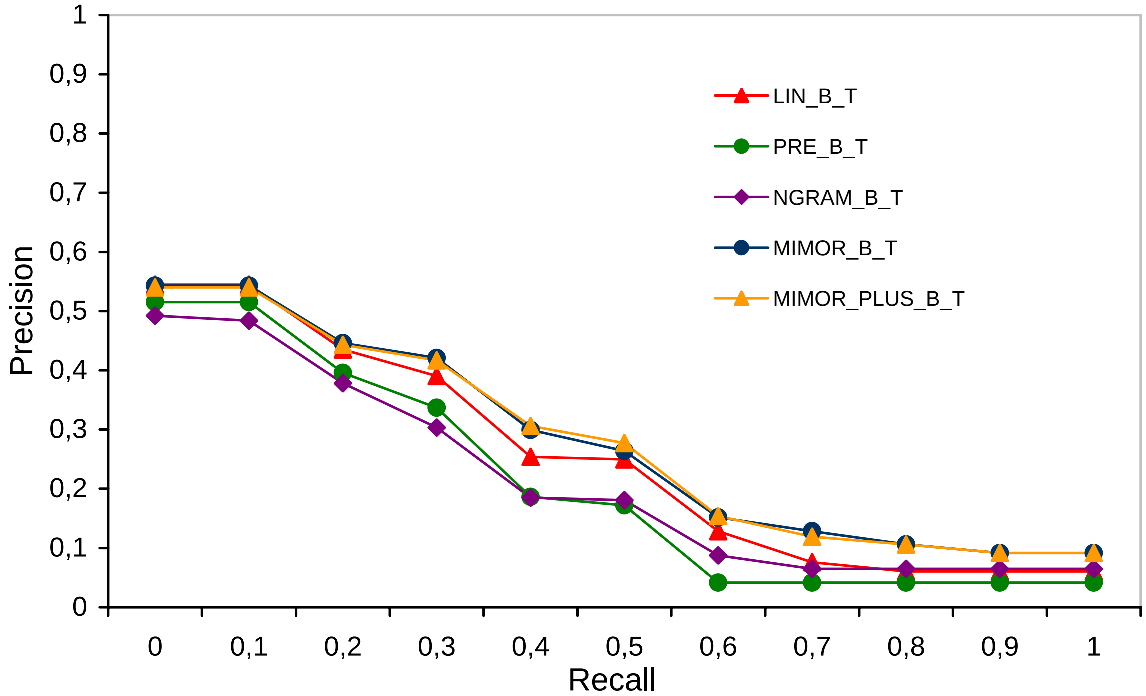

Fig. 4: Recall-Precision-Graph for Boolean DB2 runs with DER SPIEGEL

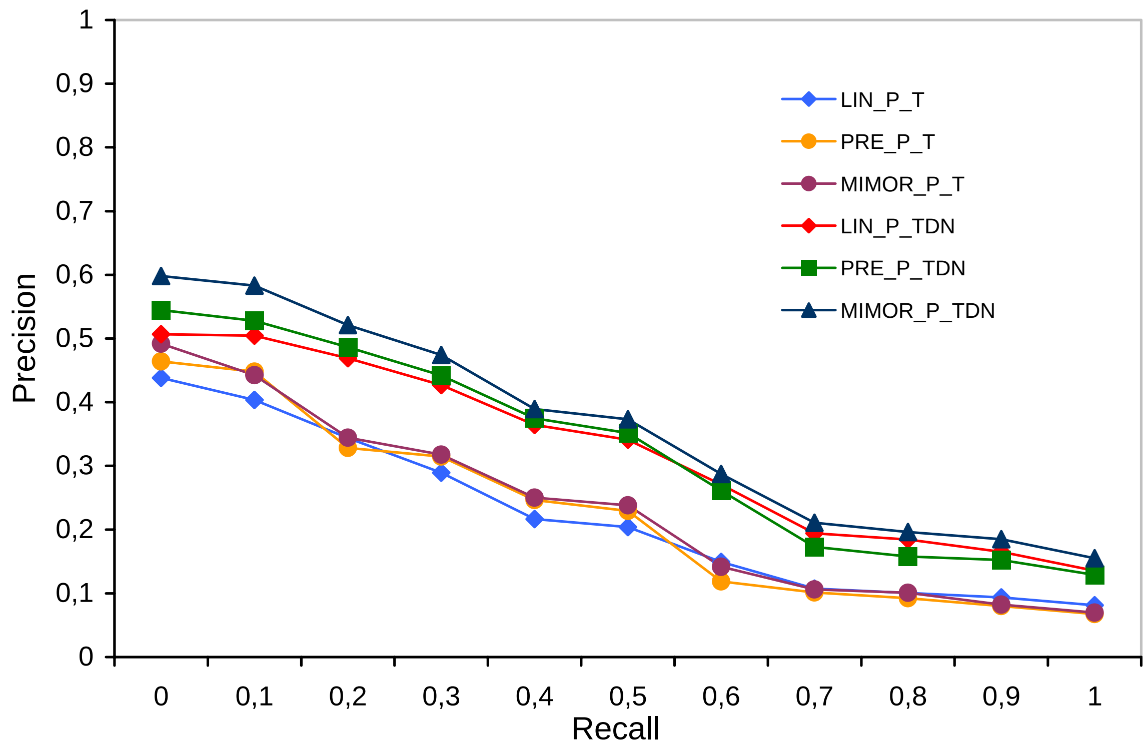

Fig. 5: Recall-Precision-Graph for Probabilistic DB2 runs with DER SPIEGEL

4.3 Interpretation of the Results

The results gained from CLEF within cross-language IR and further experiments generally provide evidence for the positive impact of fusion at certain points of the IR process. On the other hand, different combinations of IR systems, models, languages, techniques etc. lead to different results. This fact encourages the combination of fusion approaches with machine learning. According to the results, fused systems or resources should not be added together without reflection but a controlling mechanism should be applied which registers which component good results are due to and vice-versa.

5 Outlook

This paper gives a summary of MIMOR showing its evolution from the model to the system over time. The open design of MIMOR turned out to be a well reflected decision and led to interesting findings with respect to IR at a practical level. The evaluations in which MIMOR participated showed that context plays a very important role and accepted theoretical statements had to be revised by empirical facts. The learning facility of MIMOR which is a very important feature could not be tested yet with CLEF's evaluation environment. Only a long-term study could give insights to the performance of this device.

6 References

Baeza-Yates, R.; Ribeiro-Neto, B. (eds.) (1999). Modern Information Retrieval. Harlow et al. Addison-Wesley.

Carpineto, C., de Mori, R., Romano, G. und Bigi, B. (2001). "An Information-Theoretic Approach to Automatic Query Expansion." In: ACM Transactions on Information Systems. 19(1), 1-27.

Fingerhut, M. (ed.) (2002). ISMIR 2002: 3rd Annual International Symposium on Music Information Retrieval. 13.-17. October 2002. IRCAM Centre Pompidou, Paris.

Fox, E.; Shaw, J. (1994). "Combination of Multiple Searches." In: Harman D. (ed.)(1994). The Second Text Retrieval Conference. 1994. NIST Special Publication, 500-215.

Hackl, R. (2003). Mehrsprachiges Information Retrieval im Rahmen von CLEF 2003. Master Thesis Information Science. University of Hildesheim, Germany.

Hackl, R.; Kölle, R.; Mandl, Th.; Womser-Hacker, Ch. (2003). "Domain Specific Retrieval Experiments at the University of Hildesheim with the MIMOR System." In: Peters, Carol; Braschler, Martin; et al. (ed.) (2003). Advances in Cross-Language Information Retrieval: Third Workshop of the Cross-Language Evaluation Forum, CLEF 2002, Rome, Italy, September 2002. Berlin et al.: Springer [Lecture Notes in Computer Science 2785] 343-48.

Hackl, R., Kölle, R., Mandl, Th., et al. (2004). "Multilingual Retrieval Experiments with MIMOR at the University of Hildesheim." In: Peters, C., Braschler, M., et al. (eds.) (2004). Evaluation of Cross-Language Information Retrieval Systems. Proceedings of the CLEF 2003 Workshop. Berlin et al.: Springer [LNCS], to appear.

Harman, D. (1992). Relevance feedback revisited. Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval. Copenhagen, Denmark, 1-10.

Kassem, N. et al. (2000). Designing Enterprise Applications with the Java™ 2 Platform, Enterprise Edition. Harlow et al.: Addison-Wesley.

Kluck, M., Gey, F. (2000). "The Domain-Specific Task of CLEF – Specific Evaluation Strategies in Cross-Language Information Retrieval." In: Peters, Carol (ed.) (2000). Cross-Language Information Retrieval and Evaluation. Workshop of the Cross-Language Information Evaluation Forum (CLEF) Lecture Notes in Computer Science, Vol. 2069, Berlin Heidelberg New York: Springer-Verlag, 48-56.

Kwok, K.L., Chan, M. (1998). "Improving Two-Stage Ad-Hoc Retrieval for Short Queries." In: Proc 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '98). Melbourne, Australia, 250 – 56.

Li, Q. (2002). Anbindung von DB2 an die Implementierung des Fusionsansatzes MIMOR. Master Thesis Information Science. University of Hildesheim, Germany.

Mandl, T., Womser-Hacker, C. (2001). "Probability Based Clustering for Document and User Properties." In: Ojala, T.(ed.) (2001). Infotech Oulo International Workshop on Information Retrieval (IR 2001). Oulo, Finland. 2001, 100-07.

Mandl, T. & Womser-Hacker, C. (2003). "Learning to Cope with Diversity in Music Retrieval." Journal of New Music Research 31,1., 133-42.

McCabe, M.C., Chowdhury, A., Grossmann, D. et al. (1999). "Unified Framework for Fusion of Information Retrieval Approaches." In: Proc. Eighth ACM Conference on Information and Knowledge Management (CIKM). Kansas City, Missouri, USA, 330-34.

Peters, C., Braschler, M., Gonzalo, J., et al. (2001) (eds.). Evaluation of Cross-Language Information Retrieval Systems. Proceedings of the CLEF 2001 Workshop. Berlin et al.: Springer [LNCS Vol. 2406].

Peters, C., Braschler, M., Gonzalo, J., et al. (2003) (eds.). Advances in cross-language information retrieval. Proceedings of the CLEF 2002 Workshop. Berlin et al.: Springer [LNCS Vol. 2785].

Plödt, A. (2003). Multilinguales Information Retrieval: Verfahren zur Anfrageoptimierung im Kontext des Cross-Language Evaluation-Forum (CLEF). Master Thesis Information Science. University of Hildesheim, Germany.

Savoy, J. (2002). "Report on CLEF-2002 Experiments: Combining Multiple Sources of Evidence." In: Peters, C. (ed.) (2002). Cross Language Evaluation Forum: Results of the CLEF 2002 Evaluation Campaign. Working Notes CLEF-Workshop. Rom, 19.-20.9.2002, 66-90.

Scheufen, J.-H. (2002). Modularisierung des Retrievalprozesses zur funktionellen Integration heterogener IR-Komponenten. Master Thesis Information Science. University of Hildesheim, Germany.

Vogt, C., Cottrell, G. (1998). "Predicting the Performance of Linearly Combined IR Systems." In: Proc 21st Annual International ACM SIGIR Conf on Research and Development in Information Retrieval (SIGIR '98). Melbourne, Australia, 190-96.

Voorhees, E. M. (1998). Variations in relevance judgements and the measurement of retrieval effectiveness. Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval. Melbourne, Australia, 315 – 23.

Voorhees, E., Harman, D. (2001) (eds.). The Ninth Text Retrieval Conference (TREC-9). Gaithersburg, Maryland, USA: NIST Special Publication. National Institute of Standards and Technology. Nov. 2001.
http://trec.nist.gov/pubs/ trec9/t9_proceedings.html [Access September 2004].

Witten, I., Frank, E. (2000). Data Mining: Practical Machine Learning Tools and Techniques with JAVA Implementations. San Francisco: Morgan Kaufman.

Wolff, C. (2000). "Effektivität von Recherchen im WWW: Vergleichende Evaluierung von Such- und Metasuchmaschinen." In: Knorz, G., Kuhlen, R. (eds.) (2000). Informationskompetenz – Basiskompetenz in der Informationsgesellschaft. Proceedings 7. Intl. Symposium für Informationswissenschaft. (ISI 2000). Darmstadt, Nov. 8.-10. 2000. Konstanz: Universitätsverlag [Schriften zur Informationswissenschaft Bd. 38], 31-48.

Womser-Hacker, C. (1997). Das MIMOR-Modell. Mehrfachindexierung zur dynamischen Methoden-Objekt-Relationierung im Information Retrieval. Universität Regensburg, Informationswissenschaft, Habilitationsschrift.