

CAREER: Multimedia Analysis and Retrieval System (MARS)

Sharad Mehrotra

Department of Information and Computer Science, University of California at Irvine

Contact Information

Sharad Mehrotra, 424 Computer Science, University of California at Irvine, Irvine, CA 92697-3425

Phone: (949) 824 5975 Fax: (949) 824 4012 Email: sharad@ics.uci.edu URL: <http://www.ics.uci.edu/~sharad>

Project URL: <http://www-mars.ics.uci.edu/>

Project Award Information

- Award Number: IIS: 9734300
- Duration: 4 years, (June 1st 1998 – May 30th 2002)
- Current Year: 3rd year
- Title of the project : CAREER: Multimedia Analysis and Retrieval System

Keywords

Multimedia analysis, Multimedia information retrieval, relevance feedback, query refinement, multidimensional indexing, uncertainty in databases, dimensionality curse, query processing.

Project Summary

The goals of the MARS project are to design and develop an integrated multimedia information retrieval and database management infrastructure, entitled *Multimedia Analysis and Retrieval System* (MARS), that supports multimedia information as first-class objects suited for storage and retrieval based on their content. Specifically, research in the MARS project is categorized into the following four sub-areas each of which contribute to the development of the integrated infrastructure. (1) *Multimedia Content Representation*: extraction of multimedia content and content-based representation of multimedia objects in databases. (2) *Multimedia Information Retrieval*: content-based multimedia retrieval techniques including multimedia retrieval models and interactive query refinement techniques. (3) *Multimedia Feature Indexing*: that overcomes the high-dimensionality and non-Euclidean nature of feature data to efficiently support retrieval based on feature similarity. (4) *Multimedia Database Management*: techniques to effectively and efficiently incorporate content-based retrieval of multimedia information into structured database processing.

Goals, Objectives, and Targeted Activities

In the first year, our primary objective was (1) to explore techniques for multimedia object and query representation, and multimedia information retrieval, and (2) to develop mechanisms for efficient indexing of highly multidimensional feature spaces. In the second year we continued to explore techniques for multimedia object and query representation and multimedia information retrieval with integrated models that can span multiple media types for integrated retrieval and enhanced interactive browsing. We also progressed in integrating indexing techniques with multimedia retrieval by enhancing them with support for relevance feedback. In the third year, our primary goals are exploring further extensions to the data model and to develop query processing techniques to support content-based retrieval within database management systems. The overall goal of the project is to develop an integrated database management and information retrieval system that provides native support for content-based retrieval over multimedia data sets.

Publications and Products

The MARS project to date has resulted in 38 publications in journals and conferences. Summary of our achievements is described below.

- *Video Representation*: Motivated with the role of a table of content (ToC) used in accessing a book, we developed algorithms to structure a video into a set of scenes which represents a video ToC. This ToC can be used to support effective user access (browsing and retrieval) of video.
- *Indexing Highly Multidimensional Spaces*: We developed a hybrid tree data structure to support high-dimensional feature indexing in multimedia databases. The hybrid tree combines the positive aspects of bounding

region (BR)-based data structures (e.g., R-tree, SS-tree) and space partitioning (SP) based approaches (e.g., KDB-tree, hB-tree) into a single search structure to achieve superior performance and scalability to high dimensionalities than either of the two approaches. We also proposed a technique for local dimensionality reduction (LDR) as an approach to high dimensional indexing which exploits local correlation in data to increase indexing efficiency.

- **Multimedia Information Retrieval:** We made significant progress on multimedia information retrieval specially on adapting relevance feedback as an approach to learning the user's information need in image databases. With relevance feedback, the user is relieved of the burden to state their exact information need based on the feature sets supported in the system. Instead, the system learns both the query representation, and the distance/similarity measures that a user has in mind. We made progress on multiple-media-type integrated retrieval, which with aid from relevance feedback allows for a new modality of retrieval more akin to power browsing, using which the user can use search results based on one media type as browsing access points to another type of media iteratively.
- **Supporting Multidimensional Data Structures as Access Methods:** We have developed the first granular locking approach to concurrency control in multidimensional data structures. Our solution provides a scalable solution to integrating multidimensional data structures into DBMSs as an access method.
- **Supporting a Content-based similarity matching model in databases:** We developed a similarity based matching technique that supports content-based retrieval in databases and also enables relevance feedback.

Publications since last IDM workshop

1. Kaushik Chakrabarti, Michael Ortega-Binderberger, Kriengkrai Porkaew, and Sharad Mehrotra, "Similar Shape Retrieval in MARS", Proceeding of IEEE International Conference on Multimedia and Expo, New York City, NY, USA, July 30 - August 2, 2000.
2. Kaushik Chakrabarti and Sharad Mehrotra, "Local Dimensionality Reduction: A New Approach to Indexing High Dimensional Spaces", VLDB, Cairo, Egypt. September 10-14, 2000.
3. Kaushik Chakrabarti, Kriengkrai Porkaew, and Sharad Mehrotra, "Efficient Query Refinement in Multimedia Databases", Proceedings of the 2000 IEEE International Conference on Data Engineering (ICDE), Feb., 2000.
4. Michael Ortega, Sharad Mehrotra, Kaushik Chakrabarti, and Kriengkrai Porkaew, "WebMARS: A Multimedia Search Engine", Appeared in Proc. IS&T/SPIE's Annual Symposium of Electronic Imaging, January, 2000.
5. Eamonn Keogh, Kaushik Chakrabarti, Michael Pazzani and Sharad Mehrotra "Locally Adaptive Dimensionality Reduction for Indexing Large Time Series Databases", 2001 ACM SIGMOD Conference on Management of Data, May, 2001.
6. Kaushik Chakrabarti, Minos Garofalakis, Rajeev Rastogi and Kyuseok Shim "Approximate Query Answering Using Wavelets", VLDB Journal (Special Issue, Best Papers of VLDB 2000), No. 3, 2001. September, 2001.
7. Eamonn Keogh, Kaushik Chakrabarti, Michael Pazzani and Sharad Mehrotra "Dimensionality Reduction for Fast Similarity Search in Large Time Series Databases", Knowledge and Information Systems Journal (to appear).
8. Kaushik Chakrabarti, Minos Garofalakis, Rajeev Rastogi and Kyuseok Shim, "Approximate Query Answering Using Wavelets", Appeared in Proceedings of the VLDB Conference, Cairo, Egypt. September 10-14, 2000.
9. Kaushik Chakrabarti, Kriengkrai Porkaew, Michael Ortega and Sharad Mehrotra, "Evaluating Refined Queries in Top-*k* Retrieval Systems", Submitted for publication.
10. Michael Ortega-Binderberger, Kaushik Chakrabarti, and Sharad Mehrotra, "An Approach to Integrating Query Refinement in SQL", Submitted for publication.
11. Michael Ortega-Binderberger, Sharad Mehrotra, Kaushik Chakrabarti, and Kriengkrai Porkaew, "WEBMARS: A Multimedia Search Engine for Full Document Retrieval and Cross Media Browsing", Multimedia Information Systems Workshop (MIS2000), Chicago, IL, 2000

Project Impact

- This award will support in part two graduate students (Kaushik Chakrabarti and Michael Ortega-Binderberger) during the span of their PhD. The project will involve undergraduate students in building and testing various components of the MARS system. Many undergraduate and graduate students taking courses on database management and information retrieval at UCI will use the MARS prototype in their class projects.
- The PI (based on preliminary work of the MARS project) has successfully acquired grants from NASA in collaboration with Rockwell Science Center to utilize some of the indexing techniques developed to support efficient query processing in spatio-temporal data bases consisting of weather events. The objective of the project is to support a weather advisory system to help general aviation.

- Based on the MARS project, the PI has acquired software grants from Informix Corp., Oracle Corp. and Magic Software. This provides a software infrastructure for the database related research and teaching at UCI.
- The PI is in the process of filing for a patent based on indexing techniques developed as part of the MARS project.

Project References

Project web site <http://www-mars.ics.uci.edu>

GPRA Outcome Goals

Our research in the first and second years of the grant has lead to the following important discoveries:

- **Techniques to extract Table of Content from video** that group videos into semantically related scenes thereby providing a unit of representation and retrieval of video data.
- **Scalable approach to concurrency control in multidimensional data structures.** Absence of such protocols is one of the reasons why existing DBMS vendors have not incorporated these data structures. Our work attempts to overcome this limitation.
- **Hybrid tree:** a data structure that scales to very highly multidimensional feature spaces. We extended the Hybrid tree data structure to support to multi-point queries, which improves the performance of relevance feedback techniques by several orders of magnitude over the naive approach.
- **Techniques to enhance multidimensional indexing:** We proposed a technique called local dimensionality reduction (LDR) as an approach to high dimensional indexing. LDR exploits the local correlation in the data to increase indexing efficiency. Our approach improves the performance of the index structure by almost a degree of magnitude over the global dimensionality reduction (Principal Component Analysis) and original space indexing techniques.
- **Techniques for query refinement using relevance feedback for multimedia retrieval:** A relevance feedback architecture infrastructure within which several query refinement techniques were developed. Such techniques are query point movement and query expansion for feature level and inter-feature re-weighting techniques for inter-feature level. Since our work, many other researchers have also adopted and explored similar approaches.

Area Background

Many emerging applications such as multimedia digital libraries, medical image databases require accessing multimedia information based on their content. Traditional text based information retrieval techniques do not suffice due to the complexity of capturing visual content using textual description and the high degree of human effort in developing the suitable descriptions. An alternative approach is to describe images/videos based on visual properties (such as color, texture, and shape) and support retrieval based on these features. Such an approach to visual information retrieval requires interdisciplinary research in the areas of: 1) *image processing* and *computer vision* to extract salient features from the visual media that capture its content, 2) *information retrieval* to define a notion of similarity between multimedia objects, and 3) *database management* to support efficient storage of multimedia objects and retrieval based on the degree of match. Our research addresses the information retrieval and database management challenges in content-based multimedia retrieval.

Area References

1. C. Faloutsos, *Searching multimedia databases by content*, Kluwer Academic Publishers, Norwell, MA, 1996.
2. R. Jain, A. Pentland, D. Petkovic, *NSF ARPA workshop on visual information management systems*. Cambridge, MA, June 1995.
3. IBM, Query By Image Content at: <http://wwwqbic.almaden.ibm.com>
4. Robert Korfhage, *Information Storage and Retrieval*, John Wiley and Sons, 1996
5. C. Zaniolo, et. al., *Advanced Database Systems*, Morgan Kauffman Publishers, 1997

Potential Related Projects

There are a large number of projects that are potentially related to MARS since MARS cuts across the image and multimedia processing, information retrieval and database management disciplines. Examples of other related projects funded by the IDM program include the work by: Prof. R. Jain (UC, San Diego), Prof. S.F Chang (Columbia), Prof. Sistla, C. Yu., Ben Arie, and O. Wolfson (UIC), Prof. Hellerstien (UC, Berkeley), Prof. Ramakrishnan (Wisconsin), Prof. Ozsoyoglu (Case Western), Prof. Manjunath (UCSB), Prof. W. Chu (UCLA), and Prof. Zhang (SUNY, Buffalo).