



The 2nd International Conference on Integrated Information

CrossMedia: supporting collaborative research of media retrieval

Péter Mátételki, László Havasi, Márton Gergó, András Micsik, Ákos Kiss, Tamás Szirányi, László Kovács *

*MTA SZTAKI, Computer and Automation Research Institute of the Hungarian Academy of Sciences
Lágymányosi 11, Budapest 1111, Hungary*

Abstract

The goal of our new e-science platform is to support collaborative research communities by providing a simple solution to jointly develop semantic- and media search algorithms on common and challenging datasets processed by novel feature extractors. Querying of nearest neighbor (NN) elements on large data collections is an important task for several information or content retrieval tasks. In the paper a flexible framework for research purposes is introduced for testing features, metrics, distances and indexing structures. The core part of the content based retrieval system is the LHI-tree, a disk-based index scheme for fast retrieval of multimodal features. Additionally, we compare LHI-tree to FLANN, an effective implementation of ANN search and show that LHI-tree gives similar list of retrieved images.

© 2013 The Authors. Published by Elsevier Ltd.
Selection and peer-review under responsibility of The 2nd International Conference on Integrated Information

Keywords: image retrieval, e-science, research collaboration

* Corresponding author: András Micsik Tel.: +36 1 279 6248; fax: +36 1 279 6200.
E-mail address: andras.micsik@sztaki.mta.hu

1. Introduction

Multimedia information systems are becoming increasingly important with the advent of multi-sensor networks, mobile phone data capture and increasing number of multimedia databases. Since visual, auditory media and the adherent information require large amounts of memory and computing power for storage and processing, there is a need to efficiently index, store, and retrieve the visual information from multimedia/cross-media databases [1]. Experiments based on the above system are limited by several real life scenarios:

- E How to process the extremely increased information quantity? The sophisticated processing and index building methods need significantly more time than real-time. The retrieval system will not be able to follow the incoming data flow.
- E How to update a database with data from different modalities? Real life sensor networks or multimedia contents built for several modalities need various feature extraction and indexing methods.
- E How to insert novel features? Several features are based on heuristics, thus there is no fixed dimensional (vectorial) form.

CrossMedia portal presents an all-in-one solution for such problems: storage and processing capacity, flexible interfaces, built in index structure and innovative user interfaces. Users form communities on the CrossMedia portal where they can jointly create, build and share search algorithms and media datasets in an iterative way. The system automatically builds indexes and also provides a testing facility as indexes are instantly included in the portal's search interface to be tested with any desired media-based, semantic, or combined multimodal input. The generated indexes can be shared among research communities for reviewing or with the public for demonstration purposes.

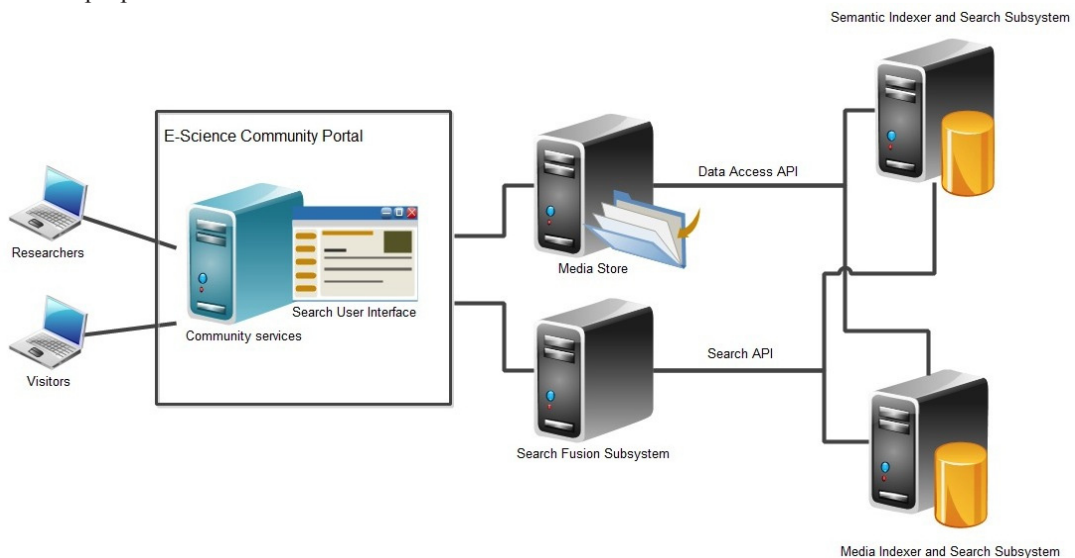


Fig. 1. Architecture of the CrossMedia e-science platform

2. CrossMedia e-science platform

The CrossMedia platform (see Figure 1) addresses many diverse tasks made accessible for the users through the portal that serves as an access point for all available services. Behind the portal we set up a distributed system consisting of multiple processing units organized in a loosely coupled service-oriented architecture. The system consists of the following subsystems:

- E The Media Store (MS) is responsible for safekeeping all searchable multimedia elements.
- E The Media Indexer and Search Subsystem (MISS) is responsible for generating index trees for a specific algorithm on a specific media set in the MS and it is also capable of executing similarity-based search on the generated index trees for any input media.
- E The Semantic Indexer and Search Subsystem (SISS) is responsible for creating semantic databases and indexes by processing media annotations of the MS entities and it is capable of executing semantic search on the built datasets.
- E The Search Fusion Subsystem (SFS) is responsible for combining the results of the MISS and SISS in case of multi-input multimodal search expressions.
- E The Search User Interface (SUI) enables users to easily create complex multimodal search expressions and to see and evaluate search results.
- E The E-Science Community Portal (ECP) is responsible for integrating and providing all the above functionalities through a stylish modern web interface and for enabling users to create and jointly work in research communities and support the collaborative work.

2.1. Managing datasets

In this area typically large sample datasets are needed for testing of search algorithms – for example we used a database containing about 5 million entities when testing our custom developed image search algorithms. As the platform aims to enable its users to collaboratively build and maintain these large datasets we had to find a sustainable solution. After evaluating many currently available solutions we realized that the shared data and permission management is far from trivial and we are in need of defining and developing a custom architecture to suit our special needs.

Our system architecture separates the portal's community management and the data management into detached, loosely coupled components: the E-Science Community Portal is in charge of all group- and community management tasks, and the Media Store is a simple but very responsive data storage, manager and retrieval component. This separation keeps the community-generated large datasets and the community management independent thus enabling us to flexibly redesign the frontend system component if necessary. The solution requires considerably less resources on the portal side as representing the large datasets only needs a low quantity of data to be synchronized, so the portal remains scalable.

Apart from simple data storage the Media Store component is the one connecting community generated (portal) content to the media processing components: the Semantic- and Media Indexer and Search Subsystems. The communication is achieved using REST API.

2.2. Protocol for instant multimodal search

The CrossMedia Search User Interface component can handle search inputs of different (semantic and content-based) modalities even in a single query expression as we use a uniform query input representation

consisting of the search index (either semantic or content-based), the search target (text data or binary media) and a weight parameter. This generic search-input definition not only enables to easily expand the search modality coverage in the future but also supports the composition of multi-input and multimodal search queries. The user-generated complex multimodal search requests are handled by the Search Fusion Subsystem. It evaluates the search expressions and divides them into single terms, forwards the individual terms to the appropriate search subsystem and aggregates the separate results into one unified result list that is sent to the Search User Interface as the final response. The protocol and also the SFS are able to generate partial (not final) results, so the user gets an instant search experience: when adding search inputs the preliminary result list is immediately displayed and it is extended or refined gradually when more precise results are available. As soon as the user composes a more complex search expression or refines an existing expression the result list immediately reflects the current state of the search parameters.

3. Disk based indexing

When the number of dimensions of a feature set increases dramatically, the methods used for low-dimensional indexing are not applicable any more. The classical kd-tree algorithm [2] is efficient in low dimensions, but its performance degrades by higher.

[3] uses a multiple randomized kd-tree, where it splits the data for a randomly chosen set of dimensions (e.g. 5) instead of that of the greatest variance. In [4] this random kd-tree is applied by using multiple trees, priority search on hierarchical k-means trees, and automatic parameter setting. They have demonstrated that this configuration of algorithms can speed the matching of high-dimensional vectors up by several orders of magnitude compared to linear search.

A recent method proposed in [5] is the Nearest Vector Tree which is designed for approximate nearest neighbor search in very large, high-dimensional databases. It transforms the high dimensionality search task into an efficient one dimensional space based on the combination of projections of data points to lines and the partitioning of the projected space.

The LHI-tree is similar to M-index where base points (so called pivots) are chosen randomly to reduce the high-dimensional feature vectors. A modification of this random selection is applied, where a quasi orthogonality criteria is forced during random point selection. Beyond the point selection we estimate basic statistical properties of input space from the representative sample set.

In contrast to permutation-based schemes, LHI-tree uses base points to compute reference distances for every input vector from the quantized distances. It is carried out by using AVL-trees inside the LHI-tree, connected to every base point. Input of AVL-trees are the distances of the input image to base points, while the outputs are the number of bins in which the quantized distances fell. Visually, LHI-tree contains base points as hyper-sphere centers in feature space. The indices from AVL-trees are the shells of such spheres with different radius. To assign a disk partition to a part of the feature space, we have used hashing function of quantized distances. This hash function guarantees that the near vectors are placed into the same disk partition (file). During the evaluation process two datasets were used:

- E CoPhIR: color structure, edge histogram and homogenous texture descriptors were selected. The applied distance definitions are the Euclidean [4] and pattern difference [6].
- E Caltech 101: SIFT vectors. The applied distance definition is the Euclidean.

Table 1 summarizes the performance parameters of some building and retrieval trials.

Table 1. Run-time parameters of the LHI-tree index. Number of base spheres/number of input points/dimension of input vector

	3sphs/300k 64dims	5sphs/600k 64dims	5sphs/829k 80dims	5sphs/632k 128dims
Build time (sec)	456	12194	17619	14901
RAM storage (MB)	0.912	1.248	1.432	0.945
HDD storage (MB)	80200	161000	215000	329210
Search time (msec) [†]	67	45	110	160
RAM usage during search (MB)	1.2	0.97	1.2	2.2

4. Using the portal

The CrossMedia E-Science Portal aims to satisfy the domain-specific needs of researchers primarily in the image processing scientific field. Registered users can be engaged in different research groups by creating a new society or by joining the desired groups. Available functionality in a group not only satisfies the specific needs of researchers but naturally offers other community-based collaboration facilities such as discussion forums, commenting and activity monitoring.

Every group has its own separate private space in the portal where group members share and collaboratively work with group content. Media collections (so-called albums) can be created to store sample media for the index trees to be built. Semantic and free text annotations can be attached to any media item in the albums thus enabling the Semantic Indexer and Search Subsystem to perform semantic search on these media items. Groups can also upload their developed media indexer algorithms to the portal that can later be used to create indexes. Image indexes logically consist of exactly one algorithm and one or more media collections. Defining an index on the portal interface launches a series of automated operations that produce the desired index structure in the backend and finally results in making the generated index available for building search requests on the Search User Interface.

Index generation tasks execute when creating a new index for an algorithm on a collection, when adding a new collection to an existing index or when modifying collections that are already indexed. In either case similar backend operations occur: the Media Store receives the index modification request and creates the atomic jobs that basically consist of an algorithm and a target media item to be processed by that algorithm. These jobs can be polled from the Media Store job queue and processed asynchronously by the indexer subsystems. Especially for large collections this may be a weary and relatively slow task so its results are fed back to the portal regularly to keep users updated of the current index processing state at all times. This mechanism keeps indexes up-to-date

[†] The data is stored in uncompressed format. Search times contain all additional operation such as extraction, db operations etc.

with the corresponding media collections so researchers or even anonymous users are provided with the most recent versions at all times.

Every piece of group content, including media collections, algorithms and indexes can be (separately) shared at two different levels: either with other group(s) or with the public. This creates a wide variety of possibilities for inter-group collaboration and for demonstration purposes as well, e.g. setting public visibility on an index tree makes it available for every visitor of the site. This could be used to show off the capabilities of a research group as the index (including its description that explains how the corresponding search algorithm works) becomes available for any user and can be queried with a user-selected media item. A group can also share collections with other groups so they can test their (private) algorithms and indexes on a common dataset. A trickier scenario can also be imagined in case a group shares one of its indexes but keeps the corresponding collections and algorithms private. This setup enables other groups to use the shared index to evaluate search but in the meantime it will not let other groups to use either the corresponding algorithm to index their own collections or the corresponding collection to be indexed with other algorithms.

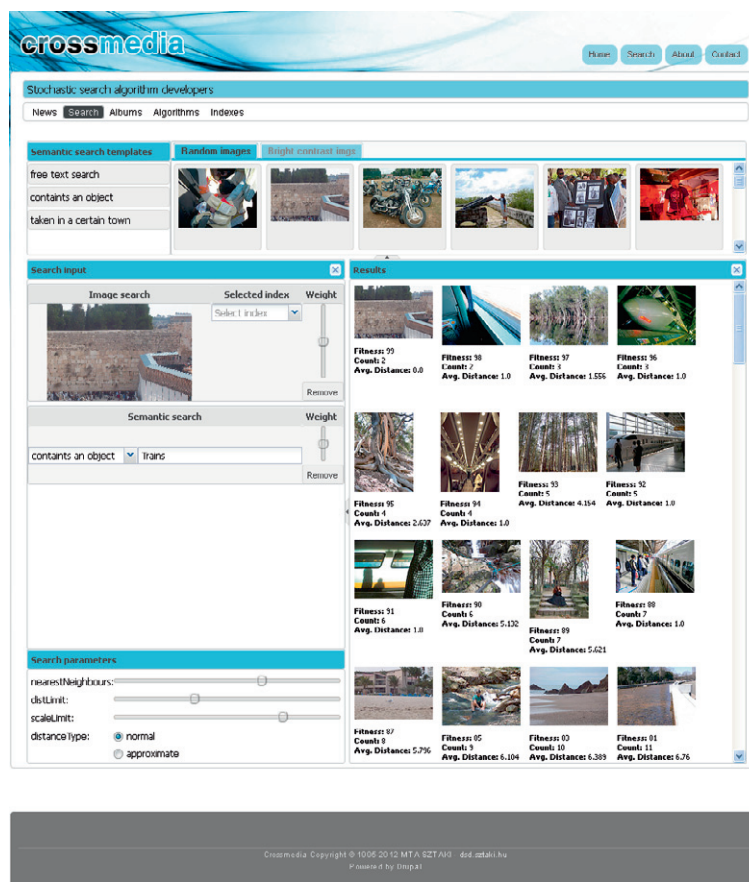


Fig. 2. Search user interface

For testing image descriptors and semantic indexes the portal provides an intuitive Search User Interface (see Figure 2) for every group and also one public interface for collections and indexes that are published site-wide. Group SUIs allow group members to assemble complex, multipart search queries. Every part of a query consists of an input media item chosen from the accessible collections, an index and a weight parameter. An index can either be content-based or annotation-based (semantic). Users have the possibility to create multi-input queries using just one (content-based or semantic) index type or they can even compose multi-input and multimodal queries where both content-based and semantic indexes are involved. The Search Fusion Subsystem generates the unified result lists using internal weighting mechanisms. Users have the ability to tune the weighting by providing positive or negative feedback on a result item regarding their opinion on moving that item forward or backward in the result list. The SUI is implemented as a platform- and browser independent web-application using the Sencha platform based on JavaScript and AJAX technologies.

5. Conclusions

We introduced CrossMedia e-Science Community Portal where different kinds of feature and distance definitions can be integrated into the disk-based indexing scheme (LHI-tree). LHI-tree is the initial indexing structure of the portal. The goal of the portal is to provide a flexible framework for research purposes on multimodal database mining and retrieval tasks. We showed that the proposed retrieval engine could achieve acceptable speed for searching with very low memory and CPU loads.

By using the CrossMedia e-Science Community Portal researchers gain the possibility to jointly work in groups and to collaborate with other research communities. Furthermore, the distributed backend infrastructure ensures the scalable and fast manipulation of indexes, while the user interface provides a testing and evaluation framework for their research. Thus, the domain-specific functionality of the portal enables them to easily produce testable results in the field of image processing and to compare or re-use each other's results.

References

- [1] Geetha, P. and V. Narayanan (2008). A Survey of Content-Based Video Retrieval. *J. Comput. Sci.*, 4: 474-486.
- [2] Friedman, J. H., Bentley, J. L., and Finkel, R. A. (1977). An Algorithm for Finding Best Matches in Logarithmic Expected Time. *ACM Transactions on Mathematical Software*, 3, 209–226.
- [3] Silpa-Anan, C. and Hartley, R. (2008). Optimised KD-trees for fast image descriptor matching. in *proc. CVPR 2008*.
- [4] Muja, M. and Lowe, D. G. (2009). Fast approximate nearest neighbors with automatic algorithm configuration. In *proc. VISAPP 2009*.
- [5] Lejsek, H., Asmundsson, F.H., Jonsson, B.T., Amsaleg, L. (2009). NV-Tree: An Efficient Disk-Based Index for Approximate Search in Very Large High-Dimensional Collections. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31(5), 869 - 883.
- [6] Eidenberger, H. (2003). Distance measures for MPEG-7-based retrieval. *ACM Multimedia Information Retrieval Workshop, Berkeley, USA, 2003*.