Integrating Content-Based Visual Access Methods into a Medical Case Database

Henning Müller, Antoine Rosset, Jean-Paul Vallée, Antoine Geissbuhler

University Hospitals of Geneva, Division of Medical Informatics

Abstract

In the computer vision domain, content-based access methods to all forms of multimedia data are a hot research topic. A large number of tools have been developed to find documents in multimedia repositories and to manage the (visual) information that has been created, for example by the Internet. Although no general breakthrough has been achieved with respect to searching diverse databases with automatically extracted features, the techniques have gained acceptance in a few well-defined domains such as image agencies (Corbis) and trademark research.
In the medical domain, the number of digital images produced and used for teaching, diagnostics and therapy is rising in a similar way as in other domains. Still, there are only a few image retrieval systems that use automatically extracted visual features for content-based access to medical image databases.
This article describes the use of an open source image retrieval system (GIFT) that has been adapted for the use with medical images using the image case database CasImage that has been developed and maintained by the University Hospital of Geneva and that is in routine use.
The first results show the potential of this technique to retrieve similar cases from a case database. This is an important task for teaching and might also become important for diagnostics using case-based reasoning, for example. For the use as a diagnostic tool, it is foreseen to specialize the visual features for the domain of lung image retrieval, using high resolution computed tomography (HRCT) images.

Keywords
Content-based image retrieval, Medical mage databases, Medical case databases, Visual access methods

1. Introduction

Since the early 1980s, tools have been developed to permit the access to documents based on their content [1] and starting from the 1990s a number of performant tools have been available as research prototypes (Photobook [2], PicToSeek, [3]) and as commercial products (QBIC [4], Virage [5]). These tools to manage multimedia data without using keywords or databases searches with fixed fields were necessary due to the high amount of visual and multimedia data that has become available through the Internet and also through the wide use of digital consumer cameras. Textual annotation of most of these data is too expensive if not impossible.
Most of these systems use automatically extracted, global colours and textures as features and sometimes local features or layout features are taken into account. Some retrieval systems allow the segmentation of the images and then use local features extracted of these regions, such as shapes or relations between the segments. The most prominent of such systems might be Blobworld [6].
A very good and comprehensive overview of the field of content-based image retrieval including its problems such as the semantic gap and the sensory gap can be found in [7].
In the medical field, the production and management of large amounts of visual data is also daily routine and the number of digital images used for diagnostics, therapy and teaching is rising. The use of content-based retrieval techniques in medicine in general [8,9] as well as in PACS systems [10] has been proposed frequently. Still, there is only small number of projects where actual implementations with real medical images using automatically

extracted visual features have taken place such as the IRMA [11] and ASSERT [12] projects.
The second section of this article will explain the *GIFT*/*Viper* system that has been used as the basis for our development. The third section then explains the integration of this content-based search tool into Casimage, a medical image case database developed at the University Hospitals of Geneva. It will also give a short introduction of a planned project with respect to lung image retrieval. The conclusions will then discuss the outcome of this article and future ideas.

2. The *GIFT*/*Viper* framework

The GNU Image Finding Tool (*GIFT*, [13]) is the outcome of the *Viper* project (Visual Information Processing for Enhanced Retrieval, [14,15]) at the University of Geneva. A demonstration version with a Java and PHP interface can also be tried out from the *Viper* web page.
Viper's retrieval mechanisms are based on techniques used commonly in textual information retrieval such as inverted files, frequency-based weights and relevance feedback. The feature space is extremely large (>85,000 possible features) and sparsely populated with each image containing around 1,000-2,000 features. This distribution is meant to model the distribution of words in textual databases.

2.1 Structure

In general, *GIFT* contains a number of separated components where each block can be exchanged with another technique. This component-based structure allows for an easy change of components to adapt the system to certain specialized application domains.

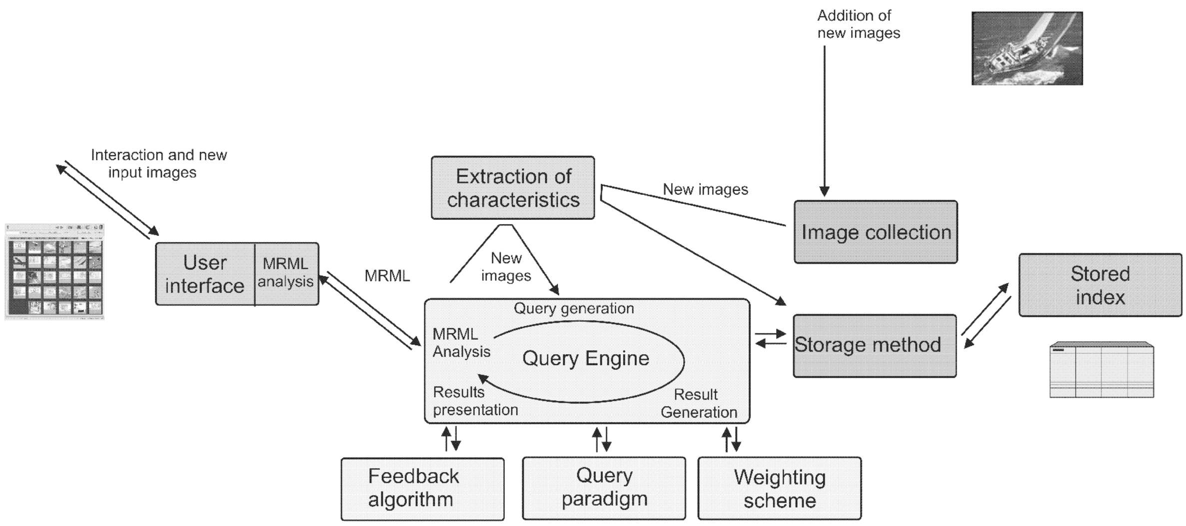

Figure 1: The component-based structure of the GIFT retrieval system

Image features can be changed relatively easy and by this, the system can be adapted to various application domains. An extremely important part is the separation of the user interface and the actual search engine using *MRML*, the Multimedia Retrieval Markup Language. This allows the system to be integrated into different applications with a standard communication interface. More about the distributed framework can be read in [16].

2.2 *MRML* – Multimedia Retrieval Markup Language

MRML [17] is a standard query markup language for content-based image retrieval systems. It is based on XML, so it is human-readable and standard parsers such as expat or Xerces can be used. *MRML*-communications have the structure of a remote procedure call: the client connects to the server, sends a request, and stays connected to the server until the server breaks the connection. The server shuts down the connection after sending the *MRML* message, which answers the request. This connectionless protocol has the advantage of easing the implementation of the server. *MRML*, in its current specification and implementation state, supports the following features:

- request of a capability description from the server,
- selection of a data collection classified by query paradigm; it is possible to request collections which can be queried in a certain manner,
- selection and configuration of a query processor, also classified by query paradigm; *MRML* also permits the configuration of meta-queries during run time,
- formulation of Query By Example (QBE) queries,
- transmission of user interaction data.

All documents are addressed as URLs, which can be file URLs on a local hard disk. A simple query in *MRML* with one relevant and one irrelevant query image looks as follows:

```
<mrml session-id="1">
 <query-step session-id="1"
             resultsize="30"
             algorithm-id="algorithm-default">
  <user-relevance-list>
   <user-relevance-element
    image-location="http://viper.unige.ch/1.jpg"
    user-relevance="1"/>
   <user-relevance-element
    image-location="http://viper.unige.ch/2.jpg"
    user-relevance="-1"/>
  </user-relevance-list>
 </query-step>
</mrml>
```

2.3 Visual features

The standard *GIFT* system uses a palette of 166 colours, derived by quantizing the HSV space into 18 hues, 3 saturations, 3 values and 4 grey levels. The HSV space is used because its characteristic to be closer to human perception [18] which might be less important when actually using grey-level images. Two sets of features are extracted from the quantized image. The first is a colour histogram, with empty bins are discarded. The second represents colour layout. Each block in the image (the first being the image itself) is recursively divided into four equal-sized blocks, at four scales. The occurrence of a block with a given mode colour is treated as a binary feature.

Gabor filters have been applied to texture classification and segmentation, as well as more general vision tasks [19,20]. We employ a bank of real circularly symmetric Gabor filters Three scales are used. The four orientations are: $\theta_0=0$, $\theta_{n+1}=\theta_n+\pi/4$. The resultant bank of 12 filters gives good coverage of the frequency domain, and little overlap between filters. The mean energy of each filter is computed for each of the smallest blocks in the image. This is quantized into 10 bands. A feature is stored for each filter with energy greater than the lowest band. Histograms of the mean filter outputs are used to represent global texture characteristics.

2.4 Weighting schemes

As the features in *GIFT* are not present in every image but are only present in a certain number of images, similar to words in a text, we can use weighting schemes for the importance of features similar to those used in textual information processing [21].

The two basic principles derived from information theory in our program are:

- Features that are frequent in an image describe this image well (described by *df*, document frequency),

- features that are frequent in the entire collection do no distinguish images well from each other and are thus less important (*cf*, collection frequency).

3. CasImage

CasImage [22,23] is a digital teaching files authoring environment that is integrated into the PACS environment at the University Hospitals of Geneva. It allows *importing* images into the system in the DICOM format and via DICOM communication, the *editing* of textual data that is stored with the images (including an anonymization of the images) and the access to the images via a web server or offline access via CD-ROMs containing part of the image collection.

CasImage is routinely used by the MDs at the University Hospital of Geneva. At the moment, the system contains more than 25,000 images, both grey level and colour images.

4. Results

No quantitative evaluation has been done as of yet but the first results of an integration are promising. The search with an example image delivers in general a number of images with the same anatomic region on screen. Proper quantitative evaluations need to follow.

4.1 Integration of *GIFT* into CasImage

Thanks to the *MRML* interface of the *GIFT* framework, it is relatively easy to integrate the system into various applications that use web interfaces. Figure 2 shows a simple PHP web interface with the images of the CasImage system after the executing of the query with an example image. The example query shows that most of the retrieved images are indeed similar.

Figure 2: A content-based query using the CasImage images, shown in a web browser.

We also included a DICOM import into the *GIFT* system that allows using these medical images directly and without a loss in quality (*ie.* through a JPEG compression beforehand). Our first tests have been executed with JPEG images that were exported from the CasImage system. This can lead to problems as different level/window settings when converting the images into JPEGs can lead to fairly different grey values in the images.

4.2 Adaptation of *GIFT* features

The features that are used in the standard *GIFT* system showed their effectiveness in retrieval of stock photography but for medical image retrieval other criteria need to be used. The 25.000 images of CasImage do contain a significant number of colour images (dermatology, pathology, hematology, etc), so colour can not be neglected completely. For this reason we simply use a different quantization of the HSV space using 6 Hues, 3 Values and 3 saturations plus 32 grey values. The use of a larger number of grey values seems to introduce noise and does not lead to better results with the images that we currently have. This might change when using calibrated images in DICOM format without present level/window settings such as HRCT lung images. Quantitative tests on the colour space quantization still need to be carried out to find an optimal quantization.

The texture features using Gabor filter of different scales and directions are kept for the moment. It still has to be shown whether other characteristics such as those based on cooccurence matrices, wavelets or Wold features can complement or replace the Gabor filters. Tests on different feature sets are being performed at the moment.

4.3 Lung Image Retrieval

Long-term goal of the project is not only the use of the retrieval system as a teaching tool but also its clinical use for diagnostics. For the diagnostic use it is planned to use HRCT images of the lung as a first test bed. The pathology of these images can well be described by its texture in certain regions of the image.

Goal is to have physicians to select slices that best characterize a disease and mark the region of interest in these images. The system can then be used to find similar cases in the database based on the texture and grey value characteristics. It needs to be stressed that the tool does not do any diagnostics but only provides the MD with similar cases for case-based reasoning, for example.

5 Conclusions

The first results of using a freely available image retrieval system to achieve content-based access to medical images are very promising. Although no quantitative evaluation has been done as of yet, the first results show that the system already works reasonably well with only a few changes. When using features specially developed for medical images the results promise to get even better.

Lung image retrieval of HRCT images will be the first domain that we will specialize our system for. These images exhibit texture features that describe a pathology well and they will be a good test bed for trying several texture descriptors.

Another important domain for improvement is the inclusion of not only query functionalities into *MRML* but also database functionalities. An application like CasImage can then easily manage its entire database for content-based access via *MRML* on a remote Linux server.

6. References

[1] N-S. Chang and K-S. Fu, Query-by-Pictorial-Example, *IEEE Transactions on Software Engineering* 6(6): 519-524, 1980.

[2] A. Pentland, RW. Picard and S. Sclaroff, Photobook: Tools for Content-Based Manipulation of Image Databases, *International Journal of Computer Vision* 18(3):233-254, 1996.

[3] T. Gevers, PicToSeek: A Content-Based Image Search System for the World Wide Web, *Conference on Visual Information Systems*, San Diego, USA, 1997.

[4] M. Flickner, H. Sawhney, W. Niblack, J. Ashley, Q. Huang, B. Dom, M. Gorkani, J. Hafner, D. Lee, D. Petkovic, D. Steele and P. Yanker, Query by Image and Video Content: The QBIC System, *IEEE Computer* 28(9):23-32, 1995

[5] JR.Bach, C. Fuller, A. Gupta, A. Hampapur, B. Horowitz, R. Humphrey, R. Jain, C-F. Shu, The Virage Image Search Engine: An Open Framework for Image Management, Storage and Retrieval for Image and Video Databases IV, SPIE volume 2670, San Jose, California, USA, pp. 76-87, 1996.

[6] C. Carson, M. Thomas, S. Belongie, JM. Hellerstein and J. Malik, Blobworld: A System for Region-Based Image Indexing and Retrieval, *Conference on Visual Information Systems*, Amsterdam, The Netherlands, pp. 509-516, 1999.

[7] AWM. Smeulders, M. Worring, S. Santini, A. Gupta and R. Jain, Content-Based Image Retrieval at the End of the Early Years, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(12):1349-1380, 2000.

[8] H. Lowe, I. Antipov, W. Hersh and C. Arnott Smith, Towards Knowledge-Based Retrieval of Medical Images. The Role of Semantic Indexing, Image Content Representation and Knowledge-Based Retrieval, *Annual Symposium of the American Society for Medical Informatics (AMIA)*, pp. 882-886, 1998.

[9] HD. Tagare, C. Jaffe and J. Duncan, Medical Image Databases: A Content-Based Retrieval Approach, *Journal of the American Medical Informatics Association (JAMIA),* 4(3):184-198, 1997.

[10]H. Qy and WE. Snyder, Content-Based Image Retrieval in PACS, *Journal of Digital Imaging* 12(2), 1999.

[11]D. Keysers, J. Dahmen, H. Ney, BB. Wein and TM. Lehmann, A Statistical Framework for Model-Based Image Retrieval in Medical Applications, *Journal of Electronic Imaging*, 2003 (in press).

[12]C-R. Shyu, CE. Brodley, AC. Kak, A. Kosaka, AM. Aisen and LS. Broderick, ASSERT: A Physician-in-the-loop Content-Based Retrieval System for HRCT Image Databases, *International Journal on Computer Vision and Image Understanding, Special Issue on Content-Based Access for Image and Video Libraries* 75(1,2):111-132, 1999.

[13]GIFT web page: http://www.gnu.org/software/gift/

[14]DMcG. Squire, W. Müller, H. Müller and J. Raki, Content-Based Query of Image Databases, Inspirations from Text Retrieval: Inverted Files, Frequency-Based Weights and Relevance Feedback, *Scandinavian Conference on Image Analysis (SCIA)*, pp. 143-149, Kangerlussuaq, Greenland, 1999.

[15]Viper web page: http://viper.unige.ch/

[16]H. Müller, W. Müller, DMcG. Squire, Z. Pecenovic, S. Marchand-Maillet and T. Pun, An Open Framework for Distributed Multimedia Retrieval, *Computer Assisted Information Retrieval (RIAO)*, volume 1, pp. 701-712, Paris, 2000.

[17]MRML web page: http://mrml.net/

[18]JR. Smith, S-F. Chang, Tools and Techniques for Color Image Retrieval, *Storage and Retrieval for Image and Video Databases*, SPIE volume 2670, pp. 426-437, San Jose, Ca, USA, 1996.

[19]WY. Ma, BS. Manjunath, Texture Features and Learning Similarity, *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*,San Francisco, California, USA, pp. 425—430, 1996.

[20]A. Jain, G. Healey, A Multiscale Representation Including Opponent Color Feature for Texture Recognition, *IEEE Transactions on Image Processing*, 7(1):124-128, 1998

[21]G. Salton and C. Buckley, Term Weighting Approaches in Automatic Text Retrieval, *Information Processing and Management* 24(5):513-523, 1988.

[22]Casimage web page: http://www.genisis.ch/casimage/

[23]A. Rosset, O. Ratib, A. Geissbuhler and J-P. Vallée, Integration of a Multimedia Teaching and Reference Database in a PACS Environment, *Radiographics*, 22(6):1567-1577, 2002

7. Address for correspondence

Henning Müller

University Hospitals of Geneva

Division of Medical Informatics

21, rue Micheli-du-Crest,

CH-1211 Geneva 4, Switzerland

henning.mueller@dim.hcuge.ch