

LIRe: Lucene Image Retrieval - An Extensible Java CBIR Library

Mathias Lux
 Institute for Information Technology
 Klagenfurt University
 Klagenfurt, Austria
 mlux@itec.uni-klu.ac.at

Savvas A. Chatzichristofis
 Department of Electrical & Computer
 Engineering
 Democritus University of Thrace
 Thrace, Greece
 schatzic@ee.duth.gr

ABSTRACT

LIRe (Lucene Image Retrieval) is a light weight open source Java library for content based image retrieval. It provides common and state of the art global image features and offers means for indexing and retrieval. Due to the fact that it is based on a light weight embedded text search engine, it can be integrated easily in applications without relying on a database server.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval

General Terms

Documentation

Keywords

Image retrieval, image features, image search, image indexing

1. INTRODUCTION

The area of *content based image retrieval* (CBIR) was motivated since the early 1990s by the idea to find and retrieve images independent from metadata other than extracted from the image itself. However a satisfactory solution has not been found yet, but a problem has been isolated: Researchers defined the *semantic gap* [5], which refers to the inability of a machine to fully understand and interpret images based on automatically extracted data (see for instance [7]). In current research efforts in visual information retrieval especially global features, which denote features capturing characteristics of the whole image instead of focusing for instance on segments, regions or patches, have lost part of their significance [4]. In applied research however content based image retrieval – for instance as part of a complex system – is often relying on fast, global features at least as a foundation for further research. LIRe (Lucene Image Retrieval) is an efficient library allowing researchers to integrate CBIR based on global features in an easy way.

This paper is organized in the following way: A general overview on LIRe is given in the next section and a detailed

technical description is given in the subsequent section. The demo application LIReDemo is described before the paper is summarized and some conclusions on LIRe and its applicability are drawn.

2. LIRE

LIRe is a Java library for content based image retrieval. LIRe extracts image features from raster images and stores them in a Lucene index for later retrieval. LIRe also provides means for searching the index. LIRe is intended for developers and researchers, who want to integrate content based image retrieval features in their applications. Due to the simplicity of the approach (no database and only a few lines of code are needed to integrate LIRe) it is an easy way to test the capabilities of classical CBIR approaches for single application domains. Also the integration of additional image features is possible easily to further extend the functionality of LIRe. Currently the following image features are included in LIRe:

1. Color histograms in RGB and HSV color space.
2. MPEG-7 descriptors scalable color, color layout and edge histogram, see [1].
3. The Tamura texture features coarseness, contrast and directionality, see [8]
4. Color and edge directivity descriptor, CEDD, see [2]
5. Fuzzy color and texture histogram, FCTH, see [3]
6. Auto color correlation feature defined by Huang et al. [6]

While LIRe itself is a development library there is also an additional demo package called LIRe Demo, which allows to test selected CBIR features of LIRe with a graphical user interface as shown in Figure 1. For the sake of speed indexing is done in parallel using multiple threads and for sake of simplicity photos from Flickr can be indexed instead of specifying a local image collection for testing purposes. LIRe Demo also integrates an application built on the capabilities of LIRe for creating image mosaics.

LIRe has been first released in February 2006 and is currently available in release v0.6. Both, LIRe and LIRe Demo, are available licensed under GPL online¹. In addition to

¹URI: <http://semanticmetadata.net/lire>




Figure 1: LIRe Demo loading search results.

the Java class documentation there is also a documentation wiki for users and developers giving code samples, answering FAQs and describing best practice approaches. Latest development code can be accessed on the subversion (SVN) server on <http://www.sourceforge.net>. Note that LIRe is hosted as sub project within the project *Caliph & Emir*, which provides Java based tools MPEG-7 image annotation and retrieval.

3. TECHNICAL DETAILS

LIRe is built on top of the open source text search engine Lucene, hosted at <http://lucene.apache.org>. Like in text retrieval images have to be indexed for later retrieval. The index is organized as a Lucene index, where *documents* consisting of *fields* each having a *name* and a *value* are organized in an index structure typically stored in the file system.

3.1 Indexing

Indexing is done using an implementation of the `DocumentBuilder` interface. A simple approach is to use the `DocumentBuilderFactory`, which creates `DocumentBuilder` instances for all available features as well as popular combinations of features (e.g. *all MPEG-7 features* or *all available features*). A `DocumentBuilder` is basically a wrapper for image features creating a Lucene `Document` from a Java `BufferedImage`. The signatures or vectors extracted by the feature implementations are wrapped in the documents as text. The document output by a `DocumentBuilder` can be added to a Lucene index.

3.2 Search

For search, classes implementing the `ImageSearcher` interface are used. The `ImageSearcher` either takes the given query feature or extracts the feature from a query image. It then reads documents from the index sequentially and compares them to the query image (linear search). Although the main indexing features of Lucene (e.g. an inverted list or stemming) are not employed in this kind of search, LIRe takes advantage of the efficient and fast disk access layer of Lucene, which results in lower search times compared to implementations using the embedded databases HSQLDB², which is used in Open Office, and Apache Derby³, which is also included in the Java runtime releases as Java DB. Also the use of Lucene allows indexes bigger than common RAM restrictions (e.g. smaller than 2 GB on 32 bit Java) and additional indexing of textual metadata for the images.

3.3 Features

Image features in LIRe can also be used for applications outside the provided search and indexing mechanisms. A

²URI: <http://hsqldb.org/>

³URI: <http://db.apache.org/derby/>

LIRe feature typically implements the `LireFeature` interface and therefore provides common methods for extraction from Java `BufferedImage` and the output and parsing of a textual representation of the descriptor. Furthermore the distance function is part of the feature implementation. Currently available features are given in Section 2.

4. PERFORMANCE

For testing runtime performance the well known 1,000 images SIMPLiCity data set of Wang et al. [9] has been used. All test have been done on a desktop computer with an Intel Core 2 Quad CPU with four 2.4 GHz cores and 2 GB RAM running Windows XP having Sun Java 1.6 u6 installed. The test was done in a single thread with the Java client VM. Maximum memory of the VM for the tests was set to 512 MB⁴.

DocumentBuilder	Sec. taken
Auto color correlation	413.06
CEDD	47.55
Color histogram	15.27
Color layout (MPEG-7)	17.77
Edge histogram (MPEG-7)	20.69
FCTH	60.83
Scalable color (MPEG-7)	20.89
Tamura	332.41

Table 1: Runtime for indexing the Wang dataset with 1,000 images.

Table 1 gives an overview on the time needed for indexing the Wang dataset with the mentioned descriptors. Note that auto color correlation is highly configurable and runtime heavily depends on the parameters chosen.

Search performance has been tested with the CEDD descriptor. For search in the indexed Wang data set with 1,000 images 78.1 ms were needed, the respective index had a size of 368 KB. For search in the bigger data set of Wang with 9,908 images 303.2 ms were needed, the respective index had a size of 3.54 MB.

4.1 Extending LIRe

The most common way to extend LIRe is to add additional image features. In this case the `LireFeature` interface has to be implemented. The classes `GenericImageSearcher` and `GenericDocumentBuilder` provide easy methods to implement indexing and search based on arbitrary features implementing the interface. To provide different combinations of features the `ChainedDocumentBuilder` can be used to link several different `DocumentBuilder` implementations.

5. LIRE DEMO

In addition to the Java library a demo application is also part of the LIRe project. The demo offers an easy way to test the basic functions with arbitrary image data (PNG and JPEG images supported). LIRe demo allows to index a set of images and to add images to an existing index. If no images from a local repository are given, a set of sample photos is downloaded from Flickr⁵. It then allows to browse the indexed images and to issue queries by example. The

⁴with parameters `-Xms256M -Xmx512M`

⁵URI: <http://flickr.com>

results window is shown in Figure 2. All basic integrated `ImageSearcher` implementations can be used for retrieval. Furthermore an example application for creating image mosaics (see Figure 3) based on LIRe is part of the demo.

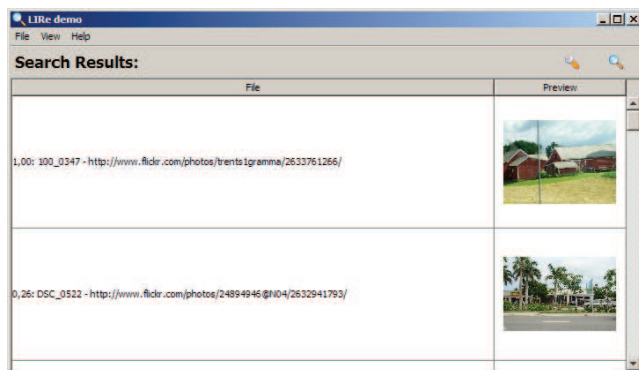


Figure 2: LIRe Demo search results.

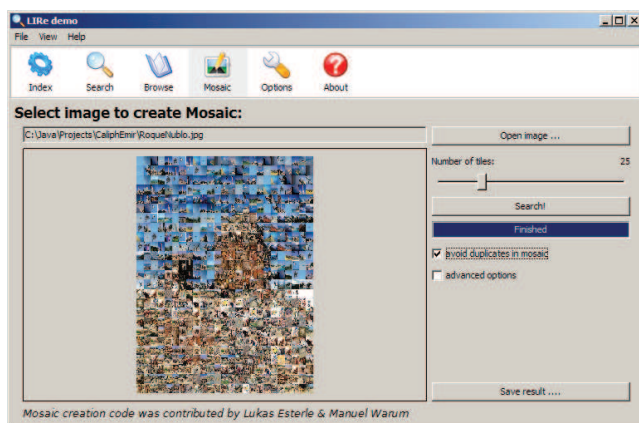


Figure 3: LIRe Demo image mosaic creation.

6. SUMMARY & CONCLUSIONS

LIRe is a light-weight Java library for content based image retrieval provided under GPL license. LIRe is fast enough to provide a platform for applied research relying on standard CBIR techniques. Due to the fact that it uses the Lucene text search engine and not a database server it is easy to integrate LIRe in existing prototypes and applications. Due to Java background of LIRe and it can be used on many different platforms. For applications retrieving images from big databases or applications with many concurrent users, the linear search approach of LIRe is not appropriate. However the open source nature of the LIRe project allows to extend LIRe towards better indexing strategies and faster than linear search.

7. ACKNOWLEDGEMENTS

We would like to thank the numerous people having contributed code or patches to LIRe or having published their work as open source allowing integration into LIRe: Arthur Pitman, Bastian Hösch, Benjamin Sznajder, Christian Penz, Daniel Pöttinger, Fabrizio Falchi, Janine Lachner, Katharina Tomanec, Lukas Esterle, Manuel Orazo, Marco Vanossi, Marko Keuschnig, Oge Marques, Roman Divotkey and Roman Kern.

LIRe and LIRe Demo are available online along with documentation on <http://sourceforge.net> and <http://www.semanticmetadata.net>.

8. REFERENCES

- [1] S.-F. Chang, T. Sikora, and A. Puri. Overview of the mpeg-7 standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 11(6):688–695, June 2001.
- [2] S. A. Chatzichristofis and Y. S. Boutalis. Cedd: Color and edge directivity descriptor. a compact descriptor for image indexing and retrieval. In A. Gasteratos, M. Vincze, and J. Tsotsos, editors, *Proceedings of the 6th International Conference on Computer Vision Systems, ICVS 2008*, volume 5008 of *LNCIS*, pages 312–322, Santorini, Greece, May 2008. Springer.
- [3] S. A. Chatzichristofis and Y. S. Boutalis. Fcth: Fuzzy color and texture histogram a low level feature for accurate image retrieval. In *Proceedings of the 9th International Workshop on Image Analysis for Multimedia Interactive Services, WIAMIS 2008*, pages 191–196, Klagenfurt, Austria, May 2008. IEEE.
- [4] R. Datta, D. Joshi, J. Li, and J. Z. Wang. Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys*, 40(2):1–60, 2008.
- [5] A. Del Bimbo. *Visual Information Retrieval*. Morgan Kaufmann Publishers, San Francisco, June 1999.
- [6] J. Huang, S. R. Kumar, M. Mitra, W.-J. Zhu, and R. Zabih. Image indexing using color correlograms. In *Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition, CVPR '97*, volume 00, pages 762–768, San Juan, Puerto Rico, June 1997. IEEE.
- [7] A. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain. Content-based image retrieval at the end of the early years. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(12):1349–1380, December 2000.
- [8] H. Tamura, S. Mori, and T. Yamawaki. Textural features corresponding to visual perception. *IEEE Transactions on Systems, Man, and Cybernetics*, 8(6):460–472, June 1978.
- [9] J. Z. Wang, J. Li, and G. Wiederhold. Simplicity: Semantics-sensitive integrated matching for picture libraries. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(9):947–963, 2001.