

# Architecture and applications of the Parallel Distributed Image Search Engine (ParaDISE)

Roger Schaer, Dimitrios Markonis, Henning Müller

HES–SO, Business Information Systems,  
TechnoArk 3 3960 Sierre, Switzerland  
roger.schaer@hevs.ch  
dimitrios.markonis@hevs.ch  
henning.mueller@hevs.ch

**Abstract:** Medical images are important to physicians for diagnosis and treatment planning. The number of images produced in medical institutions has been increasing rapidly over the past years. Making these imaging data available and allowing medical professionals to perform retrieval based on visual characteristics of images is the challenge that content-based image retrieval systems are tackling. In this paper we propose a system called ParaDISE (Parallel Distributed Image Search Engine). It allows indexing and retrieving images using both visual features as well as text. It is designed to scale well with big data sets, and it should be easily extensible with new features. The architecture of the system is described and its use in medical applications is presented. An open source release of the system is in progress.

## 1 Introduction

Medical images are a precious tool in aiding doctors to diagnose and treat a variety of diseases. [Doi05]. The amount of images produced by hospitals has been growing at an almost exponential rate over the past 20 years (as an example, the Radiology Department of the University Hospitals of Geneva produced over 200'000 images per day in 2011).

Accessing this huge amount of data is a challenge that CBIR (Content-based image retrieval) systems attempt to deal with, by allowing users to search for images based solely on their visual content (example images or image regions), without requiring the input of textual keywords. Studies show that CBIR systems can assist radiologists in diagnosis, particularly radiologists with little experience [ABWM<sup>+</sup>03, LADT09]. Novel image retrieval algorithms have been published [LADT09, MMBG04].

In recent years, retrieval systems have become increasingly complex, improving text search by using semantics, as well as utilizing modern visual representations such as Bag-of-Visual-Words [SZ03] and Vectors of Locally Aggregated Descriptors (VLAD) [JDS10]. Additionally, multimodal search (text combined with CBIR) is becoming more common, driving the need for interoperability and expandability in retrieval systems.

This paper presents the Parallel Distributed Image Search Engine (ParaDISE), an image

retrieval system that combines text with CBIR techniques. The system was designed to be scalable, meaning it should be able to handle ever-growing datasets and large image collections that are available on the Internet. It was built in a modular fashion, consisting of several independent components, thus increasing the flexibility and distributability of the system. The third main aspect of ParaDISE is interoperability, which is achieved by the use of standard protocols and data formats in order to facilitate the integration of ParaDISE with other retrieval systems.

The main driving force behind the development of ParaDISE is KHRESMOI<sup>1</sup>, a European project focused on medical information analysis and retrieval [HBGM11]. The goal of KHRESMOI is to develop a multilingual, multimodal search system for biomedical information and documents.

## 2 Related Work

Several projects have already been realized in the field of information retrieval in the medical domain and made systems available as open source. Among them is the Viper project [SMM<sup>+</sup>01], the outcome of which was the GNU Image-Finding Tool (GIFT), a CBIR system that enables users to perform "Query By Example" search operations and improve the quality of results using relevance feedback. ParaDISE uses more recent and sophisticated features than GIFT, which only uses visual information for querying and relevance feedback.

Another noteworthy project is Lire (Lucene Image Retrieval) [LC08], a library based on the Lucene text retrieval software, which provides image retrieval using various visual features related to color and texture. The system uses purely visual search and provides little support for several state-of-the-art representations (such as spatial pyramid matching or bag-of-colors), indexing parallelization or flexible index structuring.

Flexible Image Retrieval Engine (FIRE) is yet another example of a CBIR system [DKN08], also used in medical image retrieval evaluation apart from other applications. Being developed before 2007, the system does not support state-of-the-art mid-level representations (such as vectors of locally aggregated descriptors). No parallelization schema is mentioned for indexing large scale datasets, either.

Additionally, there are also several web applications available for medical image retrieval, which use various retrieval techniques, including CBIR and semantic data. GoldMiner<sup>2</sup> is a text-based radiology image search engine which allows filtering results based on imaging modality. RadMiner<sup>3</sup> is a search tool in radiology reports using semantic information. Yottalook<sup>4</sup> is a web-based radiology image search system which supports searching across multiple types of sources. These three systems do not currently use any visual information for searching. OpenI [DFAST12] allows search based on text and on image examples.

---

<sup>1</sup><http://khresmoi.eu> (accessed on 19-06-2014)

<sup>2</sup><http://goldminer.arrs.org> (accessed on 19-06-2014)

<sup>3</sup>[http://www.averbis.de/public/download/RadMiner\\_brochure\\_en.pdf](http://www.averbis.de/public/download/RadMiner_brochure_en.pdf) (accessed on 19-06-2014)

<sup>4</sup><http://yottalook.com> (accessed on 19-06-2014)

ParaDISE is geared towards a higher level of abstraction, allowing users to configure and expand every aspect of the system (search, storage, indexing, fusion, etc.). It contains many state-of-the-art descriptors and allows for parallelization in indexing and retrieval.

### 3 System Description

**Backend** The various features provided by the ParaDISE backend are distributed among four main components, shown below: The **Extractor** extracts local visual features from an image, such as Scale Invariant Feature Transform (SIFT) [Low04]. The **Descriptor** creates a vector that serves as a global representation of an image, such as Bag-of-Visual-Words. The **Storer** accesses visual indices (read/write operations). ParaDISE supports several systems to store and access the images' visual descriptors, from simple comma-separated values (CSV) files to database systems such as MySQL or Apache CouchDB. The **Fusor** combines the results of multiple query images, or multiple features of the same image. It is also designed for integrating ParaDISE with other systems.

These four basic components are integrated into two composite components (Indexer and Seeker) to provide the two main retrieval features: indexing and similarity search. The Indexer can perform serial as well as parallel indexing using Apache Hadoop [MSE<sup>+</sup>12], allowing the distribution of the indexing workload across several computers and thus reducing the overall time needed to perform the task, as well as making it easier to scale up to bigger datasets (by simply adding nodes to the Hadoop cluster). It supports building Approximate Nearest Neighbor (ANN) search indices and weighting the indices. The Seeker can use various distance and similarity metrics for the comparison of the vectors. It also supports multiple queries and relevance feedback techniques. Additionally, the Seeker takes advantage of the ANN indices created by the Indexer, helping to keep retrieval time short in spite of very large datasets (by restricting the search to a subset of images).

ParaDISE has been used in the ImageCLEF<sup>5</sup> 2013 challenge, evaluating image retrieval performance on a dataset containing images from the biomedical literature [GSdHMS<sup>+</sup>13].

**Frontend** The main features of ParaDISE are exposed in a web service layer built on top of the ParaDISE backend. The web services are based on the principles of the Representational state transfer (REST) architectural style for web applications. They use standard Hypertext Transfer Protocol (HTTP) requests, making interoperability and integration with client applications (both web-based and desktop clients) easily achievable. The JavaScript Object Notation (JSON) data exchange format is used for communication. Additional features such as the text search are also provided by the web service layer. The architecture of the web services is explained below and shown in Figure 1. The **Caption web service** is responsible for search by keyword and uses the Lucene search engine. The **Visual web service** uses the two composite components from the ParaDISE backend (Indexer & Seeker) to provide content-based image retrieval and indexing. The **Fusion web service** is responsible for combining results from different sources (text search, image search) into

---

<sup>5</sup><http://www.imageclef.org/> (accessed on 19-06-2014)

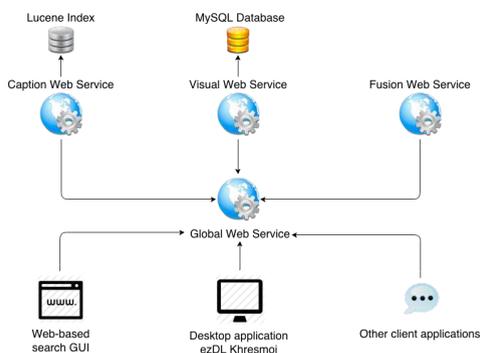


Figure 1: Overview of the ParaDISE web service layer architecture.

a single list (based on a given fusion rule), using the Fusor component from the ParaDISE backend. The **Global web service** acts as a facade for client applications. It hides the complexity of calling the underlying web services, for example by providing a general "search images" method, which can use text, visual information, or both.

## 4 Applications

**KHRESMOI** ParaDISE is used in the KHRESMOI subsystem that accesses images in the open access medical literature. This part of KHRESMOI supports text queries, queries by image example and the combination of these two for accessing the image data. It allows filtering by modality (by applying automatic modality classification methods) and is interconnected with other tools, such as the semantic search and 3D image retrieval components of Khresmoi. It supports searching into subfigures of multi-panel figures with the help of automatic compound figure detection and separation methods [CMFR<sup>+</sup>13]. ParaDISE is used as the backbone of this subsystem. Indexing a dataset of 1.7 million PubMed Central images with the parallel computing mode of ParaDISE took approximately 6 hours.

**Shambala and Leap Motion™** Whereas ezDL, used in the Khresmoi system, is a comprehensive interface, Shambala is an alternative search interface mainly targeted towards simple navigation. It is built on 3 principles : it should be easy to use (not overwhelming the user with options and parameters), interactive (allow the user to drag&drop images, update search results in real-time) and modern (use of HTML5 and Leap Motion™ technology<sup>6</sup>). Leap Motion allows the user to interact with the interface using only hand and finger gestures, without needing to use a mouse or keyboard. More details about gesture interaction with Shambala can be found in [WSMM14].

<sup>6</sup><http://leapmotion.com> (accessed on 19-06-2014)

**SearchParadise and Google Glass™** Another application prototype which uses ParaDISE, called SearchParadise, was created for the Google Glass™, a wearable computing device developed by Google<sup>7</sup>. The application allows the user to perform a search query based on a photo taken by the Glass camera and see the results on the display located in front of the right eye. This can allow physicians to get additional information on situations or lesions that they have and does not interrupt the contact with the patient.

## 5 Conclusions

Medical image retrieval can potentially benefit from novel techniques such as CBIR. Taking into account the characteristics of the domain, CBIR systems need to provide scalable and interoperable solutions. The ability to evaluate different image representations and to add new visual features needs to be supported, as well. ParaDISE, the image retrieval system described in this paper, provides a flexible solution for image indexing and similarity search operations on large datasets. It can be easily integrated with other systems and extended with new features. It can be used for very diverse client applications, including desktop applications, web-based interfaces as well as emerging technologies such as the Google Glass™. Future steps for the system include using a distributed computing framework for online operations (retrieval) that is currently only implemented for the offline indexing steps. Releasing the source code to the public is currently under way to encourage the addition of new features by third parties.

## 6 Acknowledgments

This work was supported by the EU 7th Framework Program in the context of the Khresmoi project (FP7 grant 257528).

## References

- [ABWM<sup>+</sup>03] Alex M. Aisen, Lynn S. Broderick, Helen Winer-Muram, Carla E. Brodley, Avinash C. Kak, Christina Pavlopoulou, Jennifer Dy, Chi-Ren Shyu, and Alan Marchiori. Automated Storage and Retrieval of Thin-Section CT Images to Assist Diagnosis: System Description and Preliminary Assessment. *Radiology*, 228(1):265–270, July 2003.
- [CMFR<sup>+</sup>13] Ajad Chhatkuli, Dimitrios Markonis, Antonio Foncubierta-Rodríguez, Fabrice Meriaudeau, and Henning Müller. Separating compound figures in journal articles to allow for subfigure classification. In *SPIE Medical Imaging*, 2013.
- [DFAST12] Dina Demner-Fushman, Sameer Antani, Matthew S. Simpson, and George R. Thoma. Design and Development of a Multimodal Biomedical Information Re-

---

<sup>7</sup><http://www.google.com/glass> (accessed on 19-06-2014)

- trieval System. *Journal of Computing Science and Engineering*, 6(2):168–177, 2012.
- [DKN08] Thomas Deselaers, Daniel Keysers, and Hermann Ney. Features for image retrieval: an experimental comparison. *Information Retrieval*, 11(2):77–107, April 2008.
- [Doi05] K Doi. Current status and future potential of computer-aided diagnosis in medical imaging. *British Journal of Radiology*, 78:3–19, 2005.
- [GSdHMS<sup>+</sup>13] Alba García Seco de Herrera, Dimitrios Markonis, Roger Schaer, Ivan Eggel, and Henning Müller. The medGIFT Group in ImageCLEFmed 2013. In *Working Notes of CLEF 2013 (Cross Language Evaluation Forum)*, September 2013.
- [HBGM11] Allan Hanbury, Célia Boyer, Manfred Gschwandtner, and Henning Müller. KHRESMOI: Towards a Multi-Lingual Search and Access System for Biomedical Information. In *Med-e-Tel, Luxembourg, 2011*, pages 412–416, 2011.
- [JDS10] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Aggregating local descriptors into a compact image representation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3304 – 3311, June 2010.
- [LADT09] L. Rodney Long, Sameer Antani, Thomas M. Deserno, and George R. Thoma. Content-Based Image Retrieval in Medicine: Retrospective Assessment, State of the Art, and Future Directions. *International Journal of Healthcare Information Systems and Informatics*, 4(1):1–16, January 2009.
- [LC08] Mathias Lux and Savvas A. Chatzichristofis. Lire: lucene image retrieval: an extensible java CBIR library. In *Proceedings of the 16th ACM international conference on Multimedia*, pages 1085–1088, October 2008.
- [Low04] David G. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- [MMBG04] Henning Müller, Nicolas Michoux, David Bandon, and Antoine Geissbuhler. A review of content-based image retrieval systems in medicine—clinical benefits and future directions. *International Journal of Medical Informatics*, 73(1):1–23, 2004.
- [MSE<sup>+</sup>12] Dimitrios Markonis, Roger Schaer, Ivan Eggel, Henning Müller, and Adrien Depeursinge. Using MapReduce for Large-scale Medical Image Analysis. In *2nd IEEE Conference on Healthcare Informatics, Imaging and Systems Biology (HISB)*, September 2012.
- [SMM<sup>+</sup>01] David McG. Squire, Henning Müller, Wolfgang Müller, Stéphane Marchand-Maillet, and Thierry Pun. Design and Evaluation of a content-based image retrieval system. In *Design & Management of Multimedia Information Systems: Opportunities & Challenges*, chapter 7, pages 125–151. Idea Group Publishing, London, 2001.
- [SZ03] Josef Sivic and Andrew Zisserman. Video Google: A Text Retrieval Approach to Object Matching in Videos. In *Proceedings of the Ninth IEEE International Conference on Computer Vision - Volume 2, ICCV '03*, pages 1470–1477, Washington, DC, USA, 2003. IEEE Computer Society.
- [WSMM14] Antoine Widmer, Roger Schaer, Dimitrios Markonis, and Henning Müller. Gesture interaction for content-based medical image retrieval. In *ICMR*, 2014.