Mobile Medical Image Retrieval

Samuel Duc[a], Adrien Depeursinge[a], Ivan Eggel[a], Henning Müller[ab]

[a]University of Applied Sciences Western Switzerland (HES–SO), Sierre, Switzerland;
[b]Medical Informatics, University Hospitals & University of Geneva, Switzerland

ABSTRACT

Images are an integral part of medical practice for diagnosis, treatment planning and teaching. Image retrieval has gained in importance mainly as a research domain over the past 20 years. Both textual and visual retrieval of images are essential. In the process of mobile devices becoming reliable and having a functionality equaling that of formerly desktop clients, mobile computing has gained ground and many applications have been explored. This creates a new field of mobile information search & access and in this context images can play an important role as they often allow understanding complex scenarios much quicker and easier than free text. Mobile information retrieval in general has skyrocketed over the past year with many new applications and tools being developed and all sorts of interfaces being adapted to mobile clients.

This article describes constraints of an information retrieval system including visual and textual information retrieval from the medical literature of BioMedCentral and of the RSNA journals Radiology and Radiographics. Solutions for mobile data access with an example on an iPhone in a web–based environment are presented as iPhones are frequently used and the operating system is bound to become the most frequent smartphone operating system in 2011. A web–based scenario was chosen to allow for a use by other smart phone platforms such as Android as well. Constraints of small screens and navigation with touch screens are taken into account in the development of the application. A hybrid choice had to be taken to allow for taking pictures with the cell phone camera and upload them for visual similarity search as most producers of smart phones block this functionality to web applications.

Mobile information access and in particular access to images can be surprisingly efficient and effective on smaller screens. Images can be read on screen much faster and relevance of documents can be identified quickly through the use of images contained in the text. Problems with the many, often incompatible mobile platforms were discovered and are listed in the text. Mobile information access is a quickly growing domain and the constraints of mobile access also need to be taken into account for image retrieval. The demonstrated access to the medical literature is most relevant as the medical literature and their images are clearly the largest knowledge source in the medical field.

Keywords: mobile image retrieval, medical image retrieval, information retrieval, content–based image retrieval, multimodal image retrieval

1. INTRODUCTION

There is currently a tendency of persons to be online and accessible permanently through the use of mobile Internet clients or smart phones that start having very similar capabilities to what desktop computers had only a few years ago. Information search is one of the domains where mobile access has enormously increased over the past years. An extensive review of mobile search is given in,[1] including a think tank event with a survey where academic and industry partners were asked about the scientific and commercial potential of mobile search and current tendencies in these fields. The enormous potential and impact become quickly clear after reading the text but also the problems of a very fast development with quick changes and the fact that there are few standard operating systems and that most applications are not compatible across platforms.

An analysis of mobile applications in health care is given in,[2] although this article is slightly outdated. A more recent text on mobile information system access is[3] introducing some of the problems of accessing medical data. It is very clear that medical information needs most often do not arise when being in front of a computer in

Further author information: (Send correspondence to Henning Müller, Email: henning.mueller@hevs.ch)

an office but can occur at any time when a patient is seen or when an accident happens. Quick mobile search can help here enormously and can also take into account factors such as exact location of the person automatically via a GPS (close to a hospital, in the nature, etc.).

Compared to textual information search with many commercial search engines such as Google, visual information retrieval is still a mainly scientific domain although some successful companies with visual search such as LookThatUp*. In general, image search can also be by textual keywords or automatically extracted visual features.[4] Search with visual features representing the images is often called content–based image retrieval (CBIR) as the visual content is directly used as a query. Many tools have been developed for visual information search ranging from invariant features to object detection and recognition in images.[5,6] In the past few years mainly the so–called salient–point–based features[7] or visual words[8] have been used as they have obtained best results in many evaluation campaigns and benchmarks in visual information retrieval and object recognition.[9,10] Another commercial example for a mobile visual search tool is Google goggles, which allows users to take a picture with a smartphone, upload it to a web page for similarity search, and then gives access to more data on the picture. This can be used by tourists, for example, for obtaining information on buildings and other sights.

The medical domain has often been mentioned as one of the domains with a large potential impact of image retrieval[5] and visual search has been requested for several applications in the field.[11,12] A review of a large number of systems and techniques for medical content–based retrieval is given in.[13] One of the few articles with an evaluation in a clinical setting is.[14] The article shows that particularly less experienced radiologists improve their diagnosis quality significantly. Many other articles have also shown the strong potential that image retrieval has for the medical field.

Much of the medical knowledge is currently stored in the medical literature. An increasingly large proportion of this literature is becoming available on the Internet either through Open Access publishers (i.e. BioMed Central†) or with commercial publishers making articles available several months after their publications (such as the RSNA — Radiological Society of North America). Medical articles usually contain a large amount of images and thus image retrieval from the literature has been explored several times, for example in the ImageCLEF‡ benchmark. In ImageCLEF, since 2004, around 17 research groups per year compare the performance of their visual and textual retrieval systems.[15] The benchmark also showed that generally textual information retrieval works much better than visual similarity search. On the other hand, visual classification and retrieval can well help the textual search and purely visual retrieval often has very good early precision. Several web systems exist as well that allow accessing the medical literature such as GoldMiner §[16] or the Yale Image Finder. Ideas for purely visual access to images from the literature are presented in.[17]

It thus sounds natural to combine the content of the literature with visual and textual retrieval techniques and an interface to a mobile phone for mobile information search.

Of course, there are also other knowledge sources such as the LinkedLifeData¶ that combines medical knowledge coded in the RDF (Resource Description framework) standard. These semantic knowledge sources can be another driver for making medical and also mobile search more efficient and more effective. The Khresmoi‖ research project currently aims at combining such semantic sources with the current textual and visual information search tools.

When starting to think about mobile solutions there are many technical problems, though. There is no single, stable mobile platform to develop applications for. Percentages of operating systems on the market can change very quickly and new generations of smart phones appear every few months, making it hard to choose a particular platform. Application environments to develop native applications for several platforms exist but none of them currently seems to be stable enough for production work. An important design choice is thus whether to develop a web application or a native tool. Web applications are portable between platforms but often do

*http://www.ltutech.com/

†http://www.biomedcentral.com/

‡http://www.imageclef.org/

§http://goldminer.arrs.org/

¶http://www.linkedlifedata.org/

‖http://www.khresmoi.eu/

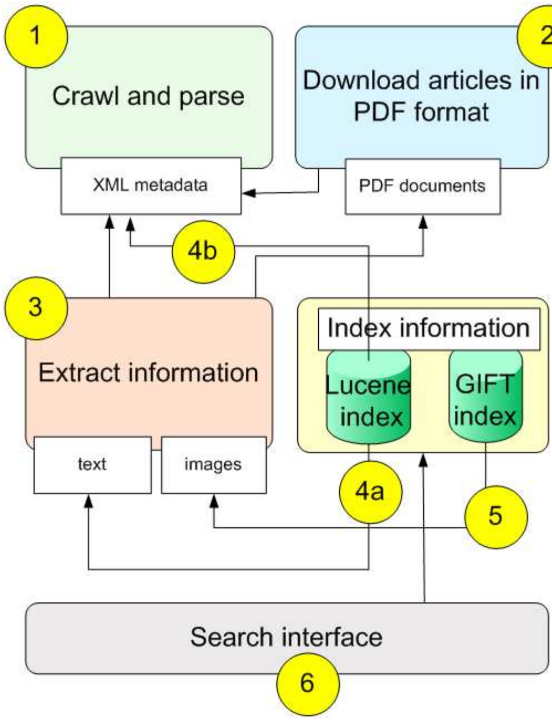

Figure 1: The existing architecture of the MedSearch system that was in large part reused for the mobile application developed. The described system includes a crawler for medical articles, a separator of the images and the text and finally the two search tools GIFT and Lucene.

not have the full functionality (such as accessing the phone's camera), whereas native applications can access all resources but can not be used on more than a single platform. Hybrid applications are a solution, where a web application is used but specific functionalities can be supported by a native application.

This article describes a hybrid application using a native part for the photo upload and a web–based applications for most of the retrieval functionality. The system relies on open source tools and openly accessible data bases of the medical literature. This article is structured as follows: Section 2 describes the basic tools and databases used for the developed application. Section 3 describes the developed solution, mainly through practical screen shots but also with a short performance analysis. Finally, Section 4 concludes the text

2. METHODS

This section describes the methods that are reused for the work described in this paper such as existing tools and databases.

2.1. Tools Reused

The system described in this paper is based on the Medsearch** system that uses the medical open access literature available from BioMedCentral indexed.[18] The components used by MedSearch are the GNU Image Finding Tool (GIFT,[19]) for visual content–based search and the Lucene†† system for text retrieval. All interface components of MedSearch are web–based (using the Glassfish application server and Java as programming language) and this is also the case for mobile MedSearch. This was the only way to make the application usable on several mobile platforms in an easy way.

**`http://medgift.unige.ch:8080/MedSearch/faces/Search.jsp?query=lung&x=0&y=0`

††`http://lucene.apache.org/`

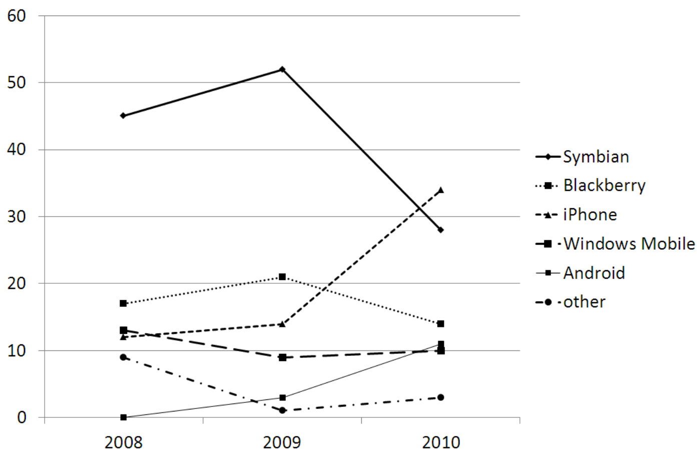

Figure 2: Operating systems used for smart phones taken from.[20]

Figure 1 shows the overall architecture of the existing system that combines open source tools for the crawling of PDF documents, the separation of the PDF documents into text and images and their indexation. The indexed data are then made available via a Java interface using the Lucene and GIFT tools for textual and visual search, respectively. The described architecture can also easily be made available as a web service and can thus be reused for the mobile search application.

To make a choice on the most adapted platform, the mobile OS report[20] was taken into account. This shows clearly the quick changes in the marked as seen in Figure 2. After a dominance of Symbian, the iPhone has become the most frequently used smart phone OS in 2010 already. Android seems to increase strongest in recent months, though.

Due to the quick changes in the market, the decision to build a web–based application was taken, but this solution does not allow to use the built–in camera. To take pictures with the built–in camera for visual information search, the cliqcliq[‡‡] native iPhone application was used. This tool can be called from a web page and deliver the images subsequently back to the web–based application. Initially, platform–independent development tools were analyzed but based on tests these tools were not mature and stable enough.

2.2. Databases Used

Two databases were indexed for the work described in this article. The first database was used for the publicly accessible interface and uses a subset of 12 journals and their articles from the BioMedCentral database[18] . All journals are related to medical informatics and image–intensive medical domains. The database contains in total 38'000 images from over 9'400 articles. A larger subset with 200'000 images was also indexed and is available but the response times are slower and it is thus not as good for demonstrations. The screen shots of the interface in this article use this database. Details of the database can also be seen in Table 1.

The second database used contains all articles of the RSNA journals Radiology and Radiographics with a total of 77'000 images. The database was used in ImageCLEF 2010[15] for the evaluation of textual and visual retrieval and in this respect it is possible to compare the performance with other state–of–the–art tools.

2.3. Evaluation methodology

The evaluation of the tools in this paper is basically based on the evaluations performed in the context of the ImageCLEF benchmark, where 17 research groups compared their tools. This gives a clear idea of the qualities and the shortcomings of the pure retrieval performance of the two employed tools. Still, in past studies it was frequently shown that the user interfaces are actually more important for obtaining a good result for a user than the pure system performance.[21]

‡‡http://www.cliqcliq.com/

Table 1: Indexed open access articles from BioMed Central.

Number of journals	24
Number of articles	9'403
Min. number of articles per journal	16
Max. number of articles per journal	2'495
Average number of articles per journal	392
Total number of images (after a cleaning step)	37'940
Min. number of images per journal	28
Max. number of images per journal	13'618
Average number of images per journal	567
Average number of images per article	4
Size of all images total	2.81 gigabytes (GB)
Average image size	77 kilobytes (KB)

3. RESULTS

This section describes the main results of the tools developed in this project. Main results are the development of the prototype described by using the various screenshots.

3.1. Integration of Components

The system described in this paper consists of a mobile client to an existing medical information access system accessing the medical open access literature. The system is web–based and can thus work on several mobile platforms although this article concentrates on the use with an iPhone and all screenshots show this operating system. General problems with image upload from the local camera (usually blocked for web access) for visual similarity search were solved using the framework cliqcliq. This framework can be called from a web application and allows to take a picture with the internal camera and return it back to the web application in our case to start a visual similarity retrieval.

This basic structure of the developed system can be seen in Figure 3, where three main layers are separated. On the right the mobile client can be seen as basically a browser–based interface that allows using all specificities of the smart phone (including the camera via cliqcliq). In the middle, web standards such as HTML (HyperText Markup Language), CSS (Cascading Style Sheets) and forms are used to transmit data to the middleware. The middleware then transmits queries to MedSearch via a simple web service, receives the results and delivers these back to the client.

To make the system available a simple browser identification was used with the standard MedSearch URL (Uniform Resource Locator). This means that a detection of a mobile browser with a small resolution automatically directs the user towards the mobile instead of the fixed interface.

3.2. Data Display

A systematic review of best practices for mobile interfaces and more particularly mobile information retrieval interfaces was performed. Several screen layouts were tested to allow using the available space in the best possible manner.

Several screen shots and a navigation adapted for touch screens are presented in Figure 4. The first screenshot shows the textual search interface when entering a query term. The second screen shows the results presentation of several articles including the first up to three images. The results presented included the title, the date of publication, the journal, the first authors, and the beginning of the abstract, allowing for a good overview. The third screen then gives a view of a single article with more images and a larger part of the abstract including links to the full text. The browser and operating system are detected by the application and the corresponding interface

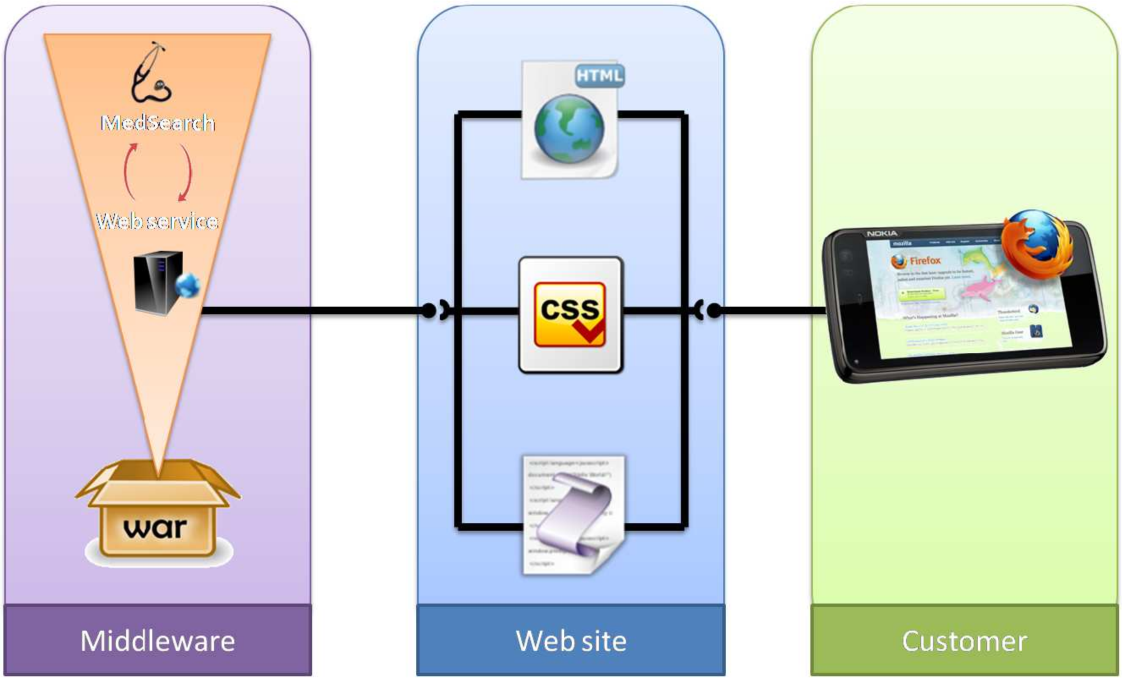

Figure 3: The three main building blocks of the architecture for creating a mobile search system from MedSearch.

(a) Search interface allows to type queries on screen.

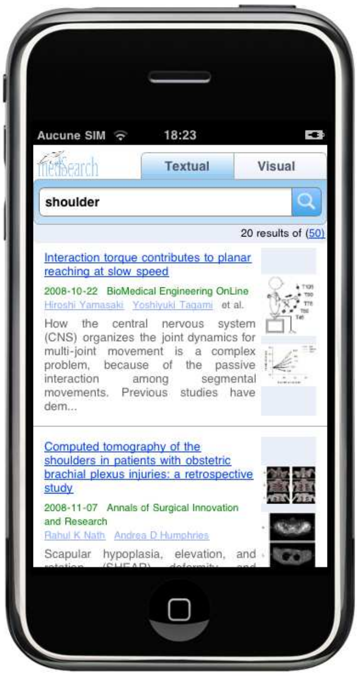

(b) Results set showing text and images in acceptable size.

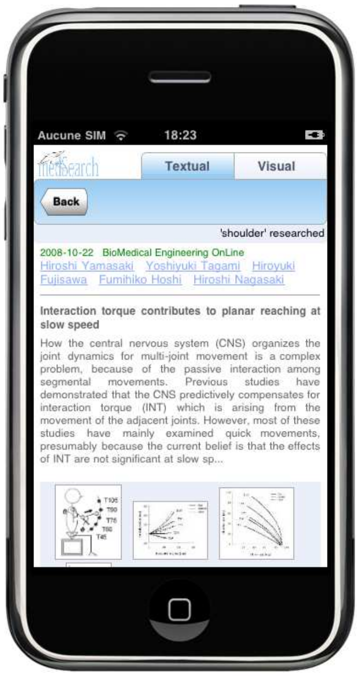

(c) Single articles can be viewed in more detail including images.

Figure 4: Examples for the interface showing the interface, a text query with results of images in an article and a single article in more detail.

(a) Images of an article with link to the web and PDF file.

(b) Browsing images of an article in larger size.

(c) Query with an image of a hand taken with the built–in camera.

Figure 5: Details of an executed results set of a query.

is then used instead of the standard MedSearch interface automatically. Navigation is via the touchscreen and very quick for browsing even long results lists.

Figure 5 shows several screens that display details of a query. In Figure 5a all images of a chosen article are shown. Links to the web page of the article as well as to the PDF exist. A different presentation of the same data is given in Figure 5b, where an image is shown full screen and the touch screen mechanism for advancing to the next image can be used efficiently. From all images it is possible to start a visual similarity query. An example for a similarity query is also shown in Figure 5c. In this example the query is taken with the built–in camera, in this case with a picture of a hand. For starting similarity queries it is also possible to request a random set of images that can then be used as query start.

Figure 6 shows that also the direction detection of the iPhone or other device can be used and that the screen is automatically adapted to this. Such a view change can help viewing pictures in particular orientations or also show the text in long lines better.

The final screenshot shows the same application but on an iPad (Figure 7). This shows that many more images can be shown in better quality with such a larger screen. In the same way as for the iPad the mobile query system can be used on other devices such as Android smart phones with a large variety of screen sizes.

3.3. Performance evaluation

No user test or direct quantitative evaluation of the retrieval quality of MedSearch Mobile was carried out on the database of open access journals but a performance comparison based on a similar database of the medical literature exists in the context of the international ImageCLEF benchmark.[15,22] In this context, 77'500 images from the scientific literature in radiology are used. 30 search tasks were developed and assessed by clinicians. 17 research groups worldwide participated in the medical task in 2009 and 16 in 2010.

GIFT is among the average or slightly above average systems for purely visual retrieval with a mean average precision (MAP) of 0.025 in 2009 and 0.023 in 2010.[15,23] For case–based tasks it was the best system in 2009 and

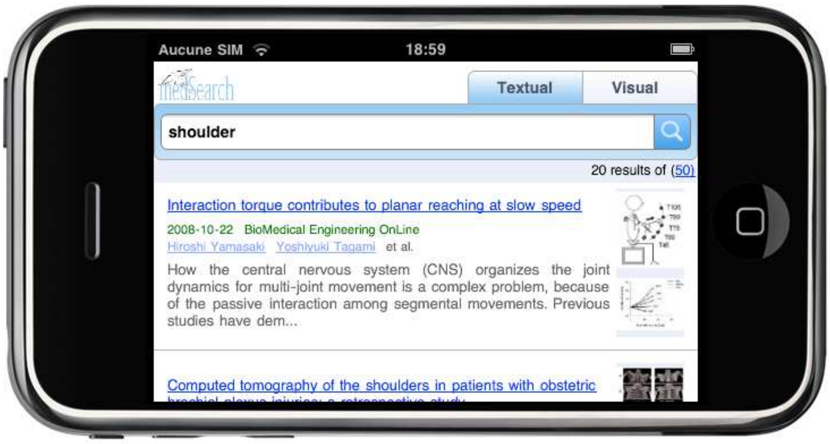

Figure 6: The system automatically adapts to the orientation and thus allows to change this depending on the content shown.

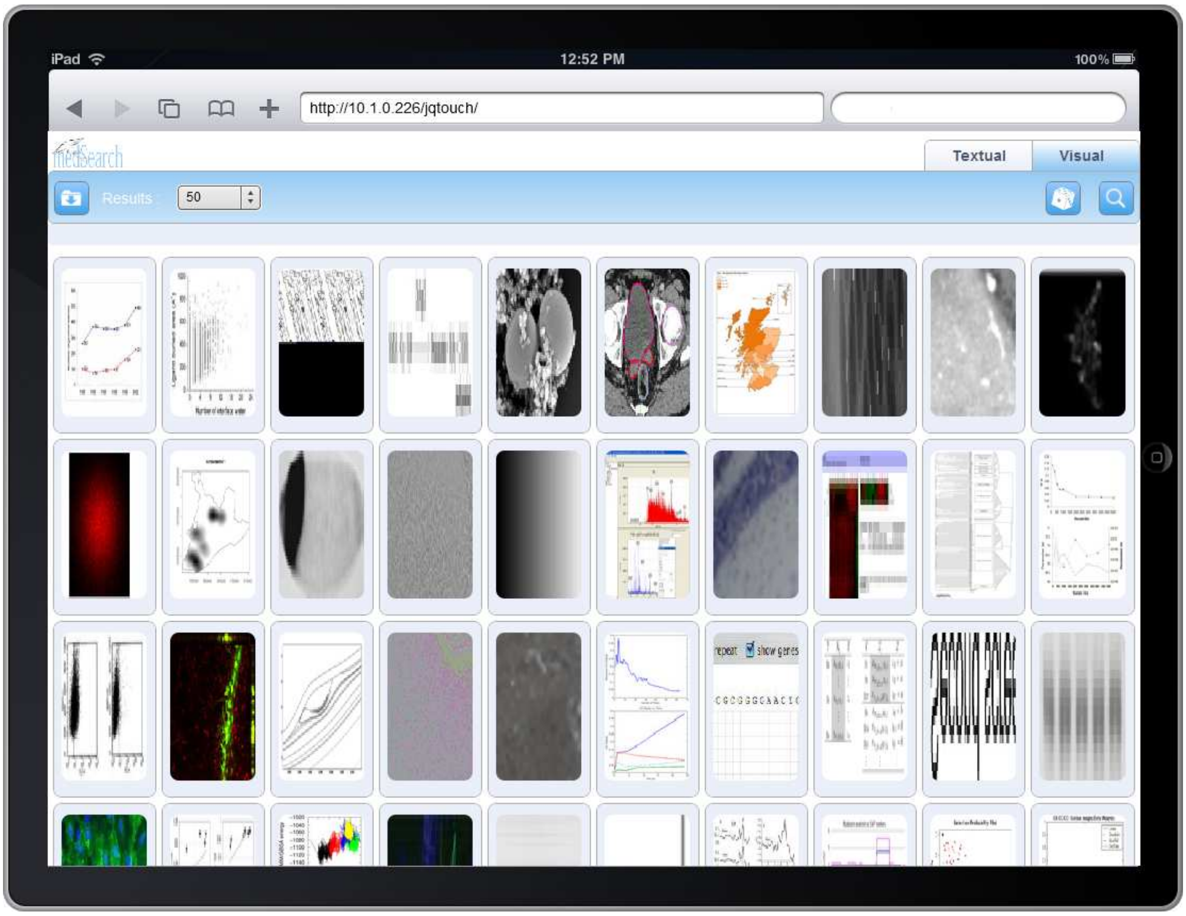

Figure 7: Version for an iPad with a larger screen that allows to show a significantly larger number of images.

2010 with a MAP of 0.0358 in 2010. Lucene was among the the very good systems for purely textual retrieval in 2009 (MAP of 0.27) and among the average in 2010 (MAP of 0.26). It had the highest early precision several times.[24] For the case–based topics Lucene had the third best performance in 2010. It is clear that purely textual retrieval quality is much better than the visual quality. Combinations between textual and visual retrieval using GIFT and Lucene led to a better early precision but a slight decrease in MAP. All this shows that the retrieval quality is sufficient for the tools to be integrated into this mobile application.

4. DISCUSSIONS AND CONCLUSIONS

Mobile applications in general and mobile information access in particular are growing fields of research. Mobile clients become reliable and lightweight and allow most desktop applications to be used despite the limited screen size and resolution (that is actually not too limited anymore). On the other hand there are currently a large number of platforms and mobile operating systems, so developing applications for one platform most often means that the result is not compatible with the development for another platform. Web applications on the other hand allow access via standard technologies. For most current platforms this works well albeit small problems such as image capture with the internal camera that is disabled for web access on an iPhone or Android, for example. To solve this problem native applications can be integrated that exist for several platforms.

This paper has presented the adaptation of a combined textual and content–based visual retrieval system for medical data from the literature of BioMedCentral. Both, visual and textual search have been adapted to the mobile environment and the small screen size and allow for an efficient work. The response times are in the area of 1 second with the current database, and the data to be communicated have been reduced as much as possible to follow these requirements.

The current system works well but is still far from perfect. It is particularly required to perform real user studies to find our what the required usability options are and how clinicians would like to use such a system and in which situations the added value would be highest. It seems also necessary to integrate semantic technologies such as terminologies and medical knowledge bases into the textual and visual retrieval systems. This is one of the goals of the Khresmoi research project that started at the end of 2010. All in all we can see that mobile information search is an extremely important domain for making medical information management more effective. Mobile medical applications will gain ground and such tools should include textual search as well as visual similarity search to navigate in databases.

5. ACKNOWLEDGEMENTS

This work was partially supported by the University of Applied Sciences Western Switzerland (HES–SO). We would also like to thank the Swiss National Science Foundation (MANY project) and the European Union (Chorus+, Promise, Khresmoi projects with grant agreements 249008, 258191 and 257528) for their partial support of this work.

REFERENCES

1. J. L. Gómez-Barroso, R. Companñó, C. Feijó, M. Bacigalupo, O. Westlund, S. Ramos, A. Jaokar, F. Álvarez, R. De Waele, G. Mateos-Barrado, and M. C. García-Jiméne, "Prospects of mobile search," Tech. Rep. 24148 EN, European Commission, Joint Research Center, Sevilla, Spain, March 2010.
2. E. Ammenwerth, A. Buchauer, B. Bludau, and R. Haux, "Mobile information and communication tools in the hospital," *International Journal of Medical Informatics* **57**(1), pp. 21–40, 2000.
3. J. C. Cortizo, D. Gachet, M. de Buenaga, M. J. Maña, and M. de la Villa, "Mobile medical information access by means of multidocument summarization based on similarities and differences," in *ACL-2008 Workshop on Mobile Language Processing*, (Columbus, Ohio, United States), June 2008.
4. P. G. B. Enser, "Pictorial information retrieval," *Journal of Documentation* **51**(2), pp. 126–170, 1995.
5. A. W. M. 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**, pp. 1349–1380, December 2000.

6. R. Datta, D. Joshi, J. Li, and J. Z. Wang, "Image retrieval: Ideas, influences, and trends of the new age," *ACM Computing Surveys* **40**, pp. 1–60, April 2008.
7. D. G. Lowe, "Object recognition from local scale invariant features," in *Proceedings of the International Conference of Computer Vision*, (Corfu, Greece), September 1999.
8. J. Sivic and A. Zisserman, "Video Google: Efficient visual search of videos," in *Toward Category–Level Object Recognition 2006*, pp. 127–144, 2006.
9. M. Everingham, A. Zisserman, C. K. I. Williams, L. van Gool, M. Allan, C. M. Bishop, O. Chapelle, N. Dalal, T. Deselaers, G. Dorko, S. Duffner, J. Eichhorn, J. D. R. Farquhar, M. Fritz, C. Garcia, T. Griffiths, F. Jurie, D. Keysers, M. Koskela, J. Laaksonen, D. Larlus, B. Leibe, H. Meng, H. Ney, B. Schiele, C. Schmid, E. Seemann, J. Shawe-Taylor, A. Storkey, S. Szedmak, B. Triggs, I. Ulusoy, V. Viitaniemi, and J. Zhang, "The 2005 PASCAL visual object classes challenge," in *Machine Learning Challenges. Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Tectual Entailment (PASCAL Workshop 05)*, *Lecture Notes in Artificial Intelligence*(3944), pp. 117–176, (Southampton, UK), 2006.
10. P. Clough, H. Müller, T. Deselaers, M. Grubinger, T. M. Lehmann, J. Jensen, and W. Hersh, "The CLEF 2005 cross–language image retrieval track," in *Cross Language Evaluation Forum (CLEF 2005)*, *Springer Lecture Notes in Computer Science*, pp. 535–557, September 2006.
11. H. D. Tagare, C. Jaffe, and J. Duncan, "Medical image databases: A content–based retrieval approach," *Journal of the American Medical Informatics Association* **4**(3), pp. 184–198, 1997.
12. T. M. Lehmann, M. O. Güld, C. Thies, B. Fischer, K. Spitzer, D. Keysers, H. Ney, M. Kohnen, H. Schubert, and B. B. Wein, "Content–based image retrieval in medical applications," *Methods of Information in Medicine* **43**, pp. 354–361, 2004.
13. H. Müller, N. Michoux, D. Bandon, and A. Geissbuhler, "A review of content-based image retrieval systems in medicine–clinical benefits and future directions," *Internation Journal of Medical Informatics* **73**, pp. 1–23, February 2004.
14. A. M. Aisen, L. S. Broderick, H. Winer-Muram, C. E. Brodley, A. C. Kak, C. Pavlopoulou, J. Dy, C.-R. Shyu, and A. Marchiori, "Automated storage and retrieval of thin–section CT images to assist diagnosis: System description and preliminary assessment," *Radiology* **228**, pp. 265–270, July 2003.
15. H. Müller, J. Kalpathy-Cramer, I. Eggel, S. Bedrick, R. Said, B. Bakke, C. E. Kahn Jr., and W. Hersh, "Overview of the CLEF 2010 medical image retrieval track," in *Working Notes of CLEF 2010 (Cross Language Evaluation Forum)*, September 2010.
16. C. E. Kahn Jr. and C. Thao, "Goldminer: A radiology image search engine," *American Journal of Roentgenology* **188**, pp. 1475–1478, 2008.
17. T. M. Deserno, S. Antani, and R. Long, "Exploring access to scientific literature using content–based image retrieval," in *SPIE Medical Imaging*, **6516**, (San Diego, CA, USA), February 2007.
18. I. Eggel and H. Müller, "Indexing the medical open access literature for textual and content–based visual retrieval," in *MEDINFO 2010 — Proceedings of the 13th World Congress on Medical Informatics*, **160**, pp. 1277–1281, IOS Press, September 2010.
19. D. M. Squire, W. Müller, H. Müller, and T. Pun, "Content–based query of image databases: inspirations from text retrieval," *Pattern Recognition Letters (Selected Papers from The 11th Scandinavian Conference on Image Analysis SCIA '99)* **21**(13–14), pp. 1193–1198, 2000. B.K. Ersboll, P. Johansen, Eds.
20. "Mobile OS/browser report," tech. rep., jTribe, January 2010.
21. W. Hersh, *Information Retrieval — A health and Biomedical Perspective*, Springer, 2003 – second edition.
22. H. Müller, J. Kalpathy-Cramer, C. E. Kahn Jr., W. Hatt, S. Bedrick, and W. Hersh, "Overview of the ImageCLEFmed 2008 medical image retrieval task," in *Evaluating Systems for Multilingual and Multimodal Information Access – 9th Workshop of the Cross-Language Evaluation Forum*, C. Peters, D. Giampiccolo, N. Ferro, V. Petras, J. Gonzalo, A. Peñas, T. Deselaers, T. Mandl, G. Jones, and M. Kurimo, eds., *Lecture Notes in Computer Science* **5706**, pp. 500–510, (Aarhus, Denmark), September 2009.
23. H. Müller, J. Kalpathy-Cramer, I. Eggel, S. Bedrick, R. Said, B. Bakke, C. E. Kahn Jr., and W. Hersh, "Overview of the CLEF 2009 medical image retrieval track," in *Working Notes of CLEF 2009*, (Corfu, Greece), September 2009.

24. J. Kalpathy-Cramer, S. Bedrick, W. Hatt, and W. Hersh, "Multimodal medical image retrieval OHSU at ImageCLEF 2008," in *Evaluating Systems for Multilingual and Multimodal Information Access*, C. Peters, T. Deselaers, N. Ferro, J. Gonzalo, G. Jones, M. Kurimo, T. Mandl, A. Peñas, and V. Petras, eds., *Lecture Notes in Computer Science* **5706**, pp. 744–751, Springer Berlin / Heidelberg, 2009.