Florentine Coats of Arms on the Web:

Experimenting with Retrieval Based on Text or Image Content*

Fabrizio Falchi [1], Fausto Rabitti[1],
Werner Schweibenz[2] & Jan Simane[2]

[1] ISTI-CNR, Pisa, Italy
[2] Kunsthistorisches Institut in Florenz Max Planck Institut, Florence, Italy

Abstract
The paper describes two different ways of accessing a collection of Florentine coats of arms, which was digitised and made available on the Web. One way is a traditional approach: Textual description and indexing. As heraldry uses a specific and complex terminology, the textual approach works well for experts but is insufficient for laypersons. Therefore an additional retrieval mode was chosen. It is based on the visual specification by users of the content of images contained in the coat of arms. This access mode allows users who are not experts in heraldry to access the collection.

1 Introduction

The *Kunsthistorisches Institut in Florenz (KHI) Max-Planck-Institut* has possessed more than 2,800 coloured drawings of Florentine coat of arms since 1901. This unique collection is a very useful instrument for identifying coats of arms on palace facades, tomb monuments, altar pieces, and so on. The high importance and usefulness of this collection were the motives for the decision to digitise the material and make it accessible on the Web (the current version is a prototype and proof of concept, available online at *http://wappen.khi.fi.it/* and *http://www.nmis.isti.cnr.it/ KHI/*). Another reason was to improve the access to the collection, which was indexed only rudimentarily on index cards. In addition to traditional textual indexing of the coats of arms, we wanted to try out retrieval based on image content because we considered the heraldic terminology to be a problem for laypersons. This claim can

* Published in: Osswald, Achim; Stempfhuber, Maximilian; Wolff, Christian (eds.) (2007). Open Innovation. Proc. 10th International Symposium for Information Science. Constance: UVK, 1-13.

be illustrated by a fairly simple example of a blazon: "Gules a Griffin with Dragon Wings, Tail and Tongue rampant Or armed, beaked, langued and membered Azure between four Roses Argent."

In this application context, we decided to explore the possibility of allowing the users to search the content of images representing the coat of arms using visual specification of this content. We believe that a very important move towards this kind of solution is to support *content-based retrieval* from the image database using *feature-based similarity access*. A feature (or *content-representative metadata*) is a set of characteristics of images, such as colour, texture and shapes. Similarity-based access means that the user specifies some characteristics of the wanted images, and the system retrieves the most relevant images with regard to the given characteristics, that is, the images *most similar* to the query. This approach assumes the ability to measure the distance (using metrics of some kind) between the query and the database images. This search paradigm:

Generalizes the information retrieval approach for textual archives, where search is based on key terms and the retrieved documents can be ranked with regard to the relevance of these key terms in the document text;

Generalizes the Web searching mechanism, as used by modern commercial systems, where searching is based on heuristics of links and domains and ranking is based on adapted text information retrieval techniques.

1.1 Description of the Material

The collection consists of two parts, on the one hand the coats of arms of Florentine families and on the other hand those of Florentine churches, hospitals and confraternities. The biggest part of the collection comprises the coats of arms of noble families, 2,512 colour drawings and 22 sketches. These drawings, which were done by an anonymous draughtsman in the 19th century, were acquired by the Institute library in 1901. A much smaller number of drawings refer to the coats of arms of Florentine churches, hospitals and confraternities. The 317 colour drawings were created by Otto Wenzel, the KHI's librarian from 1902 to 1915, based on the Priorista di Luca Chiari.

In the library's inventory of 1901, the drawings are described as "Sammlung von Wappen florentiner Familien u a." [Collection of coats of arms of Florentine families and so on], while the other part is first mentioned in the inventory of 1909 as "Sammlung von Wappen florentiner Kirchen, Hospitälern und Bruderschaften

nach Priorista di Luca Chiari" [Collection of coats of arms of Florentine churches, hospitals and brotherhoods according to the Priorista di Luca Chiari].

Each coat of arms is drawn as a triangular shield in the centre of a sheet of light brown paper of approximately 16 x 23 cm; some of the sheets have watermarks. Most of the drawings are made in pencil and watercolour, a few of them are sketched only in pencil. Below each coat of arms, the name of the family is written in black ink. In some cases, commentaries in pencil are added. The whole collection is made accessible using a handwritten card index. This index combines the names of the families with the corresponding heraldic elements. For decades, both parts of the collections were stored in three boxes, one of them containing the index. During the 1980s, the sheets were stored in 65 black ring binders while the index cards remained in boxes.

The collection has been frequently consulted by scholars and researchers to identify unknown coats of arms on buildings and artworks or for looking up a certain family's coat of arms. In this way, it has become an important resource for art historians.

1.2 Reasons for the Digitisation

For several reasons, the KHI decided to digitise this collection and make it accessible on the Web. Firstly, there were conservation issues because over time the drawings had become tattered to a certain extent due to extensive use. Therefore it seemed reasonable to digitise the material and provide access to digital facsimiles instead of the originals. Secondly, there were restrictions in physical access. The original drawings are kept in a special room with restricted office hours. The goal was to increase accessibility on-site and extend it to the public on the Web. The third reason was to improve access to the collection and the retrieval of individual coats of arms.

1.3 Project Description

The project started in 2005. The first step was to provide a formal verbal description (blazon) of each individual coat of arms in German, according to the international standards of heraldry. As heraldry applies a complex specialist terminology, the KHI decided not to rely on heraldic manuals such as Spreti [Spreti 1928-1936] but to seek the advice of a specialist. Harald Drös, Head of the Research Institute for German Inscriptions at the Heidelberg Academy of Science, took on the responsibility of supervising this work. The information from the heraldic blazon for each coat of arms was entered in the KHI's art information database (HiDA). In the database, a variety of indices were created for relevant fields, such as family names,

colours, heraldic ordinaries and charges and so on, in order to allow combined searches for the content (see figure 1).

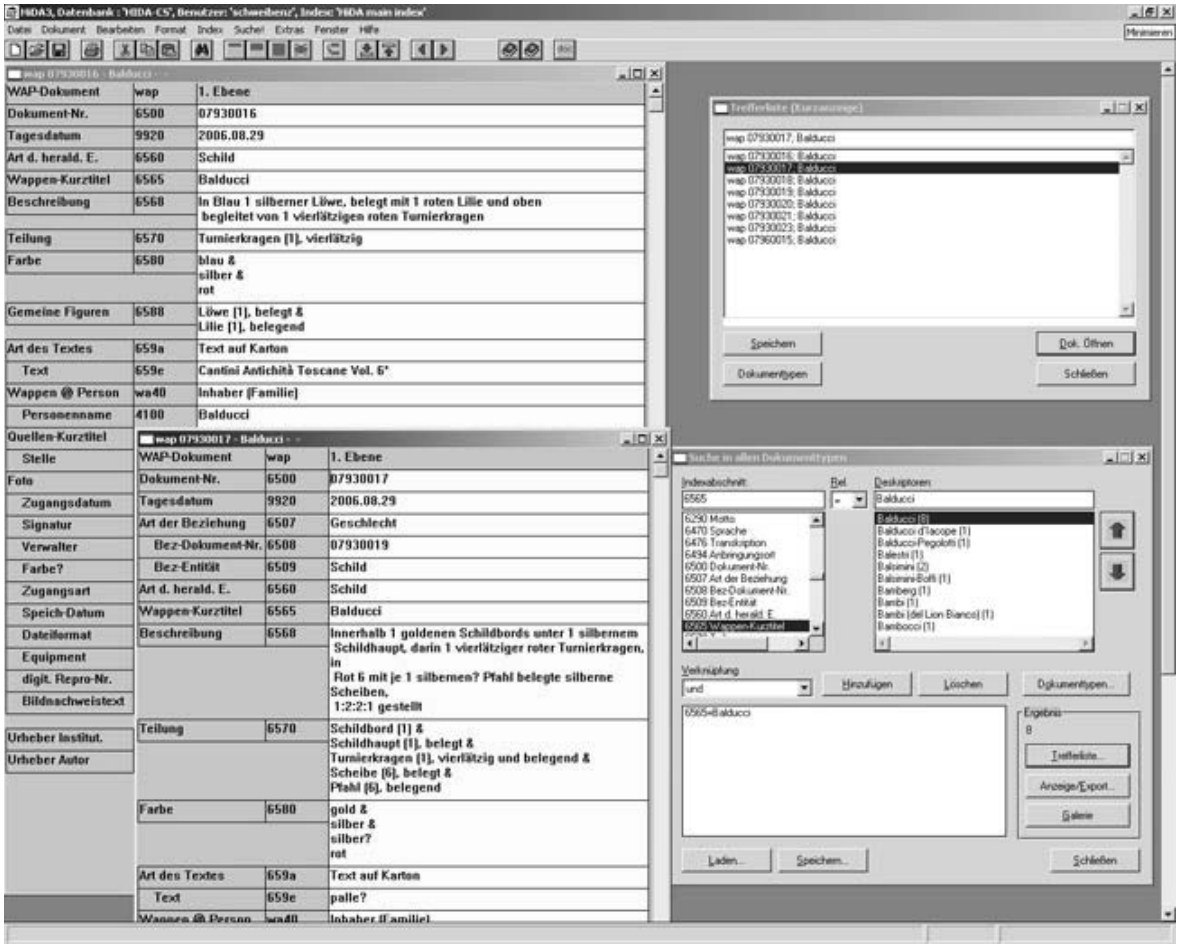

Figure 1: A heraldic blazon for a coat of arms in the HiDA database (Family name: Balducci; Heraldic charge: Lion)

In parallel, all the coats of arms were scanned in-house. The master scans were produced on a high-end scanner in RGB colour, with 16-bit colour depth and 600 pixels per inch. The height of each scan was approximately 5,500 pixels, the width approximately 4,000 pixels. This resulted in a file size of about 125 MB for the master. Each master was reduced to a working copy in lower quality with 8-bit colour depth and 300 pixels per inch for further use (for example, print). For the display on the Web, JPEGs with a width of 350 pixels are used.

The content of the art information database and the digitised images were integrated in an open-source Web database (ZOPE), which was provided by the KHI's technical partner, Fafalter GmbH in Düsseldorf. The same product has been successfully used for the KHI's digital photo library [Bieber, Schweibenz 2004, 2005]. In addition to the traditional text-based searches, the KHI decided to explore new means of access. The idea was to analyse the pictorial content of the individual coats

of arms and provide a query by image content. In order to implement this feature, the KHI looked for a partner with profound competence in this field. This was the *Istituto di Scienza e Tecnologie dell'Informazione/Consiglio Nazionale di Ricerca* (ISTI/CNR) in Pisa.

2 Text-Based Retrieval in the Coat of Arms Database

The text-based retrieval relies on the indices created in the art information database HiDA (Hierarchical Data Administrator) in German. By entering information in certain HiDA fields, a controlled vocabulary for different subjects is set up, for example, for tinctures (colours) such as "gold", "silver" or "blue", for heraldic ordinaries such as "bend", "cross" or "chief" or for heraldic charges such as "lion" or "eagle". Other fields in HiDA, such as the blazon information or the text information contained on the sheet of paper, are automatically indexed in full-text form. All indices are transferred to the ZOPE database and can be searched by entering a term in the search field on the Web interface (see figure 2) or in the advanced search (see figure 3).

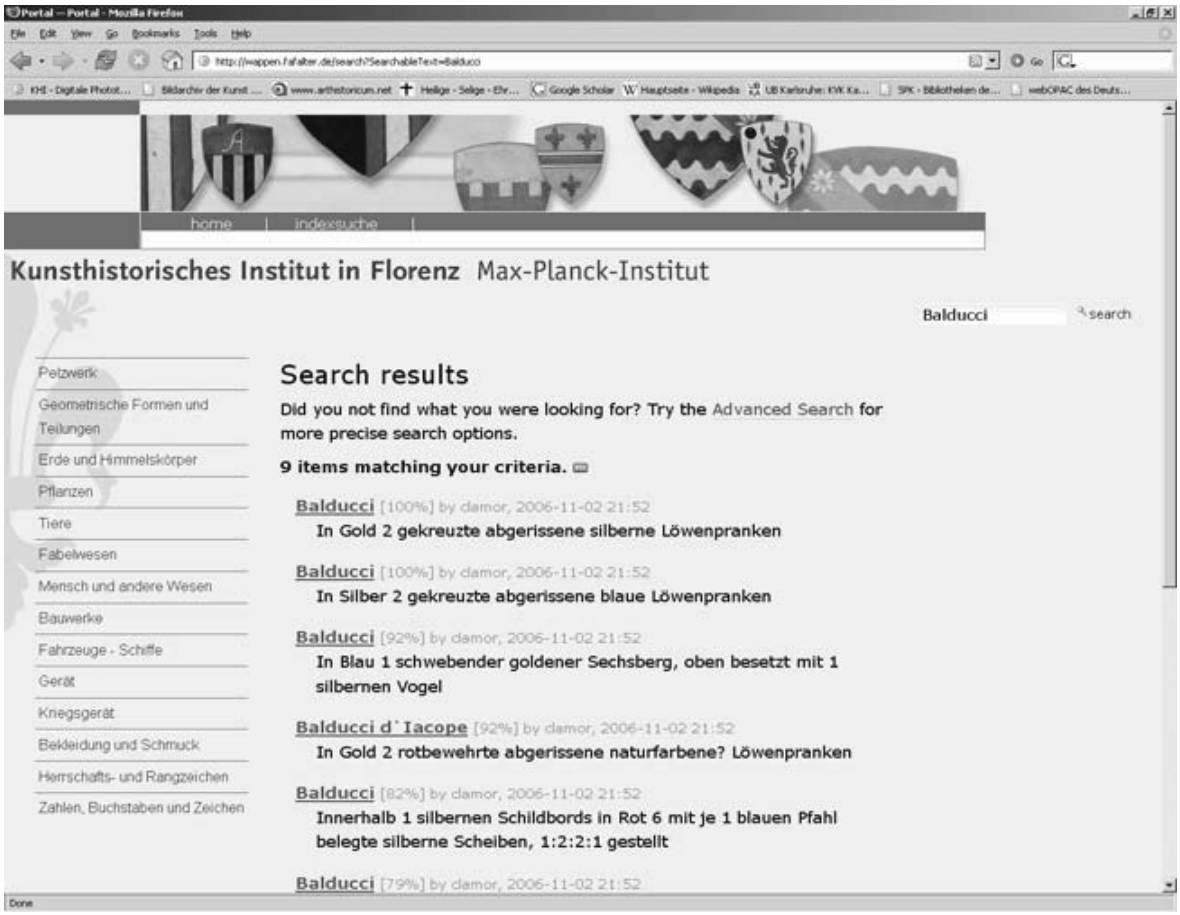

Figure 2: Live search for a heraldic charge (lion)

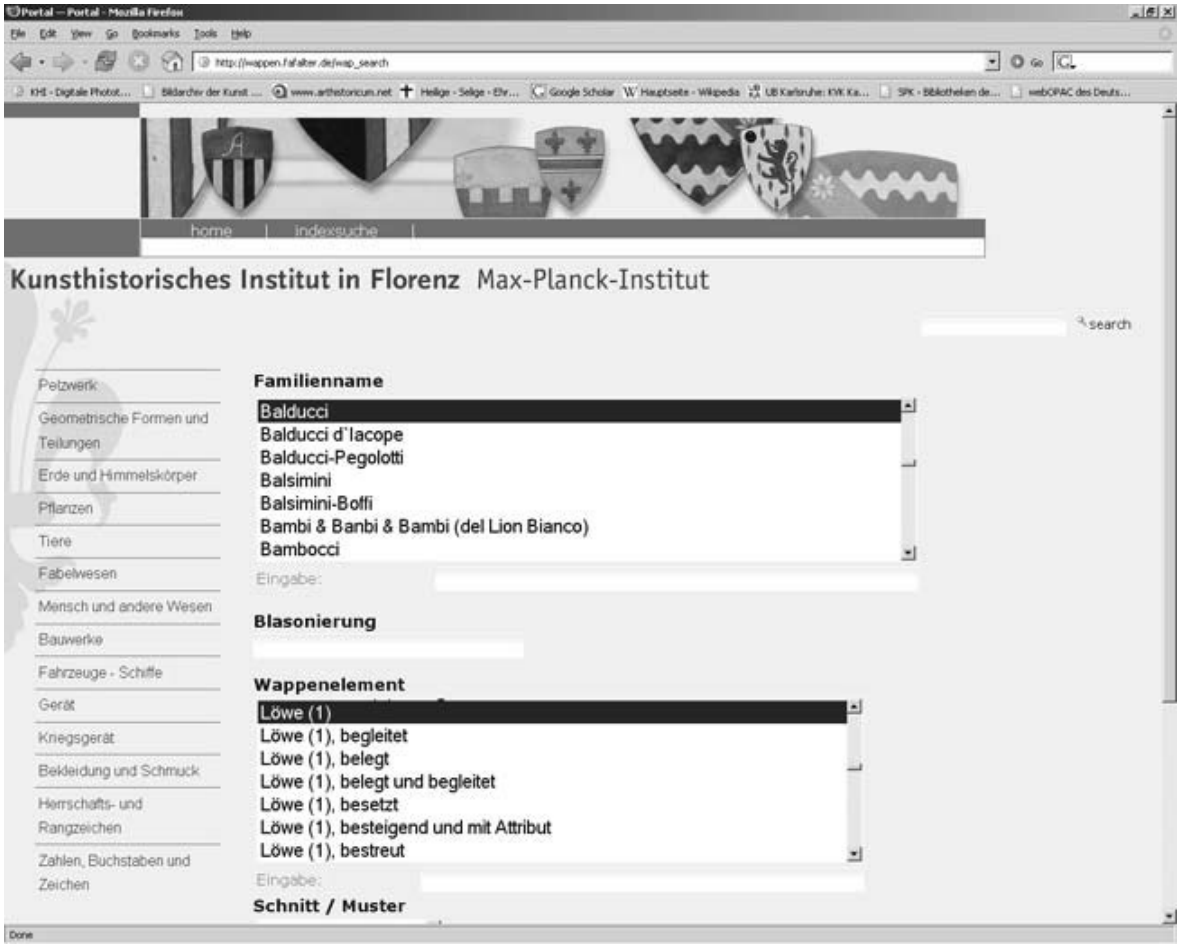

Figure 3: Combined search for a family name (Balducci) and a heraldic charge (lion)

As already mentioned above, the terminology of heraldry is very complex. Therefore we cannot expect the users to be familiar with all the specific terms. For this reason, apart from the simple and advanced search, a number of categories were added. These categories, based on a German heraldry book called "Wappenbilderordnung" [Wappenbilderordnung 1986], for example "animals," "plants" or "mythical creatures", allow users to approach the coats of arms starting from very general terminology and getting more and more specific.

Each category branches out, for example, animals into mammals and birds, and finally leads to a list of family names with an accompanying blazon. The categories are based on the indices provided by HiDA and are set up by the Content Management System PLONE using a service called "smart folders". This allows users who are non-specialists to access the coats of arms in a convenient way without having to know or use heraldic terminology. Another means of access for non-specialists is the query by image content.

3 Content-Based Searching in the Coats of Arms Image Database

We had different options for implementing the retrieval application based on the visual specification by users of the content of images represented in the coat of arms.

The earliest and most common approach is the global query-by-example paradigm. It consists of retrieving images whose visual appearance is globally similar to a selected example image. Initially proposed by Swain and Ballard [Swain, Ballard 1999], it was adopted by a vast majority of content-based image retrieval systems (for example, IBM QBIC, VisualSEEk, Virage's VIR Image Engine and Excalibur's Image RetrievalWare). However, this paradigm has limited application in our case, since it would allow users to retrieve only coats of arms similar to the one specified.

The partial query-by-example paradigm was introduced later on. This approach allows the user to explicitly select a visual component which is relevant for the query and retrieves images which contain a similar visual component. This approach proved to be more selective, hence more precise than the global query-by-example paradigm [Ma, Manjunath 1999]. This is really what is needed for our application: Searching images of coats of arms based on a visual element which they contain.

In our approach, based on a similarity search of images, it is possible to apply the relevance feedback mechanism, inspired by text retrieval, to refine the image search [Huang, Rui, Mehrotra 1997]. Among the retrieved images, the user specifies which ones are relevant and reiterates the search. By refining the similarity measure, the searched image can be reached more efficiently, taking into account the subjective preferences of the user.

A fundamental decision regarding the approach taken is the adoption of the MPEG-7 standard to represent the visual metadata, that is, the features extracted from the coat of arms images. MPEG-7 [ISO/IEC 15938], also called the Multimedia Content Description Interface, aims to cover the need for searching and retrieving multimedia information by describing the content of audiovisual objects in a standardised way. The MPEG-7 approach is to specify a set of standardised *descriptors*, that is, entities that contain the syntax and semantics of audiovisual content. For example, a descriptor could refer to the colour of an object by specifying the name of the attribute colour (such as, Color) and the type of the value (such as, a string or three integer values for RGB colours).

3.1 Image Analysis and Segmentation

Before using the images of the database in the retrieval process, it is necessary to analyse their content and identify the relevant components. This process is called segmentation.

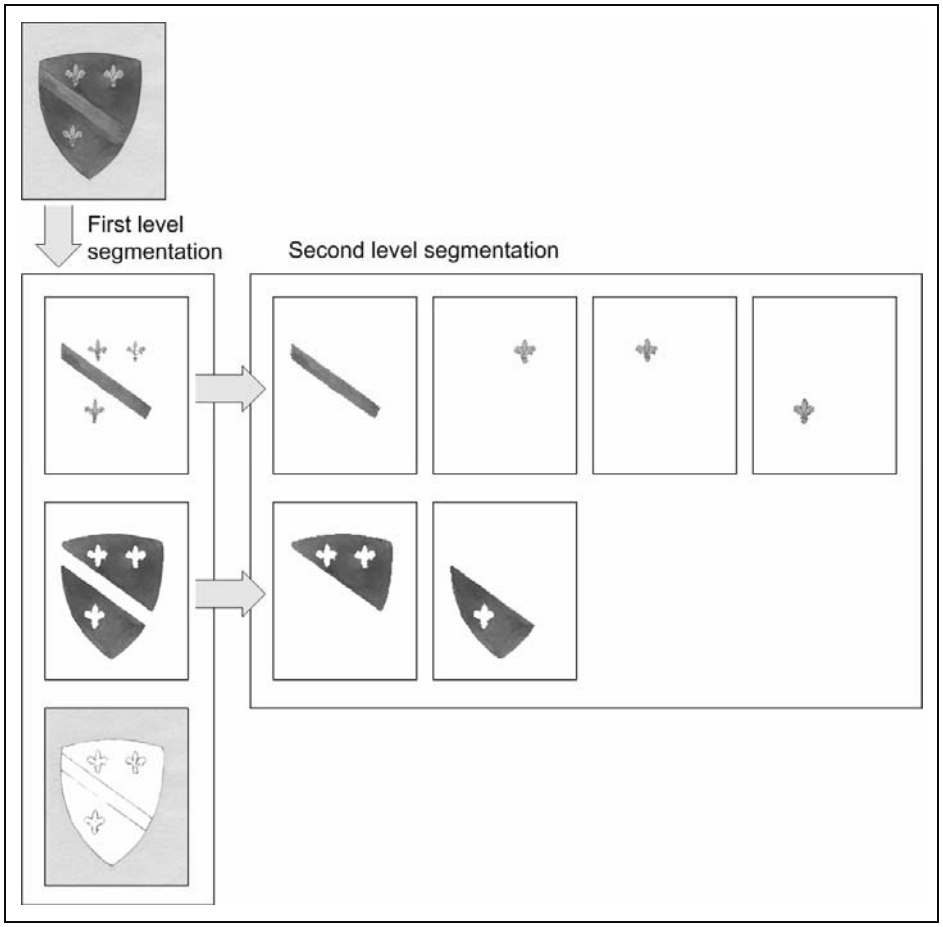

Figure 4: Two levels of segmentation

Image segmentation refers to partitioning the image into homogeneous regions. We decided to use two levels of segmentation using two different segmentation algorithms. Basically, the first one divides the image into regions based on their colour gradient. The results of this first-level segmentation are not necessarily connected regions of similar colour. These regions are then segmented again using another algorithm based on K-Means, which is an algorithm to cluster objects based on attributes trying to minimize total intra-cluster variance. The results of this second-level segmentation are connected regions. All regions obtained by both segmentation algorithms are used.

For the first-level segmentation, we used an implementation of the algorithm described in [Comaniciu, Meer 1997], developed by Dorin Comaniciu. The tech-

nique is based on the mean shift algorithm, a simple nonparametric procedure for estimating density gradients.

For the second-level segmentation, we used a novel variant of the well known K-Means-with-connectivity-constraint algorithm (KMCC), a member of popular K-Means family, which is described in [Mezaris, Kompatsiaris Strintzis 2002, 2004]. The KMCC algorithm classifies the pixels into regions taking into account not only the intensity information associated with each pixel but also the position of the pixel, thus producing connected regions rather than sets of chromatically similar pixels.

3.2 MPEG-7 Descriptors Extraction and Searching

Feature extraction was performed using an application we built based on the MPEG-7 [ISO/IEC 15938] experimentation model [ISO/IEC 15938-6:2003] of MPEG-7. The software can extract all MPEG-7 image visual descriptors [ISO/IEC 15938-3:2003]. The same application has been used on MILOS [MILOS, Amato, Gennaro, Rabitti, Savino 2004] and particularly on PhotoBook [Amato, Bolettieri, Debole, Falchi, Rabitti, Savino 2006, Amato, Bolettieri, Debole, Falchi, Rabitti Savino 2006], which is a Multimedia Digital Library Application (*http://milos.isti.cnr.it*) that we built on top of MILOS for online photo sharing.

For both complete coat of arms images and automatically selected regions, we extracted the following MPEG-7 descriptors: *ScalableColor*, *ColorStructure*, *ColorLayout*, *DominantColor*, *EdgeHistogram* and *HomogeneousTexture*.

Because of the fact that they are concerned with shape, 2 MPEG-7 descriptors were extracted only for regions: *RegionShape* and *ContourShape*.

Two more descriptors for colour were used for complete images: *ColorStructure* and *ColorLayout*. These descriptors were not used for regions because they are meaningless for non-rectangular images.

We now provide a brief description of the MPEG-7 descriptor used. *ScalableColor*, *ColorStructure*, *ColorLayout* and *DominantColor* are all to do with colours. *ScalableColor* is a colour histogram in the HSV Colour Space. *ColorStructure* captures both colour content and information about the spatial arrangement of the colours. *ColorLayout* represents the spatial layout of colour images. *EdgeHistogram* describes the spatial distribution of five types of edges. *HomogeneousTexture* characterizes the properties of texture in an image. *RegionShape* expresses pixel distribution within a region. *ContourShape* is based on the Curvature Scale-Space representation of the contour.

The result of the extraction process is an XML document. The online application developed for this project performs a linear scan of the descriptors present in the document. Given the limited number of images in the database, just a few thousand, the use of an index structure for the similarity search is not necessary and a linear search of the MPEG-7 descriptors is feasible. However, in MILOS we use M-Tree [Salembier, Sikora 2002] as an index structure for the similarity search: This allows the system to be scalable, searching tens of million of MPEG-7 image descriptors in a few seconds.

To compare the descriptors resulting from two images, we used the functions proposed in [Swain, Ballard 1991]. If the user wants to search using a combination of descriptors, distances are combined, single-feature results are accumulated using predefined (heuristic) weights.

3.3 Online Image Content Search Application

Once an image is selected (for example, from the textual description database), the application asks the user to select the region (or even the complete image) that he or she wants to use as a query to search for similar ones. The first 10 results are listed in order of decreasing similarity using *EdgeHistogram* and *RegionShape*, which we found to be the most useful descriptors for this kind of images (see fig. 5).

From the results page, it is possible to select the descriptors to be used for the search and then ask for the relative results. When the user clicks a result region, the search is performed again using the selection as the query region.
Below each results region, there is a small image of the entire coat – of arms which the region comes from. Using the *Info* link located above each region present on the results page, the user can access the textual metadata in the database. On the upper left side, the results page also reports the query coat-of-arms and the region used as the query.

4 Conclusion and Future Research

The project combines two different ways of accessing a collection of Florentine coats of arms, one being traditional and text-based, the other being experimental image searching-by-content. With this combination, we hope to satisfy the needs of both experts and laypersons in the field of heraldry as experts can use the heraldic terminology for a text-based search while laypersons can use the content-based search.

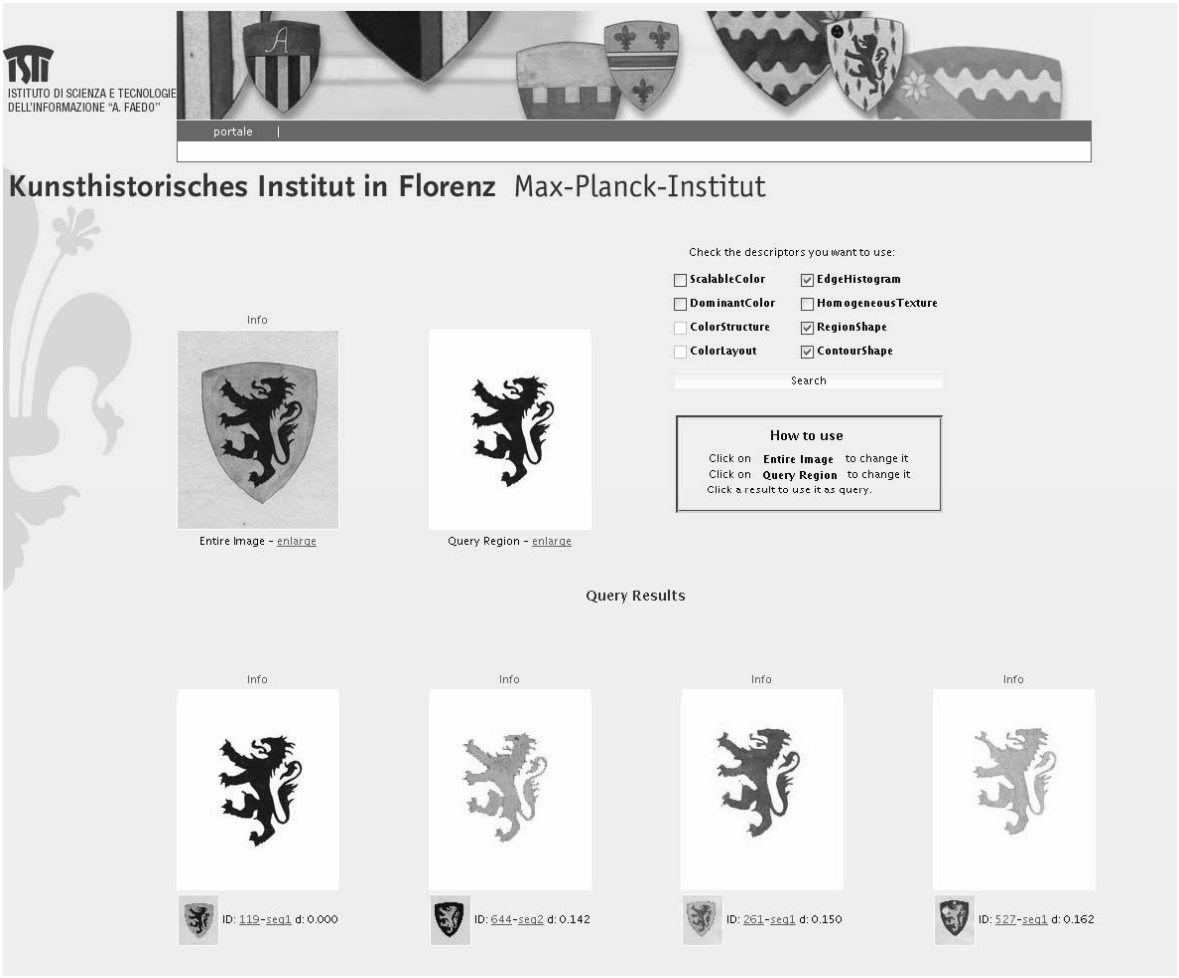

Figure 5: Search interface

The text-based approach is based on controlled vocabulary from heraldic vocabulary (colours, ordinaries, charges, and so on) and vocabulary from free-text fields such as blazon information and text information. The index search allows users to combine entries in selected fields, such as family name, blazon, ordinaries and charges. To support the users, the correct terminology for the fields is shown in a scroll field. Nevertheless, we think it difficult for non-experts to practice efficient retrieval based on text alone, due to the complex nature of the heraldic terminology which is only available in German.

An alternative approach, which does not need terminological knowledge or knowledge of the German language on the part of the users, is an image-based search. We have described the implementation of an image searching-by-content application in which a user can retrieve images of coats of arms by indicating some of their visual components. We will have to investigate how this approach is accepted by users, on the basis of future experience and compared to the text-based approach. In any case, the main advantage of this approach is that the procedure (segmentation, MPEG-7

feature extraction, content indexing) is completely automatic, once the images are digitised, and so can be applied to huge image databases with a minimum of effort.

5 Bibliography

[Amato, Gennaro, Rabitti, Savino 2004] G. Amato, C. Gennaro, F. Rabitti, and P. Savino. Milos: A multimedia content management system for digital library applications. In European Conference on Digital Libraries, ECDL 2004, Bath, UK, September 12-17 2004, 2004.

[Amato, Bolettieri, Debole, Falchi, Rabitti, Savino2006] G. Amato, P. Bolettieri, F. Debole, F. Falchi, F. Rabitti and P. Savino: Using MILOS to Build a Multimedia Digital Library Application: The PhotoBook Experience. ECDL 2006: 379-390

[Amato, Bolettieri, Debole, Falchi, Rabitti Savino 2006] G. Amato, P. Bolettieri, F. Debole, F. Falchi, F. Rabitti and P. Savino: Using MILOS to build an on-line photo album: the PhotoBook. SEBD 2006: 233-240

[Bieber, Schweibenz 2005] C. Bieber and W. Schweibenz. Images of Italian Art on the Web. The Database of the Photo Library of the German Art Institute in Florence. In: New Library World 106 (9/10) 2005. 430-440.

[Bieber, Schweibenz 2004] C. Bieber and W. Schweibenz. Digitalisierung und Langzeitarchivierung von Fotonegativen. In: Information zwischen Kultur und Marktwirtschaft. Proceedings des 9. Internationalen Symposiums für Informationswissenschaft (ISI 2004) Chur 6.-8. Oktober 2004. Hrsg. von Bernard Bekavac, Josef Herget, Marc Rittberger. (Schriften zur Informationswissenschaft 42). Konstanz: UKV. 125-140.

[Comaniciu, Meer 1997] D. Comaniciu and P. Meer. Robust Analysis of Feature Spaces: Color Image Segmentation. IEEE Conf. on Comp. Vis. and Pattern Recognition, pages 750--755, Puerto Rico 1997.

[Huang, Rui, Mehrotra 1997] T. Huang Y. Rui and S. Mehrotra. Content-based image retrieval with relevance feedback in mars. IEEE International Conference on Image Processing (ICIP), 1997.

[ISO/IEC 15938] ISO/IEC. Information technology – Multimedia content description interfaces. 15938.

[ISO/IEC 15938-6:2003] ISO/IEC. Information technology – Multimedia content description interfaces. Part 6: Reference Software. 15938-6:2003.

[ISO/IEC 15938-3:2003] ISO/IEC. Information technology -Multimedia content description interfaces. Part 3: Visual. 15938-3:2002.

[Ma, Manjunath 1999] W. Y. Ma and B. S. Manjunath. Netra: A toolbox for navigating large image databases. Multimedia Systems, 7(3):184{198, 1999.

[Mezaris, Kompatsiaris Strintzis 2004] V. Mezaris, I. Kompatsiaris and M. G. Strintzis, "Still Image Segmentation Tools for Object-based Multimedia Applications", International Journal of Pattern Recognition and Artificial Intelligence, vol. 18, no. 4, pp. 701-725, June 2004.

[Mezaris, Kompatsiaris Strintzis 2002] V. Mezaris, I. Kompatsiaris, and M. G. Strintzis, "A Framework for the Efficient Segmentation of Large-format Color Images", Proc. IEEE International Conference on Image Processing (ICIP 2002), vol. 1, pp. 761-764, Rochester, NY, USA, September 2002.

[MILOS] MILOS: Multimedia dIgital Library for On-line Search. http://milos.isti.cnr.it/.

[Salembier, Sikora 2002] P. Salembier and T. Sikora. Introduction to MPEG-7: Multimedia Content Description Interface. John Wiley & Sons, Inc., New York, NY, USA, 2002.

[Spreti 1928-1936], Enciclopedia storico-nobiliare italiana: famiglie nobili e titolate viventi, riconosciute dal R°. Governo d'Italia / promossa e diretta dal Marchese Vittorio Spreti, 1928-1936.

[Swain, Ballard 1991] M. Swain and D. Ballard. Color indexing. International Journal of Computer Vision (IJCV), 7(1):11{32, 1991.

[Wappenbilderordnung 1986] Wappenbilderordnung. Symbolorum Armorialium Ordo, hrsg. vom Herold, Verein für Heraldik, Genealogie und verwandte Wissenschaften zu Berlin. Bearb. von Jürgen Arndt und Werner Seeger. Neustadt an der Aisch: Bauer & Raspe, 1986.