A User-Centered Approach to Content-Based Retrieval of Medical Images

Ernesto Damiani*, Giovanni Righini*, Rajiv Khosla
Dept. of Computer Science – La Trobe University
Bundoora, Melbourne, Australia
Phone: +49-2408-94580, Fax: +49-2408-94582
email: khosla@cs.latrobe.edu.au
*Università di Milano, Polo di Crema
Via Bramante 65, 26013 Crema, Italy
Phone: +39-0373-898240, Fax: +39-0373-898253
email: edamiani@crema.unimi.it

ABSTRACT: In this paper we describe query support to a distributed collection of medical images as an example of a *user-centred* retrieval architecture. Most conventional multimedia databases for medical applications execute content-based indexing *off-line*, independently from their query execution mechanism. In our user-centred approach, image data are interpreted immediately before they are queried, generating suitable content descriptions on-line. Since the requirement of on-line operation prevents costly computations, only rough estimates of content related indexes are computed; this imprecision is dealt with by exploiting fuzzy query support, in the line of the application of fuzzy query languages to multimedia databases

KEYWORDS: multi-agent architectures, fuzzy multimedia databases, on-line indexing

1. Introduction

Intelligent multimedia information retrieval involves systems that enable users to create, index, present and interact with information within and across heterogeneous media such as text, speech, non-speech audio, graphics, imagery, animations, and video.

In the field of image databases, while older systems relied on the straightforward technique of providing textual *descriptions* to be stored together with the images, a number of more recent approaches focuses on using *color*, *hue*, *texture* or *shape* [Fal94] as the basis for image indexing and querying. Such systems try to exploit directly perceived pictorial features to model and reproduce the user evaluation of similarity between relevant features; the outcome of this similarity evaluation is then used for retrieving images from a multimedia archive.

However, while promising from the purely technological point of view, many of these innovative systems (see for instance [Bin94]) explicitly renounce incorporating in the query language naming or coding methods already familiar to user communities in specific application domains, as they turn out very difficult to be mapped to mathematically satisfying definitions of similarity.

Some other approaches to image databases, in order to attain effective context-based indexing, manage *feature data*, which can be seen as points in a suitable *feature space* [Chi94]. These features are classically described as a probability distribution function, rather than as a single point. However, such systems need a probabilistic *similarity measure* to support image classification, requiring complex trade-offs between efficiency and effectiveness of retrieval [Bar95].

Moreover, all systems mentioned above are *monolithic*, i.e. they integrate indexing, search and storage facilities at the same location, requiring the whole multimedia database to be stored in the same place and available for indexing at the same time.

In this paper we follow a different approach, describing query support to a distributed collection of medical images as an example of *user-centered* retrieval architecture for multimedia data.

Most conventional multimedia databases for medical applications execute content-based indexing off-line, independently from their query execution mechanism. In our user-centered approach, image data are interpreted immediately before they are queried, generating suitable content descriptions on-line. Since the requirement of on-line operation prevents costly computations, only rough estimates of content related indexes are computed; this

imprecision is dealt with by exploiting fuzzy query support, in the line of the application of fuzzy query languages to multimedia databases proposed in [Dub99].
Tele-health applications, like all other applications involving distributed components, require standardization both at the logical and at the physical interface levels to ensure modularity and cooperation among different systems. The application presented in this paper relies on the MASIF standard [Mas99] for CORBA-compliant agent-based systems, proposed by the Object Management Group.
The paper is organized into six parts: in the following Section, the characteristics of the problem are first outlined. Section 3 describes a reference model for distributed image retrieval, while the fourth section presents a user-centered solution to some of the classification and retrieval problems presented in the paper.
In Section 5 a sample application for distributed classification and retrieval of medical images is presented in some detail, and Section 6 briefly comments on the implementation of an agent-based system complying to the MASIF standard. Finally, Section 7 draws the conclusion.

2. Characteristics Of The Problem

According to a standard definition agreed in 1990 among the Ministries of Health of the European Union, *Telematics for Health Care* is "the integration, monitoring and management of medical data, including those used for personnel education, by means of communication systems providing ready access to medical databases and/or the advice of remote experts" [Tel98].
In the framework of such *telehealth* systems, searching and querying image collections is considered a routine activity in the medical field and constitutes an important support to differential diagnosis and clinical research. This application field lies at the borderline between telehealth and multimedia data analysis and processing, and several layered standards for image representation and management have been developed, originally aimed at radiological images transmission, and later extended to other fields such as dermatology and dentistry. Examples of such domain-specific layered standards for storage and transmission of medical images are DICOM/3 [CEN98] and SPI [SPI98].
In this paper we shall focus on the area of *user anatomy* and *pathology*, the latter addressing illnesses involving user cellular tissues. While a broad range of laboratory techniques are available in this field, pathologists are well aware of the increasing importance of consulting remote image libraries illustrating specific pathologies in the framework of distributed telehealth systems.

2.1 Image libraries and diagnostic systems

Many image-based diagnostic applications are currently being developed, such as tumor detection systems for full-digital images. In this field, even simple techniques for content-based indexing of images are expected to have a considerable impact on the accuracy and efficiency of the diagnostic process. Image processing for tumor detection poses two fundamental problems:

- *Shape enhancement* It is not rare that tumors have very weak contrast against their background and how to detect “suspicious regions” in digital images is one of the key points to establish reliable system.

- *Feature extraction* Expert knowledge must be exploited in order to extract features that best characterize pathologies such as malignant tumors.

The former problem has been tackled by means of adaptive filters. Such filters [Kab96] can be very effective in enhancing approximately rounded opacities no matter what their contrasts are. Clues for discriminating between malignant tumors and other pathologies are believed to be mostly in their boundary areas. Therefore, the detection of boundary is an important preprocessing for feature extraction. The boundary of malignant tumor is usually fuzzy, and fuzzy-based filters are often adopted to estimate boundary of suspicious region. In [Kab96] the optimal region of support of a filter was shown to coincide with the boundary of an approximately rounded convex region if the pixel of interest is in its inside. By applying it to real mammograms, it was possible to extract probability of the existence of the tumor boundary. In other words, the fuzziness of the boundary was reflected on the boundary probability obtained by the filter.
Several feature parameters can be adopted to identify malignant tumors from suspicious regions. It is widely recognized that the most effective of them reflect boundary characteristics; the shape parameter "ellipticity" was adopted as a feature in [Kab96] and will be used as a reference in the remainder of the paper.
Several experiments are documented in the literature to test the effectiveness of feature parameters. Classification experiments were conducted for instance in [Kob98], using 1313 suspicious regions, which were detected by an adaptive filter from 354 images. The number of suspicious regions corresponding to malignant tumours is 71. Adjusting

the system parameters so that 96% of malignant tumours could be identified correctly, the average number of false positives was 0.9 per image. These encouraging results justify the increasing interest in designing applications and systems for content-based indexing of medical images, as will be clarified in the remainder of the paper.

While a global framework for image sharing and processing is still in its infancy, a number of well-known medical images repertoires illustrating various pathologies are currently available in GIF or JPEG format over the Internet, though few of them are stored in fully-fledged image databases. Indeed, most of these collections are only accessible through standard Web sites, each of them maintained by a different medical institution.
Moreover, while the medical community has since long agreed on a basic textual coding for image contents (namely *SNOMED*, the *Systematized Nomenclature of User and Veterinary Medicine*), no standard technique for indexing or retrieval of medical images based on color or shape is currently available or planned.
Medical users are therefore accustomed to using SNOMED patterns as search and classification codes for images depicting pathologies.
Indeed, retrieval of images tagged by a SNOMED code is easily and efficiently performed through string *pattern* matching; Figure 2 presents some sample queries that can be posed using SNOMED pattern-based language [Mau96].

SNOMED
(Systematized Nomenclature of User and Veterinary Medicine)

```
T: Topography

      T=2*   Breathing Apparatus
      T=28*  Lungs
      T=281* Right Lung
```

Figure1 Sample SNOMED patterns

In the following Sections, a technique for transparently superimposing hybrid agent-based search to SNOMED compliant Web-based collections of medical images will be presented as a user-centered alternative to storing such images in a conventional, monolithic multimedia database.
Moreover, a sample content-based on-line indexing technique will be described.

3. A Reference Model for Distributed Image Retrieval

Before describing in detail the functionalities of our agent-based architecture for indexing and retrieval of medical images, we briefly outline its reference model at the highest level of abstraction, in order to clarify the use of the terminology.

Abstract Model
Roles:
Customer
Broker
Supplier
Actions
Classify
Search
Deliver

Tab. 1 Roles and actions of the image classification and retrieval reference model

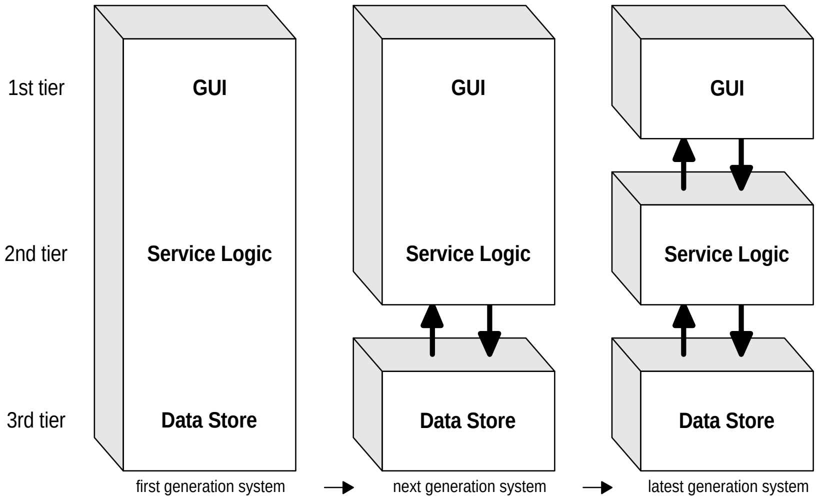

Figure2 An architectural view of roles

The core concepts of our simple reference model are the roles of *customer* (the medical researcher querying the system), *broker* (our Broker Agent for multimedia data retrieval) and *suppliers* (the sites storing collections of medical images*).* To these roles, the actions of *search*, *classify* and *deliver* are associated.
In Figure 2, an architectural view of this distinction is given, outlining how all service logic can be located in the Broker Agent, while customer and supplier respectively provide a graphical user interface and fast conventional data storage, possibly made accessible through a Web site.
In fact, in our system the Broker Agent acts as a supplier of images to the customers, and as a classification and distribution mechanism for suppliers wishing to make their images accessible over the Internet.
A basic assumption of our application design is that suppliers *do not index* images, other than providing their standard SNOMED codes; it is left to brokers to compute and store content representations of online images, in order to be able to locate the image required by the customers.
The function of our Broker Agent is therefore providing a path whereby the customer may find and obtain from a supplier a set of images offering the required characteristics to the highest possible degree.

4. A User-Centered View of Medical Images Collections

The basic idea underlying our approach is putting the user's (in this case, the medical researcher) perception of the information space (i.e., the collections of medical images available over the Net) at the center of the knowledge organization process.
The basic assumption we rely on turns out to be the same of SNOMED coding: the medical researcher organizes a mental hierarchical model of the information domain via a limited number of general features, and an agent-based broker should be able to fully comprehend and utilize such a user-centered model.
Our methodology [Cas99] starts with a *decomposition* phase, where a *Decomposition Agent* is used employing a small number of coarse-grain input features in order to identify a hierarchy of abstract classes.
In the present setting, SNOMED hierarchical encoding, in addition to being familiar to the user, provides an easy and effective way to decompose the medical images domain.
Moreover, it has the additional advantage of being familiar to the user community. Therefore, we use SNOMED domain decomposition as the guideline for designing our system, while intervals and ordinal indexes are computed *on-line* on the basis of image content.

4.1 From Domain Model to Decision Support

The SNOMED-based taxonomy identified by the decomposition phase presents a model of the whole medical images information space that is both simple and familiar to the user. However, this user-centered model is not related to the solution of any particular search problem. Our following step, the *Control Phase*, uses a *Control Agent* to execute a further specialization on the basis of finer-grain features whose values can be drawn on an *ordinal* or an *interval-based* scale.

In the first case, values will belong to any ordinal domain such as the integers, while in the second case the feature will be associated to a fuzzy linguistic variable. In our present application, *fuzzy linguistic variables* with trapezoidal elements such as the one depicted in Figure 4 will be used to represent features of image content.

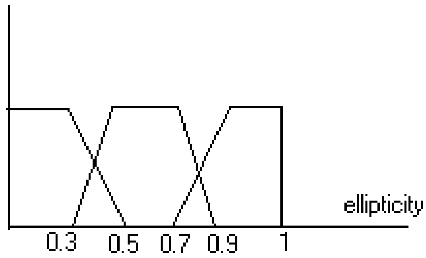

Figure4 Fuzzy elements for ellipticity linguistic variable

As mentioned above, a basic assumption of our design is that such content-related features, though an integral part of the domain model, need not be stored together with the images, but can be computed on-line by the Broker Agent. Therefore, their ordinal values need not be available in the decomposition phase. However, as we shall see in the next Section, the classes obtained in this phase will be directly involved in the decision support process of our image classification and retrieval system. It is interesting to observe that the SNOMED-based domain model developed in the previous phase allows us to easily deal at this level with the user's implicit knowledge about the (sub) domain.

For instance, the medical user's linguistic knowledge about a pathology of the uterus may well include the fact that it involves an "abnormal" ellipticity of black *blobs* in the image corresponding to pathologically suspect cells aggregates. In this case, the meaning of "abnormal" depends on the image being compared with other images of the same kind and not, say, with the image of a lung.

In our model, being a class feature, the fuzzy linguistic variable `ellipticity` may well have different definition intervals for different classes of images, thus dealing with the implicit different meanings of `abnormal`.

4.3 The Decision Phase

The Broker agent, in association with user input, to compute the Decision classes for a specific user query, will then use classes developed in the Control Phase. In our approach, this third step involves a hybrid Decision Agent, using the search and identification technique (Rule-based decision system, Fuzzy matching, Neural Networks,..) which is more suited to the specific multimedia information available from the supply sites.

In the present application, a novel selection technique based on fuzzy queries [Bos97] has been used. As we shall see, user input must be pre-processed and fuzzified in order to be used by the Decision Agent.

5. A Sample Application

The medical user-centered design of our architecture requires the broker to compute and retain information about the available images in the format of a hierarchy. Classes at the higher levels exhibit general and structured features based on a nominal scale, while lower level classes present ordinal and interval-based features.

To decompose the chosen sub domain on the basis of the nominal scale, the Decomposition agent exploits the available body of domain knowledge, in this case the SNOMED encoding.

At the Control level, however, intervals and ordinal values must be computed on the basis of image content.

In this Section, we outline how such a computation can be carried out exploiting a broker-managed repository storing *descriptors* of the images available on the network.

5.1 Descriptor-based Image Classification

The repository holding image descriptors is a structured collection of simplified descriptions of image properties, in the line of [Dam97]. We consider that such a hierarchical repository to be associated with each application sub domain, i.e. to each specific prefix of SNOMED coding, such as, for instance `T=3*` for the cardiac system.

In such a repository, the O-O control level classes are stored as a set of *fuzzy relations*, which are defined by applying an imprecise classification criterion on a crisp relation through a fuzzy predicate. For such a fuzzy relation, built on a set of domains D_i, every t-uple is supplied with a membership degree μ_R, from 0 to 1, interpreting how this t-uple satisfies a fuzzy predicate **P** applied to the relation **R**.

In the simplest case, the repository is a single fuzzy relation whose attributes are: *OID*, *feature*, *fuzzy element*, *weight*. Each feature corresponds to a linguistic variable, such as `max_blob_ellipticity`, having several fuzzy elements.
To each fuzzy element of each feature a *weight* is associated, describing to which extent the object offers the corresponding property. From a syntactic point of view, features are expressed by *nouns* whereas *adjectives* describe fuzzy elements.
Figure 4 shows a simplified fuzzy relation describing the properties of two sample medical images. The actual data model currently employed in our system will be discussed in detail in the next Section.

Object Id	**Feature**	**Fuzzy Element**	**Weight**
1	Max_blob_ellipticity	Slightly abnormal	.3
1	Max_blob_ellipticity	Normal	.8
1	Max_blob_magnification	High	0.7
2	Max_blob_magnification	High	0.8
2	Max_blob_ellipticity	Normal	1

Figure 5: An example of fuzzy descriptor relation

The Broker's Control Agent computed the above table on-line after executing a standard noise reduction based on fuzzy techniques ([Ste98], [Rus95]), by applying to two images the definition of the `max_blob_ellipticity` and `max_blob_magnification` linguistic variables (Figure 3).
Our classification procedure involves, for each image, five main steps:

- First, the image is converted to a standard RGB representation
- Secondly, RGB encoding is converted to the well known *hue-based* representation. Hue is a linear scale expressing the darkness of each pixel as a function of its RGB values.
- Thirdly, a simple hue threshold is applied to identify the main "dark" blobs in the image.
- Then, a deterministic convex-hull algorithm is used to determine the crisp ellipticity value associated to the biggest black blob in the image [Cas98],[Cas99]. This kind of *ellipse fitting* algorithms fit an ellipse to a set of points, which are all roughly associated with the ellipse while allowing for some tolerance to scatter. Current implementation of the algorithm allows us to set a maximum amount of time for index calculation per image (namely, 5 seconds).
- Finally, applying the linguistic variable definition corresponding to the SNOMED code associated to the image fuzzifies the crisp value obtained in the previous step.

It should be noted that step 1 and 2 could in principle be avoided, as GIF or JPEG images collected from Web sites could be easily converted to a gray-scale representation before computing the threshold. After conversion to gray-scale, however, images to be indexed are large, rectangular arrays of pixels with different gray-levels. Within this array of pixels there may be several blobs. In order to apply an ellipse-fitting algorithm to such blobs, it would be necessary to segment out points belonging to each of them. It is well known that, in principle, it is not possible to identify points belonging to a blob on the basis of gray levels, if the blobs are superimposed on a background of much smaller rectangles of varying gray level.
Several *Shell-clustering* techniques are currently available [Dav90], providing shape finding approaches that use fuzzy, possibilistic, and robust methods to find unknown numbers of parameterized curves from sparse and noisy edge data.
However, due to time constraints imposed by on-line indexing, in this application a simple threshold must be applied to identify black areas in the image, and ellipse fitting must be applied to these areas only. Hue representation proved to be more effective than others for the available images; the availability of fast conversion software allowed for choosing a representation more sensible to the chromatic property of the original image.
It should also be noted that the linguistic variables used for classification are part of the system's knowledge base for a certain value of the SNOMED code.
The linguistic variables' definitions for a given sub domain are a part of the domain knowledge stored by the corresponding broker. It is also important to observe that this computation takes place as the suppliers' sites are periodically *polled* with simple HTTP connections, allowing the Broker Agent to take updates into account.
Finally, it should be noted that, while the simple convex-hull deterministic algorithm (see the next section) proved to be reasonably fast and accurate, other techniques could be evaluated in order to be put at the Broker Agent's disposal as alternative tools to be selected according to time or accuracy constraints. A back-propagation neural network [Yim99] with a single hidden layer can usually map inputs to outputs for arbitrarily complex non-linear systems, and could be

used in the framework of our system. Application of Kohonen's self-organizing maps [Koh95] is another interesting possibility.

5.2 Ellipticity evaluation

In order to satisfy the time constraints imposed by on-line indexing, the computation of blob ellipticity is performed in a "quick and dirty" manner, leaving it to the query execution engine to deal with data imprecision introduced by the indexing procedure. Here is the pseudo-code of the naive algorithm employed for computing ellipticity:

```
Point:=ComputeBlobCenter(Blob)
MAxEll:=0;
InitialSlope:=0;
For step:=0 to maxstep do
Ratio:= (GetSegmentLength (Point, InitialSlope+step)/ GetSegmentLength (Point,
InitialSlope+step+π/2));
if MaxEll<Ratio then MaxEll:=Ratio;
od;
Return MaxEll;
```

Figure 6 Pseudo code of the indexing algorithm

The employed naive algorithm simply computes the center of the blob and then samples pairs of orthogonal straight lines passing through it with angular coefficients from 0° to 360° degrees (measured in the image local coordinate system).
The sampling step is a parameter of the algorithm as it can be tuned according to time constraints and/or image characteristics. It is usually set at $\pi/10$. On each line, the length of the segment belonging to the blob (i.e., whose hue values are all above threshold) is computed. Then the highest ratio between a pair of orthogonal segments is returned as a rough estimate of the blob's ellipticity.
Finally, we remark once again that this procedure does not take into account blob identification, which is obtained by applying the threshold; this of course could mean "losing" some blobs whose hue is not different enough from the image background, or computing some "false positives" for the opposite reasons. For this reason, user feedback is collected: the user is given the possibility of manually modifying the index value for any of the images returned in the query result.

5.3 Querying the Broker Agent

User interaction is made through a User Agent, whose graphical interface is shown in Figure 7.
From the user point of view, four main elements are involved in the interaction:

- A medical **user connected to Internet** from a certain machine, who requires a service from the Broker Agent system. We assume that the client has previously contacted a “Master Broker”, which contains lists of various URL of known Broker in the SNOMED domain under consideration.
- A **Broker site** where all classes defining the GUI interface are stored, together with the Broker Agent itself and its relational repository, including the characteristic values of the various fuzzy predicates. The Broker receives fuzzy requests from users who obtained the GUI connecting to the Broker Web site. On the basis of this information, performs the division (according to the fuzzy predicates) on the repository and returns the best matching images to the user.
- An **Adapter site** that supports a fuzzy adapter system whose role is to dialogue with Web servers storing medical images and maintain a coherent view of fuzzy predicates according to some images' content-related properties.
- Several **Web servers** that actually provide the images whose content is described in the Trader base. No fuzzy issue characterizes these sites: the server programmer or installer only specifies the SNOMED code for each image as a part of the HTML or XML markup of the page holding it.

.
The first interaction step requires the medical user to select a specific Broker agent, i.e. a part of the domain model built in the Decomposition phase, again via SNOMED codes. SNOMED provides a complete decomposition of the information space: a hierarchy of Brokers could well rely on the domain decomposition provided by SNOMED.
In a second step, the user must detail the features of the desired image.
User input is collected through a graphical interface in order to identify a fuzzy predicate for each interval-based or ordinal feature specified by the user.

When provided by the user, an absolute value in the definition universe of the selected predicate is also selected (for instance `max_blob_magnification =2`).
This results in a query language whose simple grammar is given in Figure 7 in *Backus-Naur Form*.

Productions:

```
statement:= noun [':' adjective] ['=' value]
noun:= TERM |REGX
adjective:=TERM
value:=REAL
```

Terminal symbols:

`TERM:` a string coded in ISO-LATIN-1
`REGX:` a string coded in ISO-LATIN-1, including wildcards '?' and '*'
`REAL:` a real number

Figure 7 The query language grammar

While the simple grammar given above is rather self-explanatory, a comment should be made on the fact that it does not explicitly involve any fuzzy number or linguistic modifier. Indeed, in a user-centered approach to image querying it would be rather unrealistic to require the medical researcher to be conversant in, or even knowledgeable about, fuzzy query techniques.
Here, all the user needs to know is that since classification based on content related parameters is performed on-line, it is forcedly raw, i.e. values are to be taken as approximate.
Therefore, while the query interface requires the user to specify crisp values for content-related queries, query result will automatically include also images whose indexes are reasonably "near" to the value specified by the user.
Of course, many content-based indexes other than `ellipticity` and `magnification` of black blobs can be used for querying; to this aim, we assume the availability of a Thesaurus, i.e. a medical *controlled vocabulary* allowing features about images in a given sub domain to be uniformed through a naming discipline, in order to deal with a standard context-dependent vocabulary [Dam97].
SNOMED itself provides a fully featured glossary [Mau96] that allows both Brokers and clients to use a subdomain-specific language to express features, without any explicit reference to the fuzzy model. Fuzzy elements and their membership values, i.e. the internal knowledge representation, are only computed and dealt with inside the broker.

5.3 Computation Of The Decision Classes

With reference to the previous example, a medical user could request to the Broker Agent an image having the following features: a `max_blob_magnification` of 2 and a `max_blob_ellipticity` of 0.8 (the reader should note that this latter value is not a fuzzy membership value, but the crisp geometrical parameter expressing ellipticity).
User input filtering computes a list of properties, each one associated with a certain fuzzy predicate and weighted by a value between 0 and 1. Once again, we remark that these values are obtained inside the Broker Agent, by transformation of crisp values specified by the user according to the linguistic variable definition determined at the Control level.
The processed input defines a *fuzzy request* to the broker, which is nothing but another fuzzy relation like the one below.

max_blob_magnification	average	1
max_blob_ellipticity	abnormal	0.9

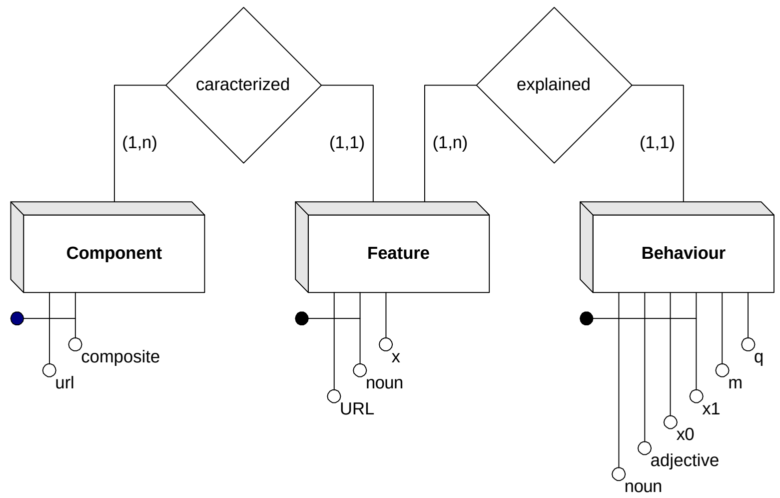

Figure 8 Relational schema of the broker's repository

Executing a *fuzzy division* ([Bos97]) between the user query and the Feature table in the broker's repository allows the Decision Agent to compute a ranked list of candidate images, whose URLs is returned to the user.

5. 4 The Division Operation

In order to describe in detail the internal operation of the Broker, we now recall some basic definitions about relational division. Let us consider two relations R(X,A) and S(Y,A) where A, X, and Y point out sets of attributes. The division of R by S, denoted R[A/A]S is a relation on X, which can be defined as follows:

$$x \in R[A/A]S \text{ if } \forall a \in S[A], (x,a) \in R.$$

Following [Bos97] we now examine the extension of the division to fuzzy relations.

The operation of division of R by S can be considered as a set inclusion:

$$x \in R[A/A]S \Leftrightarrow S[A] \subseteq \Gamma^{-1}(x) \text{ , with } \Gamma^{-1}(x) = \{a, (x,a) \in R\}$$

By clicking on an URL, the user directly accesses the Web site supplying the corresponding image.

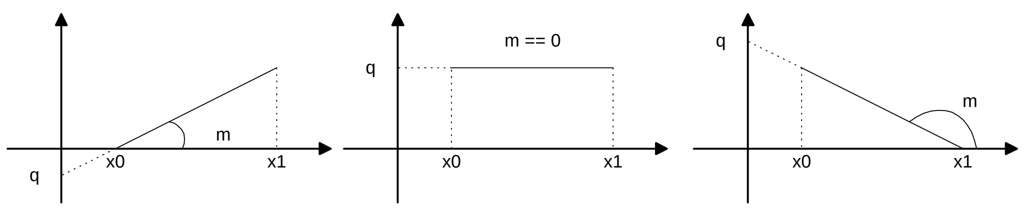

Figure 9: Membership function storage format used by the broker agent

In order to compute such a list of images, our Decision Agent exploits the Broker Agent's internal repository, whose relational schema is depicted in Figure 8.
The entity *Component* has two attributes: *URL* is a reference to the image on the Web and *Composite* is the HTML document holding the image. This entity gives access to the Web pages holding the images composing a query result. Note that the URL alone would not be sufficient, as an image may well appear in multiple Web pages.
As mentioned above, image features are stored in the entity *Feature*, much in the same way as shown in Figure 4. However, instead of explicitly storing fuzzy membership values, in our system we use entity *Behavior* to store definitions of trapezoidal fuzzy elements as shown in Figure 9.

In our current release, the linguistic variables have three fuzzy elements each, but (as is usually the case) the superposition in their definition involves the fuzzy elements two by two. Thus, in this case the maximum amount of fuzzy offer properties having non-zero membership for each server is 6.
In order to compute membership values for the content-related indexes the Broker actually uses the parameters describing the shapes of the fuzzy elements' membership functions
Thus, if we call *inf* the inferior boundary, *sup* the superior boundary and *max* the value for which the membership function yields 1, we obtain for the first fuzzy element of the `ellipticity` feature:

μ_{low}: $[0 .. 1] \to [0,1]$
$x \to 0$ if $x \geq$ sup or $x \leq$ inf
$x \to 1$ if $x =$ max
$x \to (x - \text{inf}) / (\text{max} - \text{inf})$ if $x \in [\text{inf},\text{max}]$
$x \to (\text{sup} - x) / (\text{sup} - \text{max})$ if $x \in [\text{max},\text{sup}]$

However, there is no need to explicitly evaluate such functions, as the Broker Agent can compute fuzzy membership values at run time as simple SQL queries.
With reference to the membership functions storage technique described in Figure 10, a sample query to the Broker Agent repository is reported below:

```
SELECT (b.m*f.x+b.q) AS weight
From Feature As f, Behaviour as b
WHERE b.x0<=f.x and f.x <=b.x1
```

The above SQL query computes the membership values (`weight`) exploiting the definition of linguistic variables held by the Broker Agent in order to transform crisp values into the fuzzy format required to perform the division.
Besides selecting the target site, the medical user can choose the aggregation operator the Decision Agent will use to compute the division.

Available choices include fuzzy AND (min), fuzzy OR (max) and any mean-based aggregation operator (AVG). While leaving to the user the choice of the operator to be employed may at first sight seem improper, the resulting system operation is indeed easy to understand.
In fact, the choice of the AND operator results in retrieving images possessing at the highest possible degree *all* the features required by the user. Images lacking even one of the requested features will be left out. The OR operator selects images possessing (at the highest possible degree) at least one of the features required by the user, regardless if they possess the other features or not. Finally, the AVG operator corresponds to an intermediate behavior between AND and OR.
While the present choice of aggregation operators could in principle be extended to widen the spectrum of system behavior [Bos97] in order to fully support any user-selected retrieval semantics, it should be noted that currently supported operators are easily executable using standard SQL queries.

5.4 Collecting User Feedback

As the on-line indexing system described above is admittedly prone to errors, user feedback plays an important role in the framework of our system, corresponding to our methodology's Validation phase.
Following the approach described in [Bel99], the user is given the possibility of manually modifying the index value for any of the images returned in the query result, in order to eliminate invalid data.
Users are required to input an approximate value, to be cached by the Broker and used in future query sessions involving the same image.
As a help to providing this estimate, users are shown the wrongly indexed images in a window where they can point out four segment endpoints to be used by a client applet to compute the ellipticity value to be transmitted to the Broker Agent.

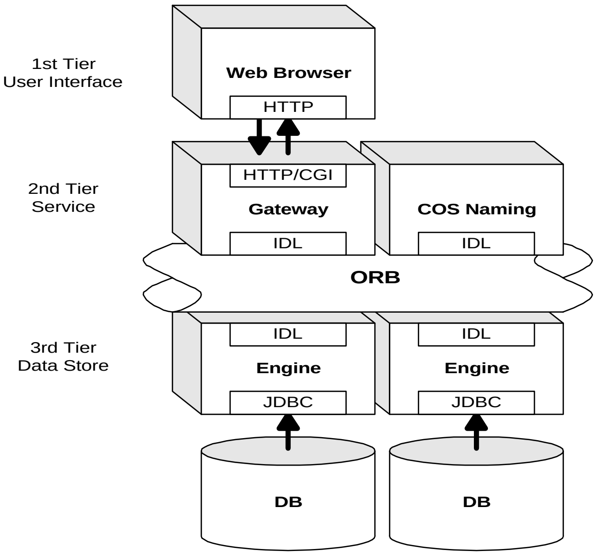

Figure 10: Current prototype architecture

6. System Architecture

Having presented the design of a general and flexible Broker Agent system based on our approach, we are now ready to proceed to describe its implementation architecture.

The system includes procedures for classification and retrieval as a part of the Broker Agent, as well as a User Agent based on a graphical user interface. The architecture of the current implementation of our image retrieval system is depicted in Figure 10. Both design and implementation are based on CORBA-based *Mobile Agent System Interoperability Facility (MASIF)* proposed by General Magic, IBM, GMD and the Open Group [Mas99]. In order to standardize design and implementation of O-O agent-based distributed execution environments, the Object Management Group (OMG) decided to promote MASIF in the framework of the *CORBA* (*Common Object Request Broker Architecture*) software architecture [Yan96].

CORBA reference architecture, called the *Object Management Architecture (OMA)* upon which applications can be constructed. OMA attempts to define at a high level of abstraction all facilities needed for distributed object-oriented computing. It consists of four components: an *Object Request Broker (ORB)*, *Object Services (OS)*, *Common Facilities (CF)*, and *Application Objects (AO)*.

The core of the OMA is the ORB component, a transparent communication bus for objects that let them transparently make requests and receive responses from other objects. In other words, the ORB allows client and server objects, possibly written in different languages, to co-operate. It intercepts calls and is responsible for finding an object that can execute them, pass it the parameters, invoke its methods and return the results. Invocations can be done either statically at compile time or dynamically at run time with a *late binding* of servers.

Objects Services specifications define a set of objects that perform fundamental functions such as naming services, life cycle services, transaction services or trader services. Generally speaking, they augment and complement the functionality of the ORB, whereas *CORBA Common Facilities* provide services of direct use to application objects. Finally, *Application Objects* implement distributed services to be invoked by clients through the ORB bus.

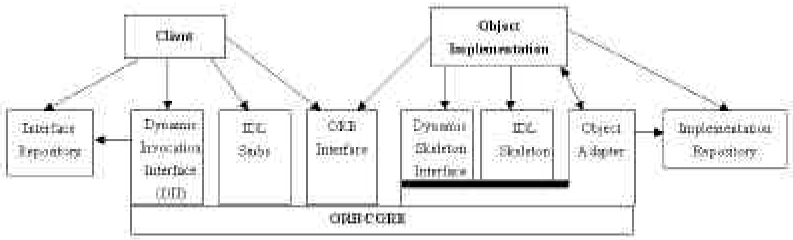

Figure 11: Structure of ORB communication

The prototype described in this paper is entirely implemented in Java, using CORBA support provided by JDK 1.2. It currently indexes images from several European medical sites, as well as some local demo databases; all Broker Agents provide a CORBA-compliant IDL interface to be used by User Agents in order to query the system.

On the client side, the user must first connect to the Broker site. His browser then downloads an HTML page including two Java applets (*TraderInitApplet.java* and *BestResult.java*), that is, Java classes implementing the predefined interface *java.applet.Applet*, together with the other MASIF classes directly called by these two applets. Such classes are used as interfaces for exchanging messages with the Broker Agent.

Since the purpose of the Broker is to help the client to select a picture through a GUI interface, the applet *TraderInitApplet* first displays the window enabling the user to select the SNOMED code he or she is interested in.

Then the GUI displays a frame (corresponding to the file *PropertiesFrame.java*) to let the user make his selection over available features. All this information (i.e., the SNOMED pattern, the semantics and the corresponding operator, the selected features and the crisp values) compose a fuzzy request, that the *TraderInitApplet* passes on to the Broker to compute the division, after having contacted it using the naming service of the MASIF/CORBA architecture. For each picture regarding the subject, the Trader computes its degree of satisfaction relatively to the fuzzy query (using the *FuzzyCalc* class). Finally, it returns the *BestResult* applet information on the best matching service: the address of the server where the picture can be found, its name, and its final degree (see Figure 12). The client browser then connects to the target site to see the retrieved picture, by means of a third applet, *TraderImageViewer*.

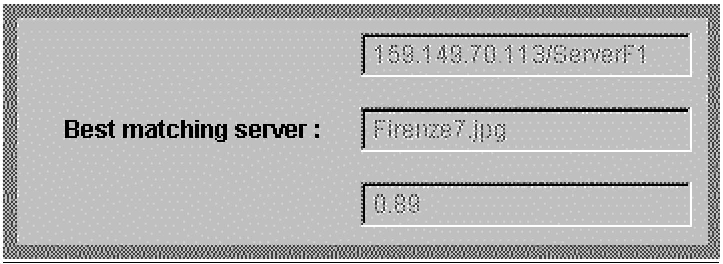

Figure 12 The Broker Agent answers the user query showing the best target site

7. Conclusions

In this paper we have presented a MASIF-compliant system for medical images on-line indexing and retrieval based on a user-centered approach. The following table summarizes the methodology of our *Medical Images Retrieval Model* as a layered architecture:

Medical Image Retrieval Reference Model	**Phase of the Methodology**
Structure of a *image description* to be supplied by customer	Input filtering and pre-processing
Domain Specific Knowledge SNOMED-based classification of images in the domain of interest	**Decomposition** *Domain model* as a hierarchy of abstract classes. Domain specific metadata.
System Functional Units Fuzzy query to the broker's database.	**Control** Alternative Fuzzy/NN component identification techniques, based on multimedia data and domain specific metadata
Candidate Selection Presentation of candidate images	**Decision** Candidate classes to be proposed to medical user
Transaction Network transfer of the chosen images/Collection of user feedback on the images' indexing	**Validation** Optional validation of the decision phase

As we have seen, our approach allows for providing a seamless connection between a user-centered domain decomposition, based on nominal scale features, to ordinal and interval-based features suitable to be processed by intelligent agents. Besides developing a full-featured system for on-line indexing of pathology images [Cas99], we explored a number of other applications of the approach, including on line analysis and retrieval of satellite images for crop control [Ahm99]. Many measures of comparison of descriptions of objects have been proposed and studied in given frameworks or domains of applications. B. Bouchon-Meunier, M. Rifqi and S. Bothorel [Bou96] proposed a general classification consisting in four main kinds of measures of comparison, depending on the purpose of their utilization: they define measures of *satisfiability, resemblance* and *comparison* that on one hand, can be considered as measures of similarity, and on the other hand measures of dissimilarity.

In a future release, a general measure of dissimilarity will be used to evaluate to which extent an object is different from a reference, helping the user to choose objects that are closer to the reference than the others. Secondly, as far as the design of the architecture is concerned, we are currently working on the servers-to-Broker dialogue in order to support intelligent load balancing techniques.

However, several other research directions could be also pursued.

First of all, the mechanism of retrieval of servers can be enlarged and complemented on the basis of features the images should *not* have, in order to eventually decide between those that would have a same level of fitness after the division. Indeed, the phase of division searches in the repository for images that have as many as possible of the desired requirements. On the contrary, the optional phase of differentiation, through a measure of comparison, should compute a similarity or a dissimilarity measure to eliminate images that have features the user did not ask for.

Acknowledgements The authors wish to thank Maurizio Lestani, M.D., Maurizio Pea, M.D. and their colleagues at the Institute of Anatomy and Pathological Histology, Medical University of Verona, Italy, for their precious assistance on medical issues. Thanks are also due to Letizia Tanca and Patrick Bosc for their encouragement and valuable suggestions. Ruben Castelletti implemented the first prototype of the system.

6. References

[Ahm99] M. Ahmed, E. Damiani, A. Tettamanzi "A Fuzzy System for Crop Control", in B. Reusch, ed., Proc. of the 6th Intl. Conf. on Computational Intelligence, Dortmund, 1999, LNCS 1625

[Bar95] J. Barros, J. French, W.Martin, P. Kelly "A System for Indexing Multi-spectral Satellite Images for Content-Based Retrieval", Proc. of IS&T/SPIE: Storage and Retrieval for Image and Video Databases III, San Josè, CA, 1995

[Bel99] C. Bellettini, E. Damiani, M.G. Fugini "Collecting User Feedback in Software Reuse Systems", Proc. of the SSR Workshop on Software Reuse, co-located with IEEE Intl. Conf. on Software Engineering, Los Angeles 1999

[Bin94] E. Binaghi, I. Gagliardi, R. Schettini "Image Retrieval Using Fuzzy Evaluation of Color Similarity", Intl. Journal of Pattern Recognition and Artificial Intelligence vol.8, n.4, 1994

[Bos97] P. Bosc, D. Dubois, O. Pivert and H. Prade, "Flexible Queries In Relational Databases - The Example of The Division Operator", Theoretical Computer Science, vol.171, 1997

[Bou96] B. Bouchon-Meunier, M. Rifqi, S. Bothorel, “Towards General Measures of Comparison of Objects”, Fuzzy Sets and Systems, vol.84, 1996.

[Cas98] R. Castelletti, "Un servizio per la gestione distribuita di immagini di interesse patologico", Tesi di Laurea in Scienze dell’Informazione, Dip. Scientifico e Tecnologico, Università di Verona, 1998 (in Italian)

[Cas99] R. Castelletti, E. Damiani, G. Righini, R. Khosla "A User Centered Approach to Image Retrieval", in B. Reusch, ed. Proc. of the 6th Intl. Conf. on Computational Intelligence, Dortmund, 1999, LNCS 1625

[CEN98] CEN/TC251 Project Report, available from http://miginfo.rug.ac.be:8001/index.htm

[Chi94] T. Chiueh, "Content-based Image Indexing", Proc. of the 20th VLDB Conference, Santiago (Chile), 1994

[Dam97] E. Damiani, M.G. Fugini "Fuzzy Identification of Distributed Components", in B. Reusch, ed. Proc. of the 5th Intl. Conf. on Computational Intelligence, Dortmund, 1997, LNCS 1226

[Dav90] R. N. Dave, "Fuzzy Shell-Clustering and Applications to Circle Detection in Digital Images," International Journal of General Systems vol.16 no.4 (1990)

[Dub99] D. Dubois, H. Prade, F. Sedes, "Fuzzy Logic Techniques in Multimedia Databases Querying: A Preliminary Investigation of the Potentials", Proc. of IFIP Working Conference on Database Semantics (DS8) Roturoa, New Zealand, 1999

[Fal94] C. Faloutsos, W. Equitz, M. Flickner, W. Niblack, D. Petrovic, R. Barber "Efficient and Effective Querying by Image Content", Journal of Intelligent information Systems vol. 3 n. 3/4, 1994

[Kho97] R. Khosla, T. Dillon, "Engineering Intelligent Hybrid Multi-Agent Systems", Kluwer Academic Publishers, 1997

[Kho98] R. Khosla, “User-Centered Virtual Machine of Problem Solving Agents, Intelligent Agents, Software Agents and Objects”, Proc. of the 3rd IEEE Symposium on High Assurance in Intelligent Systems, Washington, Nov. 13-14 1998
[Koh95] T. Kohonen "Self-organizing maps", Series in Information Sciences, Vol. 30, Springer 1995, 2nd ed. 1997
[Kob96] H. Kobatake, M. Murakami “Adaptive Filter to Detect Rounded Convex Regions: Iris Filter”, Proc. Int. Conf. on Pattern Rec., Vol.II, pp.340-344, 1996.
[Kob98] H Kobatake, H Takeo, S Nawano “Tumor Detection System for Full-digital Mammography” Proc. of the 4th International Workshop On Digital Mammography Nijmegen, the Netherlands, 1998
[Mau96] F. Mauri, A. Beltrami, V. Della Mea *Telepathology Using Internet Multimedia Electronic Mail*, Journal of Telemedicine and Telecare, vol.6 n.3 1996
[Mas99] MASIF: A CORBA-based Mobile Agent System Interoperability Facility available from http://www.omg.org/masif
[May97] M. Maybury (ed.) “Intelligent Multimedia Information Retrieval”, MIT Press, 1997
[Rus95] F. Russo, G. Ramponi,” A Fuzzy Operator for Enhancement of Blurred and Noisy Images”, IEEE Trans. on Image Processing, vol. 4 (8), 1995
[SPI98] “The SPI standard: a White Paper”, available from http://www.siemens.de/med
[Ste98], I. Stephanakis, G. Stamou, S. Kollias “A Fuzzy Associative Operator for Enhancement of Noisy Images By Detection of Local Multiresolution Features”, Proc. of JCIS '98, Rayleigh, NC, USA, 1998
[Tel98] Telemedicine Research Center, "An Introduction to Telemedicine", available from http://tie.telemed.org/tiemap.html
[Yan96] Z. Yang, K. Duddy “CORBA: A Platform for Distributed Object Computing”, ACM Operating Systems Review, vol. 30, 1996.
[Yim99] T. Yiming, F. Ping, Z. Yong, *Y*. Siqian "Evaluating Nugget Sizes of Spot Welds by Using Artificial Neural Networks", in B. Reusch, ed. Proc. of the 6th Intl. Conf. on Computational Intelligence, Dortmund, 1999, LNCS 1625
[Zha94] J. Zhang, D. A. Norman, “Distributed Cognitive Tasks”, Cognitive Science, pp. 84-120, 1994