We are IntechOpen,
the world's leading publisher of
Open Access books
Built by scientists, for scientists

4,800
Open access books available

122,000
International authors and editors

135M
Downloads

Our authors are among the

154
Countries delivered to

TOP 1%
most cited scientists

12.2%
Contributors from top 500 universities

CLARIVATE ANALYTICS

BOOK CITATION INDEX

INDEXED

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Content Based Retrieval Systems in a Clinical Context

Frederico Valente, Carlos Costa and Augusto Silva

Additional information is available at the end of the chapter

1. Introduction

Nowadays, digital images and protocols stand as a cornerstone of most modern health-care systems where they are used to provide important data and insights into the inner workings and ailments of the human body. The recent appearance of new modalities, the devices responsible for data acquisition, such as the fMRI[1] and the MDCT[2], produce copious amounts of data [1]. This, coupled with recent advances in storage technology has had as consequence an explosion in the amount of data produced at medical imaging institutions. For instance, the Geneva Hospital alone produced, in 2006, over 50000 images per day and such numbers are steadily rising [2]. Given the prevalence of digital images and protocols in the medical arena, we are, nonetheless, still a long way from fully taking advantage of the potential brought up by this digital revolution. The current data explosion makes it troublesome to a practitioner to sift through the imaging repositories while searching for data relevant for his context. This means we have the data, but not the information, which should be readily available to the experts in the area. In fact, data overload has been reported as a problem by practitioners from medium to large imaging institutions [3].

The current methods of data search, such as the ones provided by the standard query mechanisms present in Digital Imaging and Communication in Medicine (DICOM) are sub-optimal, relying on template matching over a limited number of textual fields [4] (which fields are available depend on the specific software backend), and can consequently be improved upon. It is expected that, by providing more refined and robust methods of searching the large image repositories that currently exist, diagnostic accuracy and efficiency can be improved and more accurate and useful Computer Aided Diagnosis (CAD) tools be devised.

---

1 Functional Magnetic Ressonance Imaging

2 Multi-detector computed tomography

A promising approach to solve the data explosion problem is to integrate computer-based assistance into the image querying and storage processes. This brings us into the topic of Content-Based Image Retrieval (CBIR). At its core, CBIR are a set of techniques to extract relevant pieces of information directly from an image or multimedia object itself with minimum (ideally none) intervention from a human.

The overarching goals are to improve the efficiency, accuracy, usability and reliability of medical imaging services within healthcare enterprises by analyzing content extracted directly from raw image data.

2. Picture archive and communication systems

In a medical imaging institution, such as a hospital or a clinic, the set of technologies employed through the processes of archiving, visualizing, acquiring and distributing medical images over a computer network (see Figure 1) is commonly referred to as a Picture Archive and Communication System (PACS). PACS have evolved tremendously since when, as early as 1972, Dr. Richard J. Steckel implemented a minimal imaging system, comprising not much more than a scanner next to a film developer for digitalization of radiographs, a communication protocol to transmit those images and a video monitor to receive and display them [5]. No more than a proof of concept back then, but fast forward to current days, however, and a properly integrated hospital, or an enterprise PACS implementation is now a major undertaking that requires careful planning and several million dollars of investment [6]. Such investment is often required since large PACS commonly have to handle more than 20000 radiological procedures per year each procedure comprising potentially hundreds of distinct images. This means around 10 Terabytes of imaging information are stored per year [7]. It is in such situations that Content Based Image Retrieval systems are expected to provide the largest benefits.

PACS are still a very active field of research where the ever-changing requirements and the desire to provide more efficient services coupled with new ideas. Figure 2 shows a chronological view of the different challenges that arose and of some of the problems in which the research community is currently focusing. The push for CBIR enabled PACS has gained momentum since the late 90's up to the present day, however, even nowadays there are very few such systems currently powering medical institutions.

2.1. Digital imaging and communication in medicine

A major step in the direction of modern PACS was given circa 1985 with the creation of an earlier form of what would become the current DICOM standard. This protocol stands as one of the key protocols involved in medical imaging systems. We can consider it as the glue that holds the equipment and software developed by multiple companies together. It is an expansible, object oriented protocol with support for multiple imaging modalities and respective structured reports that also allows for private data to be embedded in its objects. Of great importance is the fact that it defines how medical image data and correspondent meta-

data is to be stored, retrieved, and transmitted thus enabling communication between devices manufactured by distinct entities within a PACS.
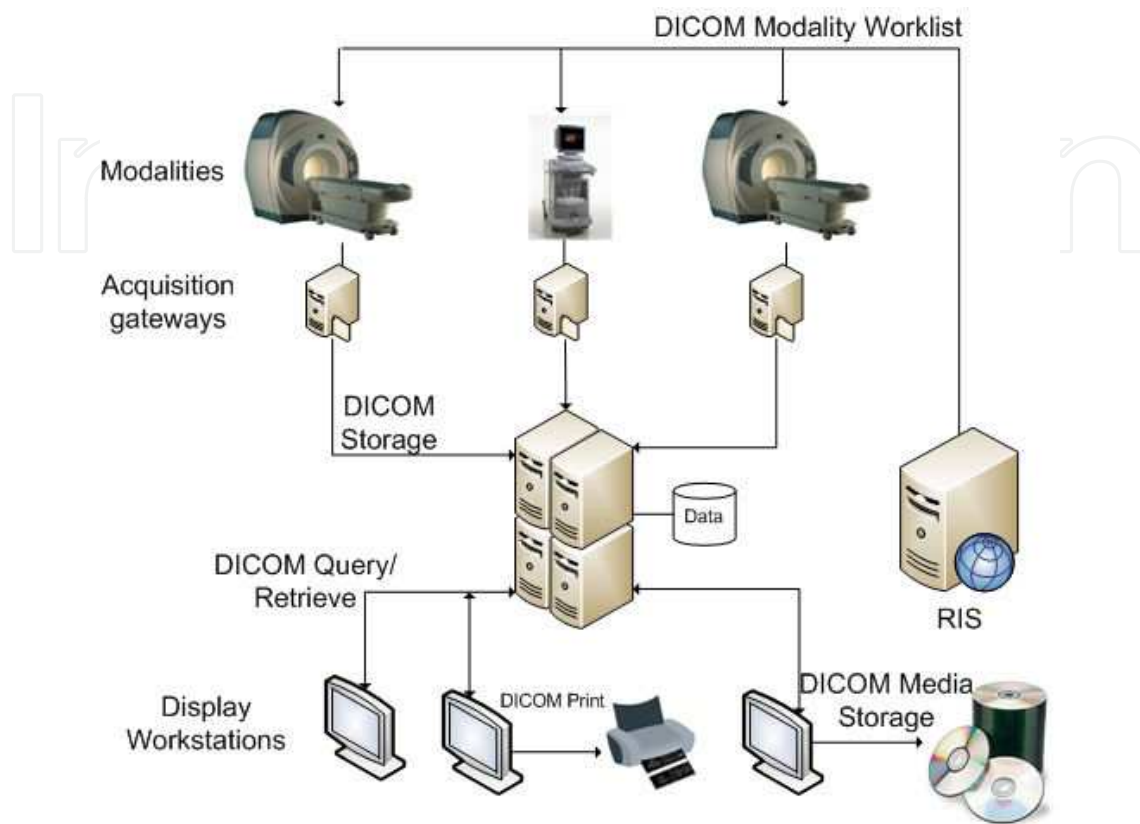

Figure 1. Outline of a PACS infrastructure comprising the most common components in an imaging institution
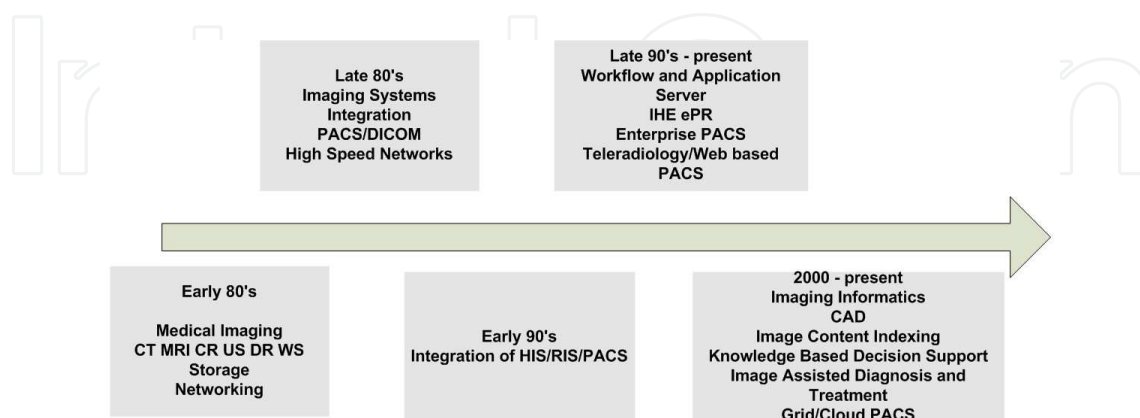

Figure 2. Evolution of PACS research and current trends

The protocol, first proposed by the National Electrical Manufacturers Association (NEMA), in 1983, and currently in its third version, was in itself a major contribution to the exchange of structured medical data [8]. With most available medical equipment providing embedded DICOM support, large sets of medical data have been produced in the DICOM file format. DICOM controls proper image display, it allows a large set of image post-processing from multi-planar reconstruction to the more advanced perfusion analysis, virtual colonoscopy and volume segmentation. This protocol can also be leveraged to enable a PACS-independent way of performing computer-aided diagnosis [9] and knowledge extraction [10]. In practice, much of the ease and flexibility radiologists enjoy today at work is due to this protocol.

However, while DICOM is an open standard, it was created with an eye on the future. As such, it is not set on stone and addends are continuously being added to support new services and modalities. DICOM uses an object oriented approach and its functionality can be extended. This is of great advantage since it allows bridging CBIR with PACS expanding on the DICOM protocol proving the extra functionality with minimal changes in infrastructure.

In DICOM, all data is organized in a patient, study, series and images hierarchy. These are viewed by DICOM as objects with a set of properties or attributes. The definitions for these objects and attributes are standardized according to predefined Information Object Definition (IOD). We can see IODs as templates for objects describing how each particular data object is constructed from attributes. It is then DICOM's group responsibility to maintain a list of all standard attributes and ensure consistency in their naming and composition [8]. The attributes comprise information regarding dates, radiation dosages or any other data of interest. Even image or video data are encoded within a DICOM object as a particular attribute (the attribute (0x7FE0, 0x0010) stands for the pixel data element).

3. Content based image retrieval systems

In its broadest sense, CBIR are systems that help users find similar content to a given image in large and potentially multi-modal repositories [11]. Even extremely large image archives, with often limited textual annotations, can be managed by CBIR as it allows navigation by visual content as opposed to keyword search or the more common form of direct patient/ series searching. An automated approach based on content extraction has the advantage that it needs no manual tagging of images, the features employed are extracted automatically as part of the dataflow, and has the potential to discriminate even very fine details that escape the practitioner. Using information from DICOM Modality Worklists similar studies can be made available to a practitioner without the need for a manual query. Even in the presence of textual information (or rich enough DICOM metadata) content-based methods can potentially improve retrieval by offering additional insight into medical image collections [12]. It is important to note that, while striving to retrieve similar images, CBIR systems, unlike CAD systems, do not attempt to provide a diagnosis.

3.1. Searching information in content based image retrieval systems

Searching relevant data is a fundamental operation in CBIR. In relational databases the search procedure is applied to structured data, that is, numerical or alphabetical information that is searched for exactly. More sophisticated searches such as range queries on numerical keys or prefix searching on strings still rely on the concept that two keys are, or are not, equal. In order to guarantee query performance, traditional databases assume that there exists a total linear order on the keys which is used to establish indexes over the tables. That total order is something that does not arise naturally when dealing with unstructured high-dimensional spaces [13]. However, content-based retrieval relies heavily on similarity queries performed over them [14], hence a similarity function can be defined that establishes that ordering in relation to the source image.

When a query is performed with a source image, every element matches the input with a similarity value. If performed naively, without resorting to advanced indexing techniques, the outcome of similarity query is a permutation of all database content. That is, the elements are rearranged from the highest similarity value to the lowest [13]. This is a behavior that is not desirable. Assuming the similarity is properly defined there are two canonical types of queries that are of interest[3]:

- Range Query: Where we want to retrieve all elements that are closer than a given distance to the query content.

- Nearest Neighbor Query: Where we want to retrieve a certain number of the elements most similar to the query.

In order for a practitioner to perform a search he must provide input to the CBIR. Unlike in traditional query systems text is not used. Several approaches have been explored:

- Query by example – In this type of query a user merely provides a sample image and, relying on its analysis, the engine will provide the user with a set of similar images (see Figure 3).

- Query by region – From an image, the user selects a region of interest comprising the characteristics he is interested in. It is then up to the CBIR engine to retrieve images that share those same characteristics.

- Semantic query – A type of keyword query. However, it is not based on existent metadata but instead relies on mappings between the low level features extracted from an image and a high-level concept. An example would be a search for "micro-calcifications in a fatty tissue breast". Due to its complexity (it is still an unsolved problem) this types of query are only present in research systems.

- Query by sketch – Instead of using an image as source for a query, the user draws something alike what interests him. This methodology has been used to search for works of art in museums and images in the internet, but we know of no use-case in a clinical context.

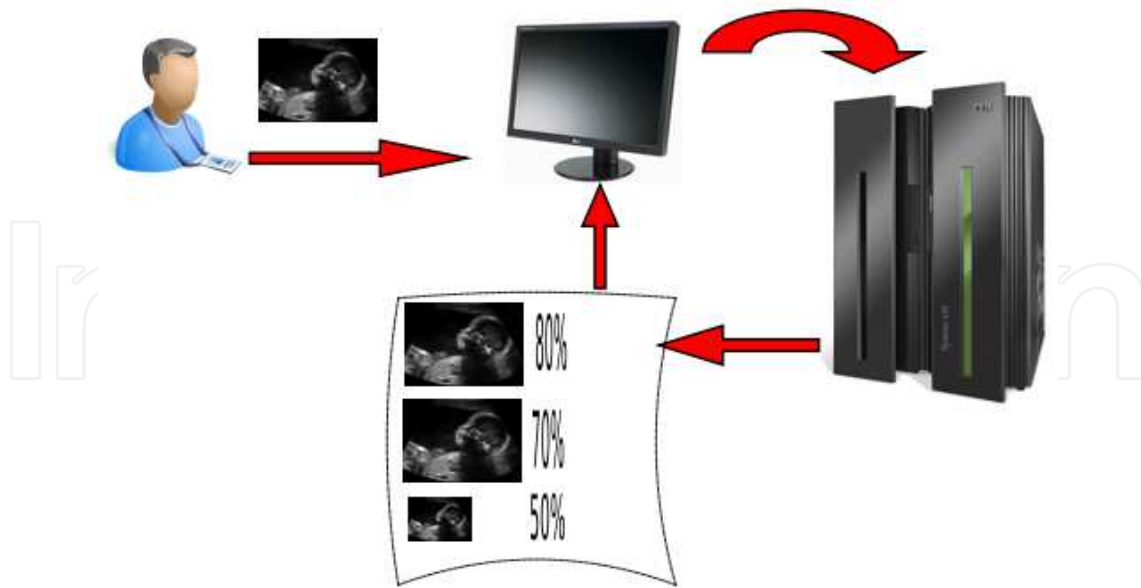3 Other, more complex types of query can be expressed over these.

Figure 3. Query by example

3.2. Content based image retrieval systems in a clinical context

The push for the usage of CBIR of systems in a clinical context comes from their success in other areas where they have been successfully applied to handle large quantities of data. A recent example is Google's "search by image"[4] functionality that operates according to the query-by-example paradigm.

Several scenarios exist where medical practitioners can benefit from the use of these types of system. A key functionality that is of value to radiologists assessing medical images is the ability to provide them with a set of similar images, already diagnosed, thereby aiding them in their process of interpretation by quickly providing them with a second opinion. This proves to be orders of magnitude faster than the current mechanisms provided to manually browse the archives. The potential for this type of assisted interpretation is motivated not only by time constrains, but also by the recognition that variations in interpretation between practitioners, commonly based on perceptual errors, lack of training, or fatigue, do exist [11]. Significant inter-observer variation has been documented in numerous studies [15, 16]. Besides being a useful clinical tool, it is conceivable its use in an academic context where students can benefit from access to similar, diagnosed, data.

Selecting studies by similarity has another benefit. Considering a large repository, built over time, some of the retrieved images are bound to of some age. If a medical institution has kept track of a patient, performing more recent examinations, this data can be very useful to predict possible outcomes to an ailment's evolution. Furthermore, DICOM headers may contain a fairly high rate of errors, for example for the field anatomical region, error rates of 16% have been reported [17]. This hinders the correct retrieval of wanted images via textual search.

―――――――――――――
4 http://images.google.com/

Yet another important and useful outcome of CBIR is the possibility to bridge the semantic gap, allowing users to search an image repository for high-level image features. For instance a researcher may be interested in all studies containing a particular type or disposition of lymph nodes, or query only for images containing a particular feature. This concept expands on CBIR systems and requires that we establish a relation between the low level features employed and the high level concepts of a semantic interpretation.

3.3. Features and feature extraction

At the core of each CBIR there is a matching algorithm analyzing the similarity between the query content and the content stored in the database. However, except for the most trivial CBIR engines operating on simple content, it is not the actual content that is compared. As briefly mentioned, features extracted from the source content, are used instead. In the case of images, pixel by pixel comparisons are not commonly performed due to, not only to the computational effort involved, but also because such comparisons lacks any type of semantic meaning, are dependent on resolution and often are very sensitive to small changes. Furthermore, it is not clear which pixels from the one image correspond to pixels in another image. That said, a feature is simply a relevant piece of information, a synonym for an input variable or an attribute of an image [18], usually much smaller in size than the original data.

Thus, when there is a need to cope with large datasets, such as the ones present in medical repositories, or to deal with large inputs where most information is redundant or irrelevant, as is the case with some images, the analysis is commonly preceded by a pre-processing stage that provides a reduced representation of the original data. This step is called feature extraction and is of crucial importance for any CBIR currently deployed, as content matching operates by comparing features and only the features are indexed.

Using a feature based approach to image analysis brings several advantages to CBIR systems. Besides reducing the size of the input data, thus providing great performance improvements to the matching algorithms, its reduced representation also translates directly in a smaller storage footprint. Of great importance is that, by discarding redundant or useless information, some features can generalize a concept and allow predictive models to become both more general and accurate. Some features also map well into high level concepts (circles, nodes, shapes, nodes) and help bridge the semantic gap.

Generally a feature is represented by a set of values that can be organized into a vector. The global entropy of an image is a single real value but, on the other hand, a normalized intensity histogram or a texture descriptor can be understood as a n-dimensional vector where each index contains the probability of a pixel having an intensity value equal to that index.

In most CBIR systems, a single feature is very often not enough to fully represent the image in a way that makes possible to perform relevant queries. The usual approach is then to extract multiple features from the image and merge them into a single vector, canonically called the feature vector. The set of all possible feature vectors constitutes a feature space. Depending on the features, this can be a space with a very high dimensionality.

The types of features that can be used when designing a CBIR system are essentially limitless and new methods are continuously devised. However most features relate to the original image in a way that can be categorized as presented on table 1.

Criteria	Type	Description
Level of abstraction	Low level	Visual cues, such as color or texture extracted directly from the raw pixel data without any *a priori* information. Examples are edges, corners, contours, brightness histograms.
	Middle level	Regions or blobs obtained as result from image segmentation.
	High level	These types of features contain semantic information about the meaning of an image or the object represented. Usually require knowledge of contextual information and very often imply the use of a classification step. An example would be the number of cars present in an image or the location of nodes in a mammography.
Scope	Local	Features of this type describe a localized region of the image and are usually computed around interest points. A widely used method that uses these types of features is the scale-invariant feature transform.
	Global	Global features comprise information that somehow relates to the entire image. Image entropy is such a feature as is the color histogram.
Representation	Photometric	These are features that explore color and textural cues taken from raw pixel data. A relevant example is Gabor texture descriptors.
	Geometric	Instead of relying on color they employ shape-based cues, most features based on contours are of this type.
Domain	Binary	An on/off type of feature
	Categorical	Instead of having values in a numeric domain this features of this type are aggregated in categories. Usually high level features are also categorical.
	Continuous	Features of this type are represented by a continuous value or vector. Numerical features such as entropy are usually of this type.
	Structural	Where the feature is represented by a graph employed by structural descriptors based on segmentation.
	Vectorial	Sets of continuous values that are related amongst themselves such as histograms, space-based shape descriptors or centers of mass of clusters.

Table 1. Feature taxonomy

The relevancy of a feature is, however, highly dependent on the domain of the problem. This brings the problem of what features should be selected or are relevant in a given context. Namely, what features should be used to perform efficient CBIR in a medical environment where multiple modalities are in place? This is a topic of great interest nowadays and the subject of intensive research.

3.4. Similarity

Due to the unstructured nature of content in images, CBIR systems eschew exact matching and rely instead in nearest neighbor or range queries based on a similarity function. Hence, one of the most important tasks in both the research and development of CBIR systems is to properly define that similarity. Implicitly, one person has a clear notion of whether any two objects or images are similar. Even so, humans are also very much subject to subjective opinions and those can vary wildly. Nonetheless, when searching for reasonable similarity measures, the most obvious place to look is at the human similarity assessment. After all, when a user searches for something similar, he already has in mind his own concept of similarity, whose form is doubtlessly quite different from the metric spaces (such as the Euclidean) typically used for feature vector comparison. The similarity used by the CBIR systems should then be as similar as possible to the human concept of similarity if the results of the search are to be satisfactory [19]. Algorithmically modeling that behavior thus requires that the internal image representations closely reflect the ways in which users interpret, understand, and encode visual data. Finding suitable image representations, based on the types of features described previously is an important step towards the development of effective similarity models [14]. However, creating such algorithmic functions is complicated due to the fact that there is no single model of human similarity. Furthermore a user may have in mind a very specific type of similarity or criteria he is interested in. For instance, in a radiology setting, a practitioner may wish to place more emphasis in finding mammographs sharing a certain disposition of micro-calcifications rather than those containing the same tissue type or having a similar breast size.

Combining multiple representation models can partially resolve this problem. If a retrieval system allows parameterized or multiple similarity functions, the user should be able to select those that most closely model his or her perception [14]. This is not a trivial problem to solve by any means and similarity selection functionality is hardly present in current medical CBIR. In fact such feature is lacking in even most CBIR systems. However, within a medical institution often exist multiple modalities and the DICOM protocol offers support for all those types of distinct imagery.

3.4.1. Similarity measures

Of crucial importance in a CBIR system is the design of the similarity metrics used to match a query to the database feature vectors. Mathematically we can define these metrics as a function $f(x,x')$ that takes as arguments the set of features belonging to two distinct images and returns a value from an ordered set (such as the set of real numbers). This sorting embodies the idea that some images look more like the query than others and al-

low a content engine to return, not only the closest match, but a bundle of images arranged by similarity thus increasing the probability that the user has of finding what he is looking for. Typically, smaller values correspond to higher similarity although that depends, of course, on the specific function being used. The similarity measures employed in CBIR systems are deeply tied with the representation of the features extracted by the system. We now present some of the most applied functions.

- **Vector distance –** One of the most common similarity measures. A function of two feature vectors is defined over the feature space. These are often applied due to their conceptual simplicity. Simpler distances, such as the Euclidean, are also quick to compute, however, that is not always the case, other measures such as the EMD[5] or statistical distances can have significant complexity.

- **Shape based –** These are used when features consist of points delineating a shape boundary. The similarity between shapes is defined in terms of the transformations required to transform a shape into another.

- **Structural/Graph matching -** A class of similarity measurements that apply when the extracted features are represented by a graph. The similarity can be computed by an attributed graph-matching scheme such as relaxation schemes or combinatorial algorithms.

- **Classifier-based -** These classifiers employ machine learning techniques to classify the image as pertaining to a set of predetermined labels. This scheme does not follow the concept of a similarity function, however, in most systems the label obtained is merged with the existing feature set and a vector distance metric is subsequently applied.

In table 2 we find a list containing some of the methodologies employed for similarity measurements in various CBIR projects.

3.4.2. Relevance feedback

While CBIR systems should operate in a transparent manner, in order to increase their overall accuracy it can be desirable to allow a user to relate back to the system which results are actually relevant. Relevance feedback is the process of automatically adjusting an image query using the information provided from the expert on previously executed queries [20]. A way to achieve this goal is to expose to the user an interface that allows him to provide feedback on the relevancy of the results on a per-image basis. A new query can then be executed in order to replace non-relevant results and the feedback loop is repeated many times until the user is satisfied. A key issue is in how to effectively utilize the feedback information to improve the retrieval performance. This aspect depends on the particular implementation of the CBIR given it modifies the way the similarity computation is performed and several methodologies have been explored. An overview of these mechanisms can be found in [21].

---

5 Earth Mover's Distance

3.5. Indexing and performance

When the number of images in the database is small, as is often the case with research systems, a sequential linear search across all elements can provide an acceptable performance. However, with large-scale image databases, such as the ones present in medical systems, more efficient query mechanisms become a necessity. The search task can be significantly improved by relying on multidimensional indexing structures. Like traditional databases, the indexing of an image database should support an efficient search based on the extracted features.

The basic idea behind any indexing procedure (figure 4) is a hierarchical division of space that increases the lookup speed by removing the need to sift the entire feature space o obtain the desired resuls. Due to the nature of CBIR queries, which require quick lookup of the nearest neighbors to a data point in the feature space, the indexing structure must preserve locality.
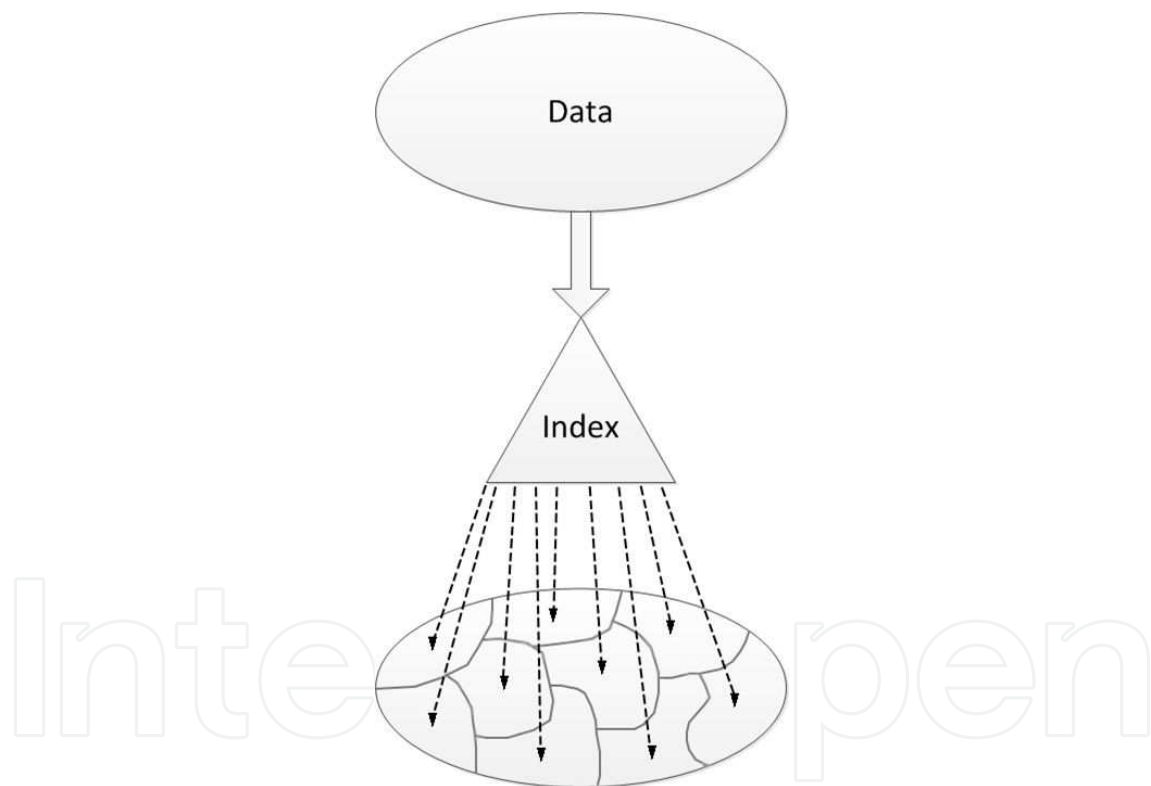

Figure 4. Indexing a feature space

The most popular class of indexing techniques in traditional databases is the B-tree family which provides very efficient searches when the key is a scalar. However, they are not suitable to index the content of images represented by high-dimensional features. Nonetheless, multidimensional indexing techniques exist. There are a large variety of multidimensional indexing methods, which differ in the type of queries they support and the dimensionality of the space where they are advantageous. The R-tree [22] and its varia-

tions are probably the best-known multidimensional indexing techniques in general purpose content retrieval engines. Other approaches are the k-d tree and variants such as the R+-tree and the R*-tree [23].

If a similarity function is at the same time a distance, and thereby a metric to the feature space, a distinct set of methods that operate in a metric space are available. These methods rely only on the definition of the distance function and make no other assumptions. Hence they prove to be very general indexing structures. A study on such methods is available in [13] and [24]. One reason these types of data-structures are not more pervasive in medical CBIR is that research is still being conducted on how to provide mechanisms in Database Management System (DBMS) that allow users to easily incorporate them into search engines.

3.6. Architectural overview of content based retrieval engines

Taking into account the presented requirements and operations for CBIR systems, in figure 5 we show how a generic architecture to a PACS-aware CBIR architecture can be designed. In this architecture the CBIR engine operates outside the PACS repository. This guarantees the integrity of the imaging repository and allows clinical operations to proceed should the CBIR engine fail.
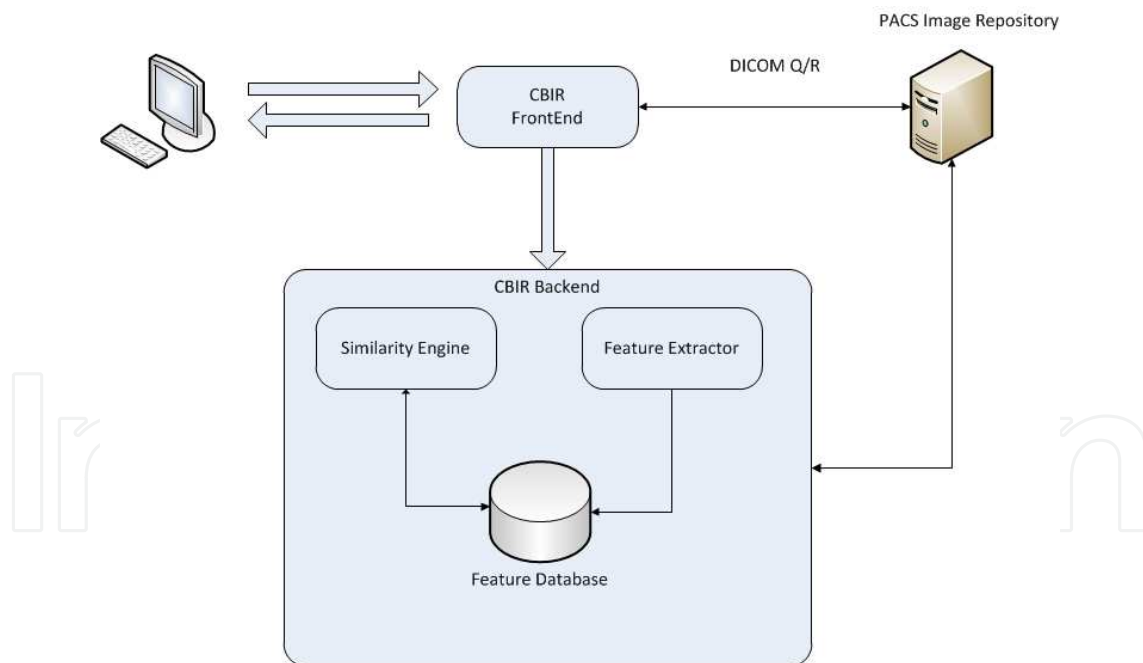

Figure 5. General architecture of a PACS enabled CBIR system

The frontend is the component in charge of receiving similarity requests. Such requests can be triggered manually, by a practitioner, or by analyzing DICOM's Modality Worklist. It replies to the requests with a list of DICOM files to be retrieved which are similar to the source image of the request.

The Feature Extractor component's, main responsibility is, like the name indicates, to extract the relevant features from an image. This behavior is triggered during an initialization procedure, when analyzing images from the PACS repository, or upon the creation of new images by the modalities. The extracted features are then passed to the Feature Database which indexes them and allows for fast nearest neighbor queries. The last major component, the similarity engine comprises the set of metrics and similarities that can be applied.

3.7. Review of CBIR applications

Most radiological CBIR applications are still in a conceptual or research stage. In table 2 we present a brief listing of such systems together with the most important aspects present in a CBIR. This table is based on a similar, more complete table presented in [11].

The features column is arranged in three categories: General, Mixed and Specialized. General features are low and middle level features, extracted from the image with no *a priori* domain specific knowledge and typically extracted with no user input. Mixed features comprise both general features and extra annotations, whether provided from a practitioner or extracted from other sources. Specialized features rely on specific knowledge about the nature and type of the dataset and are typically not automated requiring an expert to provide extra information such as regions of interest.

Ref	Features	Similarity measures	Relevance feedback	PACS Integration
[25]	General	Classifier-based	Yes	No
[26]	General	Classifier	No	No
[27]	Mixed	Classifier	No	No
[28]	Mixed	Vector distance	No	No
[29]	Specialized	Classifier	No	No
[30]	Specialized	Structural	No	No
[31]	General	Vector distance	No	No
[32]	General	Classifier	No	No
[33]	Mixed	Vector distance	No	No
[34]	Mixed	-	No	Yes

Table 2. Overview of medical CBIR systems

Of the presented systems, only [34] focuses specifically on PACS level integration, however, only the concepts and methodology are discussed.

4. Dicoogle

We have developed Dicoogle[6], an open-source PACS with support for data indexing, peer-to-peer communication and CBIR functionality. This tool complements, and may even re-

place a traditional PACS server and enhance it with a more agile indexing and retrieval mechanism [35]. Besides providing basic DICOM services such as Storage and Query/ Retrieval, Dicoogle can automatically extract, index and store all metadata present in a DICOM header (including data present in private data elements). The indexed data can then be queried using free text. A more advanced search mechanism is also provided using a rich query language based on Lucene's syntax. This syntax has support for element selection, numerical and range-based search, wildcard expansion and Boolean operators such as AND, OR and NOT. As a data-extraction tool, Dicoogle has been used in several small to medium imaging institutions. For instance, in [36] Dicoogle was used to demonstrate several inconsistencies in the handling of some DICOM attributes by the modalities and to perform a study on the radiation dosage of the patients handled at the site.

Recently, Dicoogle was extended to support CBIR over a DICOM image repository using a query-by-example paradigm (see figure 6) and following the architectural considerations exposed on the previous sections.
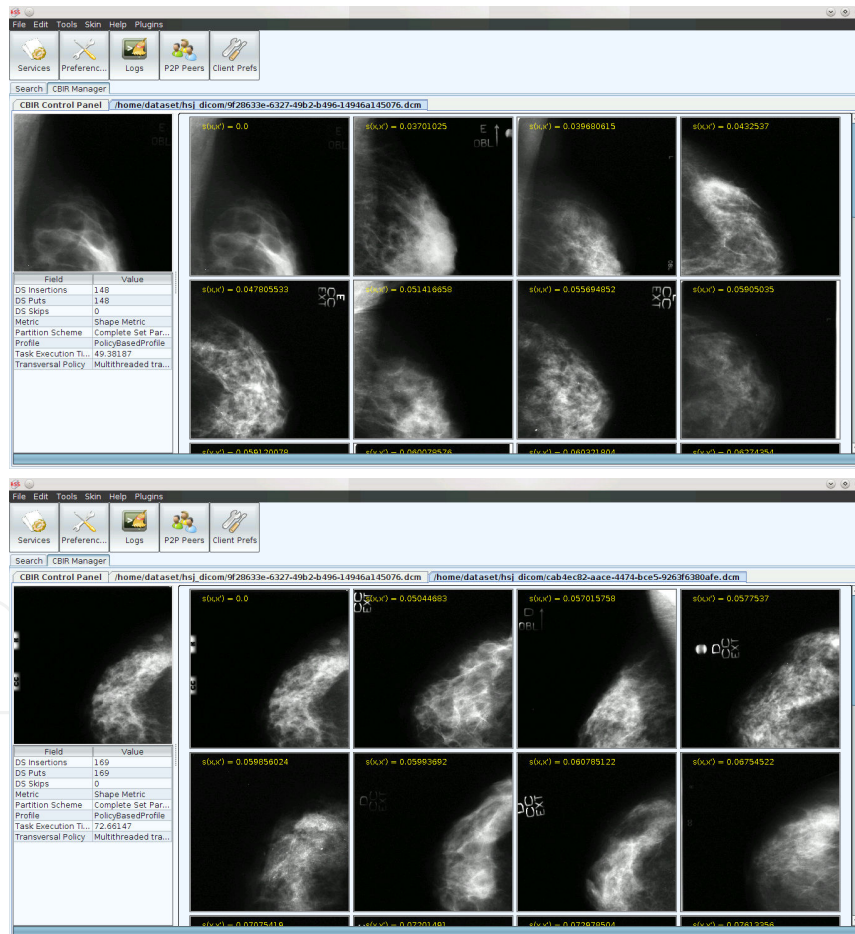

Figure 6. Dicoogle's Query by example results

_____

6 http://www.dicoogle.com

4.1. A profile-based approach to CBIR in a medical context

In the context of multi-modality institutions each modality has distinct criteria to evaluate similarity. Structures identifiable in CT scan images likely have no meaning in the context of mammograms. Similarly, a feature set apt to describe an image within a context of a modality can be entirely useless in another. Likewise for the functions that express the similarity from those features. In a multi-modal environment it seems a needless imposition to use a single set of features and a single measure for similarity, independent of context. Particularly since feature sets coupled with similarity functions can be used to highlight different aspects of an image. In the context of mammographies there is a tendency to focus on micro-calcifications to provide the relevant similarity rather, than, for instance, tissue type or size of breast.

Therefore we've separated the similarity metric from the feature extraction process and provided the user with the concept of "CBIR profiles". A profile contains information on the metric to be used and which features are required to apply it. A profile also contains hints to the indexing mechanism to limit the search space and on which modalities it can be applied. Profiles can be automatically selected using data provided by the DICOM header. Using profiles our CBIR engine allows a practitioner to specify what is of interest to him and fine tune the query if required. In figure 7 we show the dataflow of Dicoogle's CBIR engine.
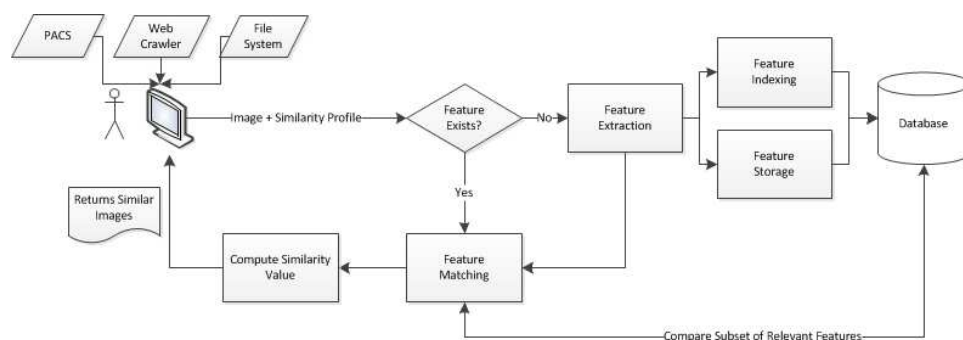

Figure 7. Flow diagram of the interactions between the distinct components of Dicoogle CBIR

5. Challenges and opportunities

In spite of its success in other areas, CBIR is still not a widely deployed technology as a decision support tool. It is the author's opinion that this is currently due to both a lack of integration with the standards that operate through medical institutions, as well as from the stringent requirements that must be fulfilled when operating in an area as critical as the health-care industry. Nonetheless, these types of systems provide enough benefits to the practitioner fully justifying the effort and research applied towards their implementation. Moving towards a clinically useful CBIR in radiology will, however, require a concertated and multi-disciplinary approach. We now point out some challenges that arise from both the general topic of CBIR and its integration with the medical imaging infrastructures.

- **PACS and DICOM integration.** We proposed an approach that complements a PACS by externalizing the CBIR and interacting through the DICOM protocol. However, in this approach, third party tools are limited to images pushed to them through the DICOM C-Move operation. This needlessly hampers the possibility of cooperation between applications as third party tools are either passive, must conform to a private API or must implement themselves an indexing and similarity mechanism.

- **Enabling multi-dimensional database systems.** Several studies exist on how to perform multi-dimensional indexing. However, databases that natively provide support for indexing multi-dimensional data-points according to an arbitrary similarity function and able to cope with large, dynamic volumes of information are, to the best of our knowledge, inexistent. This leads researchers and application developers resorting to either implement indexing mechanisms atop relational databases or to completely ignore the problem and focusing on the other aspects of a CBIR.

- **Multi-modal data integration.** To rely exclusively in pixel data imposes some limits to CBIR systems, not only in the medical arena. Gathering information from multiple sources, such as demography data of a patient from the Radiology Information System, and combining it with the extracted data has the potential to further improve clinical CBIR. A further step forward will be to move from single image analysis and retrieval and merge information from the multitudes of sources that may be present in a DICOM study. Combining this information is not trivial due to missing information, the heterogenetic nature of the data and the problems involved in relating a set of images to another in a meaningful way. This will allow CBIR engines to move to a study-based paradigm, the most common unit of search for practitioners in most modalities. Tackling this problem is likely to have the biggest impact in a clinical environment [37].

- **Lack of a Gold Standard for CBIR.** It is currently not possible to compare the performance across different medical CBIR systems. Not only their respective application domain and specific goals differ, but there is a lack of a common Gold Standard. The Image-CLEFmed[7] is one of the few platforms to evaluate and compare different systems. Other public datasets exist with annotated data, however, the fact they are scattered and the annotations provided do not follow any particular structure makes them cumbersome to use. To access the research effectively, task-related standardized databases on which distinct groups can apply their algorithms are needed. Cooperation with clinical experts is a necessity to provide the necessary relevancy assessments.

- **Towards semantic search.** Much emphasis is being placed on automatic and assisted concept extraction. This effort is directed towards bridging the semantic gap and further increase the accuracy of CBIR systems [38]. An advantage of a CBIR operating in the medical field is that semantics in the medical domain are much better defined and there is a vast accumulation of formal knowledge representations that could be exploited to support semantic search for any specialty areas in medicine [39].

---

7 http://www.imageclef.org/

This idea expands on CBIR systems and requires that we establish a relation between the low level features employed and the high level concepts of a semantic interpretation. Several strategies currently under study are the following [38]:

- Usage of object ontologies to define high-level concepts

- Employ machine learning concepts to perform the association between features and concepts

- Rely on users and perform relevance feedback for continuous learning

- Generate semantic templates (profiles) to support the association

- Rely on the meta-data or other textual information provided by the user

The major functionality enabled by semantic search is the advanced textual queries that can be provided to a practitioner. For example, queries such as "*show me blood smears that include polymorphonuclear neutrophils*"[8] should become a possibility.

6. Conclusion

In this chapter we've exposed the some of the most common methodologies employed in Content-based Image Retrieval and provided an overview of the state of such systems in a clinical context. We pointed out how CBIR, being a query mechanism more adapted to the workflow of a radiologist than the traditional string matching present in DICOM, can help improve diagnosis speed and accuracy in a clinical context.

It was shown how the creation of accurate and performing CBIR systems in a clinical context is a task hard to tackle, ripe with non-trivial challenges that arise from a multitude of factors such as the need for integration with currently deployed PACS, the need to handle multiple modalities and cope with stringent performance requirements.

Furthermore, we presented Dicoogle, our approach to bring CBIR to DICOM enabled PACS systems relying on a profile-based approach.

The goal of CBIR systems is not to replace the practitioner or, unlike CAD tools, provide automated diagnosis, but to empower the expert with tools that allow for faster and more accurate diagnosis in a workflow adapted to his needs.

Acknowledgements

---

8 An example from the ImageCLEFmed competition as shown in [2]

Author details

Frederico Valente, Carlos Costa and Augusto Silva

Universidade de Aveiro, IEETA, Portugal

References

[1] Rubin GD. Data explosion: the challenge of multidetector-row. European Journal of Radiology. 2000;: p. 74-80.

[2] Henning Muller XZADMPJsIaAG. Medical Visual Information Retrieval: State of the Art and Challenges Ahead. 2007 IEEE International Conference on Multimedia and Expo. 2007;: p. 683-686.

[3] Andriole K. Addressing the coming radiology crisis - the society for computer applications in radiology transforming the radiological interpretation process (trip) initiative. Journal of Digital Imaging. 2004;(17): p. 235-243.

[4] National Electrical Manufacturers Association. DICOM Part 4: Service Class Specifications. 2009..

[5] Steckel J. Daily x-ray rounds in a large teaching hospital using high-resolution closed-circuit television. Radiology. 1972; 105(2).

[6] Huang HK. PACS and Imaging Informatics: Basic Principles and Applications: John Wiley & Sons; 2010.

[7] Oliveira MC, Cirne W, Marques PM. Towards applying content-based image retrieval in the clinical routine. Future Generation Computer Systems. 2007 March; 23(3): p. 466-474.

[8] Oosterwijk H, Gihring P. DICOM basics.: Tech; 2002.

[9] CAD–PACS integration tool kit based on DICOM secondary capture, structured report and IHE workflow profiles. Computerized medical imaging and graphics : the official journal of the Computerized Medical Imaging Society. 2007 June; 31(4).

[10] Lehmann TM, Güld O, Deselaers , Keysers , Schubert , Spitzer , et al. Automatic categorization of medical images for content-based retrieval and data mining. Computerized Medical Imaging and Graphics. 2005 March; 29(2).

[11] Akgül B, Rubin DL, Napel , Beaulieu CF, Greenspan , Acar. Content-Based Image Retrieval in Radiology: Current Status and Future Directions. Journal of Digital Imaging. 2011 April; 24(2).

[12] Lew MS, Sebe , Djeraba , Jain. Content-based multimedia information retrieval: State of the art and challenges. ACM Transactions on Multimedia Computing, Communications, and Applications. 2006 February; 2(1).

[13] Chavez E, Navarro , Baeza-Yate , Marroquin JL. Searching in metric spaces. ACM Computer Survey. 2001; 33.

[14] Castelli V, Bergman LD. Image databases: search and retrieval of digital imagery: Wiley; 2002.

[15] Siegle RL, Baram EM, Reut SR. Rates of disagreement in imaging interpretation in a group of community hospitals. Academic Radiology. 1998; 5(3).

[16] Soffa D, Lewis R, Sunshine J, Bhargavan M. Disagreement in interpretation: a method for the development of benchmarks for quality assurance in imaging. Journal of the American College of Radiology. 2004; 1(3).

[17] Guld MO, Kohnen , Keysers , Schubert. Quality of dicom header information for image categorization. 2002..

[18] Guyon I. Feature extraction: foundations and applications (Studies in fuzziness and soft computing): Springer-Verlag; 2006.

[19] Santini , Jain. Similarity queries in image databases. In IEEE Computer Society Conference on Computer Vision and Pattern Recognition; 1996.

[20] Pinjarkar L, Sharma , Mehta. Comparative Evaluation of Image Retrieval Algorithms using Relevance Feedback and it's Applications. International Journal of Computer Applications. 2012 June; 48(18).

[21] Pinjarkar , Sharma , Mehta. Comparison and Analysis of Content Based Image Retrieval Systems Based On Relevance Feedback. Journal of Emerging Trends in Computing and Information Sciences. 2012 July; 3(6).

[22] Qian , Tagare. Optimal embedding for shape indexing in medical image databases. In Medical Image Computing and Computer-Assisted Intervention - MICCAI 2005; 2005.

[23] Beckmann , Kriege HP, Schneid. The r*-tree: an efficient and robust access method for points and rectangles. In SIGMOD; 1990. p. 322-331.

[24] Zezula , Amato , Dohnal , Batko. Similarity Search: The Metric Space Approach: Springer; 2006.

[25] Bhattacharya M, Desai B. A framework for medical image retrieval using machine learning and statistical similarity matching techniques with relevance feedback. IEEE Transactions in Information Technologie in Biomedicine. 2007; 11(1).

[26] Greenspan H, Pinhas A. Medical image categorization and retrieval for PACS using the GMM-KL framework. IEEE Transaction in Information Technologies in Biomedicine. 2007; 11(2).

[27] Lim J, Chevallet J. Vismed: A visual vocabulary approach for medical image indexing and retrieval. In Second Asia Information Retrieval Symposium; 2005; Jeju Island.

[28] Müller H ea. The Use of MedGIFT and EasyIR forImageCLEF 2005. in Accessing Multilingual Information Repositories. In Accessing Multilingual Information Repositories.: Springer p. 724-732.

[29] Petrakis E, Faloutsos C, Lin K. ImageMap: an image indexing method based on spatial similarity. IEEE Transactions in Knowledge Data Engineering. 2002; 15(5).

[30] Amores J, Radeva P. Registration and retrieval of highly elastic bodies using contextual information. Pattern Recognition Letters. 2005; 26(11).

[31] Mohammad-Reza et al. Content-based image database system for epilepsy. Computer Methods and Programs in Biomedicine. 2005; 79(3).

[32] Alto H, Rangayyan R, Desautels J. Content-based retrieval and analysis of mammographic masses. Journal of Electronic Imaging. 2005; 14(2).

[33] Kim J et al. A new way for multidimensional medical data management: Volume of interest (VOI)-based retrieval of medical images with visual and functional features. IEEE Transactions in Information Technologies in Biomedicine. 2006; 10(3).

[34] Benedikt Fischer et al. Integration of a Research CBIR System with RIS and PACS for Radiological Routine. In Proceedings of the SPIE; 2008.

[35] Costa CaFCaBLaRLaSAaOJ. Dicoogle - an Open Source Peer-to-Peer PACS. Journal of Digital Imaging. 2011; 24(5): p. 848-856.

[36] Santos , Bastião L, Costa C, Silva A, Rocha N. DICOM and Clinical Data Mining in a Small Hospital PACS: A Pilot Study. In Communications in Computer and Information Science; 2011: Springer Berlin Heidelberg. p. 254-263.

[37] Medical Visual Information Retrieval: State of the Art and Challenges Ahead. In IEEE International Conference on Multimedia and Expo; 2007; Geneva. p. 683 - 686.

[38] Ying , Zhang , Lu , Ma WY. A survey of content-based image retrieval with high level semantics. Pattern Recognition. 2007 January; 40(1).

[39] Zhou XS, Zillner S, Moeller M, Sintek M, Zhan Y, Krishnan A, et al. Semantics and CBIR: a medical imaging perspective. In CIVR '08 Proceedings of the 2008 international conference on Content-based image and video retrieval; 2008. p. 571-580.