

## Content-Based Image Retrieval in Radiology: Current Status and Future Directions

Ceyhun Burak Akgül,<sup>1</sup> Daniel L. Rubin,<sup>2</sup> Sandy Napel,<sup>2</sup> Christopher F. Beaulieu,<sup>2</sup> Hayit Greenspan,<sup>3</sup> and Burak Acar<sup>1</sup>

Diagnostic radiology requires accurate interpretation of complex signals in medical images. Content-based image retrieval (CBIR) techniques could be valuable to radiologists in assessing medical images by identifying similar images in large archives that could assist with decision support. Many advances have occurred in CBIR, and a variety of systems have appeared in nonmedical domains; however, permeation of these methods into radiology has been limited. Our goal in this review is to survey CBIR methods and systems from the perspective of application to radiology and to identify approaches developed in nonmedical applications that could be translated to radiology. Radiology images pose specific challenges compared with images in the consumer domain; they contain varied, rich, and often subtle features that need to be recognized in assessing image similarity. Radiology images also provide rich opportunities for CBIR: rich metadata about image semantics are provided by radiologists, and this information is not yet being used to its fullest advantage in CBIR systems. By integrating pixel-based and metadata-based image feature analysis, substantial advances of CBIR in medicine could ensue, with CBIR systems becoming an important tool in radiology practice.

**KEY WORDS:** Content-based image retrieval, imaging informatics, information storage and retrieval, digital image management, decision support

### INTRODUCTION

Diagnostic radiologists are struggling to maintain high interpretation accuracy while maximizing efficiency in the face of increasing exam volumes and numbers of images per study.<sup>1</sup> A promising approach to manage this image “explosion” is to integrate computer-based assistance into the image interpretation process. While substantial progress has been made in computer-aided diagnosis/detection (CAD) for lesions such as breast masses, lung nodules, and colonic polyps; current

CAD methods target very specific image features, limiting their broader application to many other scenarios where assistance with image interpretation could be beneficial.

Medical image interpretation consists of three key tasks: (1) perception of image findings, (2) interpretation of those findings to render a diagnosis or differential diagnosis, and (3) recommendations for clinical management (biopsy, follow up, etc.) or further imaging if a firm diagnosis has not been established. The potential for assisted interpretation and decision making is motivated not only by time constraints on readers, but also by the recognition of variations between readers based on perceptual errors, lack of training, or fatigue. Significant inter-observer variation has been documented in numerous studies.<sup>2</sup> For example, in mammographic interpretation, there is variation in sensitivity, specificity, and area under the receiver operating characteristic curve among radiologists.<sup>3</sup> This variation results partly from the complexity of processing the vast amounts of knowledge needed

<sup>1</sup>From the Electrical and Electronics Engineering Department, Volumetric Analysis and Visualization (VAVlab) Lab., Boğaziçi University, Istanbul, Turkey.

<sup>2</sup>From the Diagnostic Radiology, Stanford University, Stanford, CA, 94305, USA.

<sup>3</sup>From the Department of Biomedical Engineering, The Iby and Aladar Fleischman Faculty of Engineering, Tel Aviv University, Ramat Aviv, Israel.

Correspondence to: Burak Acar, Electrical and Electronics Engineering Department, Volumetric Analysis and Visualization (VAVlab) Lab., Boğaziçi University, Istanbul, Turkey; tel: +90-212-3596465; fax: +90-212-2872465; e-mail: acarbu@boun.edu.tr; URL: [www.vavlab.ee.boun.edu.tr](http://www.vavlab.ee.boun.edu.tr)

Copyright © 2010 by Society for Imaging Informatics in Medicine

Online publication 8 April 2010

doi: 10.1007/s10278-010-9290-9

to interpret imaging findings. Much of radiological practice is currently not based on quantitative image analysis, but on “heuristics” to guide physicians through rules-of-thumb.<sup>4</sup> Such heuristics can fail in a variety of circumstances where combinations of features related to diagnosis do not fit expected patterns and practitioners do not recognize the impact of such circumstances. In addition, the heuristics and their use are subject to inter- and intra-reader variability.

While CAD systems are primarily designed to enhance the perceptual component of image interpretation and have been developed for several domains of radiology,<sup>5</sup> decision support systems are much broader; they do not focus on detection, but rather on the reasoning process that radiologists go through after detecting an abnormality, often called a “finding.” Decision support integrates the imaging finding(s) with a formal model (or knowledge base) representing disease processes, ideally aiding arrival at an accurate diagnosis.<sup>6,7</sup> Radiologists always utilize broader patient-specific or demographic knowledge, such as clinical history or results of other tests, in their decision-making processes; as such, it is expected that decision support systems would incorporate these data as well.

Another emerging technique that may assist radiology interpretation is content-based image retrieval (CBIR). In its broadest sense, CBIR helps users find similar image content in a variety of image and multimedia applications. CBIR applications in multimedia can save the user’s time considerably in contrast to tedious, unstructured browsing. The role for CBIR in medical applications is potentially very powerful: in addition to enabling similarity-based indexing, this framework could provide computer-aided diagnostic support based on image content as well as on other metadata associated with medical images. However, despite its success outside medicine, CBIR has had little impact on radiology to date. Current work in image processing, medical informatics, and information retrieval domains provides building blocks that can markedly increase the relevance of CBIR to radiology practice in the future.

At present there is a substantial gap between CBIR, and its focus on raw image information, and decision support systems, which typically enter the workflow beyond the point of image analysis itself. This gap represents what we believe is a

major opportunity to develop decision support systems that integrate image features exploited in CBIR systems. With such integration, CBIR may be a starting point for finding similar images based on pixel analysis, but the process would be augmented by inclusion of image and nonimage metadata as well as knowledge models, broadening the system from “image based search” to “patient” or “case-based” reasoning.<sup>8</sup>

This article reviews the status of CBIR in radiology and highlights some challenges and opportunities to be addressed to achieve significant scientific and clinical impact in the years to come. In “[Current CBIR Technology](#)” section, we provide an overview of the CBIR technology along with a general description of its key components. In “[State-of-the-Art Medical CBIR Systems](#)” section, we present an analysis of existing medical CBIR systems. The analysis has been carried out along several axes such as the way visual features are employed for content description, the methods to measure/quantify similarity between medical images, the use of statistical classifiers in a user-interactive context, and medical image segmentation. In “[Challenges and Opportunities for CBIR in Radiology](#)” section, we first point out the challenges faced by the generic CBIR technology in the radiology domain, and then we describe some of the opportunities to advance CBIR in context. In “[Strategies for Moving Forward in the Next Decade](#)” section, we propose a strategic vision to make CBIR operationally useful in radiology in the near future. “[Conclusions](#)” section presents our conclusions.

## CURRENT CBIR TECHNOLOGY

Content-based image retrieval has been a vigorous area of research for at least the last two decades. The abundance of publications within this period reflects diversity among the proposed solutions and the application domains (see the extensive surveys in <sup>9–11</sup>). CBIR has been most successful in nonmedical domains, e.g., the QBIC system employed in the Russian Hermitage Museum (<http://www.hermitagemuseum.org/fcgi-bin/db2www/qbicSearch.mac/qbic?selLang=English>), the ALIPR (<http://alipr.com/>) system enabling automatic photo tagging and visual search on the web, and

the RIYA (<http://www.riya.com/>; <http://www.like.com/>) system for visual shopping.

A generic CBIR system has two main components. The first component represents the visual information contained in image pixels in the form of *image features/descriptors* and aims at bridging the gap between the visual content and its numerical representation. These representations are designed to encode color and texture properties of the image, the spatial layout of objects, and various geometric shape characteristics of perceptually coherent structures within the image. The second component provides for assessment of *similarities* between image features based on mathematical analyses, which compare descriptors across different images. Ideally, the computed similarity measures should at least partly parallel the similarity between images when judged by the human visual system alone. In a typical CBIR system, database images are returned and displayed in decreasing order of their computed similarity to a query image provided by the user. Thus basically whatever the application domain is, a CBIR system must provide a means for (a) describing and recording the image content based on pixel/voxel information (image features/descriptors) and (b) assessing the similarity between the query image and the images in the database.

#### Image Features/Descriptors

Image features/descriptors are derived from visual cues contained in an image. They are represented as alpha-numeric data in different formats such as vectors or graphs, which stand as compact surrogates for the visual content. One can distinguish two types of visual features. *Photometric* features exploit color and texture cues and they are derived directly from raw pixel intensities. *Geometric* features, on the other hand, make use of shape-based cues. In the following, we describe these cues in detail. Table 1 provides a summary of visual features used in the medical domain.

**Color** The use of color cues in image description dates back to one of the earliest CBIR proposals.<sup>12</sup> A global characterization of the image can be obtained by binning pixel color components (in an appropriate color space, e.g., hue–saturation–illumination) into a histogram<sup>13–16</sup> or by dividing the image into subblocks, each of which is then

attributed with the average color component vector in that block.<sup>17–19</sup> While color is one of the visual cues often used for content description,<sup>10,11</sup> most medical images are grayscale. True color-based characterization is applicable only where color photographs are used for diagnosis, such as in ophthalmology, pathology, and dermatology, or when color is used to scale flow velocities or intensity scales such as in nuclear cardiology. Thus for the majority of medical images, color features will not be useful in image retrieval.

**Texture** Texture features encode spatial organization of pixel values of an image region. The common practice to obtain texture-based descriptors is to invoke standard transform domain analysis tools such as Fourier transform, wavelets, Gabor, or Stockwell filters on local image blocks.<sup>14,15,20–22</sup> In addition, one can also derive the so-called Haralick's texture features such as energy, entropy, coarseness, homogeneity, contrast, etc., from a local image neighborhood<sup>16,20,21,23,24</sup> or utilize linear system approaches such as simultaneous autoregressive models.<sup>25</sup> In the medical domain, texture-based descriptors become particularly important as they can potentially reflect the fine details contained within an image structure. For example, cysts and solid nodules generally have uniform internal density and signal intensity characteristics, while more complex lesions and infiltrative disorders have heterogeneous characteristics. Some texture features may be below the threshold for humans to appreciate, and computers may be able to extract important texture and pattern information that is not readily visible.

**Shape** We use the term *shape* to refer to the information that can be deduced directly from images and that cannot be represented by color or texture; as such, shape defines a complementary space to color and texture. A powerful way of representing shape is through perceptually grouped geometric cues such as edges, contours, joints, polylines, and polygonal regions extracted from an image. Such a grouping can serve as a spatial layout or as a rough sketch by additional postprocessing. It has been successfully used in nonmedical domain by the CIRES (<http://amazon.ece.utexas.edu/~qasim/>) system,<sup>26</sup> which is based upon a combination of higher-level and lower-level computer vision principles. While the term lower-level refers to basic image features (color and/or texture, see

Table 1. A Summary of Image Features/Descriptors Used in the Medical Domain

Category	Representations/cues	Examples
Photometric	Grayscale and color	Histograms <sup>13-16</sup> Moments <sup>21,24</sup> Block-based <sup>17-19</sup>
	Texture	Texture co-occurrence <sup>16,20,21,23,24</sup> Fourier power spectrum <sup>21</sup> Gabor features <sup>15,20</sup> Wavelet-based <sup>14</sup> Haralick's statistical features <sup>32</sup> Tamura features <sup>18</sup> Multiresolution autoregressive model <sup>13</sup>
Geometric	Point sets	Shape spaces <sup>33</sup>
	Contours/curves	Polygon approximation <sup>34</sup> Edge histograms <sup>16,24,32</sup> Fourier-based <sup>13,16,34</sup> Curvature scale space <sup>35</sup>
	Surfaces	Level sets/distance transforms <sup>20,36</sup> Gaussian random fields <sup>37</sup>
	Regions and parts	Statistical anatomical parts model <sup>38</sup> Wavelet-based region descriptors <sup>39</sup> Spatial distributions of ROIs <sup>40</sup>
	Other	Global shape (size, eccentricity, concavity, etc.) <sup>16,17,32,41,42</sup> Morphological <sup>20,42,43</sup> Location and spatial relationships <sup>17,20,41,42</sup>

above), higher-level analysis benefits from perceptual organization, inference and grouping principles to extract information describing the structural content of an image. In cases when the image contains object(s) that can be clearly separated from the background or surroundings, these geometric cues can be used to extract powerful complete shape descriptors. This approach has empirically proven to be very successful in applications such as visual shopping <http://www.riya.com/>; <http://www.like.com/>, logo and trademark retrieval,<sup>27</sup> handwritten digit recognition,<sup>27</sup> and retrieval of three-dimensional computer models.<sup>28</sup> Single object-based shape description has also been successfully employed in several medical CAD applications such as in CT colonography<sup>29,30</sup> and lung nodule detection.<sup>31</sup>

A shape-based representation of the image content in the form of point sets, contours, curves, regions, or surfaces should be available for the computation of shape-based features. Such representations are not usually available in the data directly. Accordingly, as a first step of geometric feature computation, a suitable shape representation should be extracted from the pixel intensity information by region-of-

interest detection, segmentation, and grouping. In the medical context, the shape information is one of the strongest factors in detecting a certain disease/lesion or in understanding its evolution. Thus shape-based descriptors are likely to be useful to fulfill the fine detail requirement of medical image retrieval. However, most of the current medical CBIR systems do not exploit the full potential of the shape information as they either use indirect correlates of the shape cue such as texture measurements or employ very global and simple shape description schemes which are incapable of capturing the required classification granularity. The segmentation problem can be seen as the main obstacle toward the use of more elaborate methods for shape analysis. Objects of interests such as anatomical structures or lesions are embedded in complex and arbitrary backgrounds, in which case robust and automatic segmentation presents a great challenge.

In addition to the preceding quantitative visual features of the image content, there are *semantic* image features that can be derived from a human expert's observations—the radiologist viewing images. Radiologists describe a variety of information in images through annotations that provide

essential semantics describing image content. In addition to these, images can be associated with related information about the individual from which the image was obtained, such as laboratory reports, clinical diagnosis, demographic information, etc. These image and nonimage data, acquired automatically or through manual annotations, are very informative to radiologists in providing their interpretations and thus could provide useful information for CBIR systems as well.

### Similarity Measures

Image similarity measures generally assess a *distance* between (sets of) image features. Intuitively, shorter distances correspond to higher similarity. The choice of metric depends on the type of image features/descriptors as well as on their representation.

The simplest feature representation is a high-dimensional vector space. Metrics defined on vector spaces (e.g., the Euclidean distance) are common similarity measures. Many CBIR systems employ such vector distances due to their computational simplicity. Despite their popularity, distance definitions require the features to be continuous within a range for the computed distances to correspond to perceptually coherent similarities. This requirement is equivalent to assuming that a linear combination of two feature vectors is a valid feature vector corresponding to a valid shape (and/or a semantically similar image), which cannot be generalized. Instead, the metric used should define a convex space of *semantically similar* images. This calls for the concept of *manifolds* and manifold-learning techniques.<sup>44–46</sup> Furthermore, in cases where there are possible mismatches between corresponding dimensions of the above-mentioned feature vectors, the standard vector space metrics yield erroneous results. One such case is when the image content is described by multidimensional feature histograms, which are obtained by accumulating the count of feature vectors falling into a fixed number of bins of predefined size, this representation can be rendered more expressive by allowing the bins to vary in number and size. When comparing two such histograms, however, usual vector distances cannot capture true similarity since the bins are of different sizes and they are not aligned. The Earth

Mover's Distance (EMD)<sup>47</sup> has been proposed as a solution. EMD takes the variable sizes of the bins into account and remedying the correspondence problem by computing the optimal alignment between the two multidimensional histograms.

Alternative to vector-based description, the graph-based representation of image features is a powerful technique that is capable of representing not only local color/texture/shape features, but also their interrelations such as their relative spatial distribution within the region-of-interest. They require a generic class of special computational approaches for similarity assessment, collectively referred to as graph matching.<sup>48</sup>

In some medical applications, subtle geometrical differences between imaged structures may be of diagnostic importance to experts. In such cases, a promising approach is to define similarity through the notion of elastic deformations required to transform one shape into another.<sup>49,50</sup> The energy required to transform one shape into another is assumed to be inversely proportional to similarity.

Similarity can also be measured through the use of statistical classifiers that categorize new instances using high-level information extracted from a training set of instances with known labels. This constitutes a promising attempt to close the so-called *semantic* gap between the visual description of an image and its meaning.<sup>11</sup> The semantics of an image, defined here as its possible interpretations, is task-dependent. It varies based on what is sought. Consequently, different classifiers need to be trained for different tasks, even on the same dataset. Assuming that the images in the database are categorized, one can employ this measure of attachment as a pseudo-similarity measure between the test image and all those in the database. The limitation is that a fixed set of *labels* is required. Relevance feedback techniques that have been widely investigated in nonmedical domains<sup>51–54</sup> address this limitation by more flexible user-centric labeling schemes. In relevance feedback, during the search session, the user labels a few database items that he/she thinks relevant or irrelevant to his/her query. These labeled items, when employed with multiple-instance learning-based approaches, can serve as high-level information about what the user has in mind in that particular search session. More specifically, the query model can be iteratively updated so that the similarity measure can be refined implicitly as in

the Accio! system<sup>52</sup> or the retrieval machine can directly optimize the parameters of the similarity model via statistical learning as in.<sup>51,54</sup> Either way, the requirement of a fixed set of labels is removed since the labels are user-specific and session-dependent. While tested only in a few medical retrieval systems (see “[Relevance Feedback](#)” section), relevance feedback can potentially offer several advantages in the medical domain.

## STATE-OF-THE-ART MEDICAL CBIR SYSTEMS

Table 2 presents a summary of a large set of current systems. The table is organized in five columns addressing the most important four components of a medical CBIR system and the imaging modality the system is designed for. The modality defines the type of information that can be extracted from medical images. This limits usable image representations and features. Consequently, the modality targeted by a medical CBIR system is important. Few medical CBIR systems support multimodality.<sup>15,55,56</sup> We have included some selected medical image classification systems as their components are relevant to CBIR, which can alternatively be viewed as a classification task with respect to the query image.

### Image Descriptors

We grouped image descriptors into three operational categories: general, mixed, and specialized descriptors.

*General* descriptors, such as color, texture, pixel/gradient histograms, etc., are commonly used in all CBIR systems, irrespective of the application domain. They form generic and rather robust feature sets. Since no a priori domain-specific knowledge is exploited in their computation, they are primarily used with vector distance-based similarity analyses<sup>14,40,55,57</sup> in high-dimensional feature vector spaces or with statistical classifiers.<sup>21,24,32,40,55,57,58</sup> However, in a few cases, these general descriptors have been used with elastic deformation-based<sup>59</sup> and graph matching-based<sup>55</sup> similarity analyses. In the former case, a descriptor is continuously deformed into another so that the deformation energy serves as a similarity measure between the two descriptors. In the latter case, general descriptors describe

segmented image regions that are represented as attributed nodes of a graph.

*Mixed* descriptors not only use the general descriptors but also the annotations.<sup>15,17,56,60</sup> Manual or automatic, the type of annotations may vary from textual descriptions of image to diagnoses. Annotations are very powerful descriptors in filling the *semantic gap* in CBIR systems. For instance in<sup>56</sup>, the system (<http://www.dim.hcuge.ch/medgift/>) aims at bringing the standard information technology tools and infrastructures within the reach of medical practitioners. The strategy adopted in its retrieval implementation relies on the GNU GIFT tool (<http://www.gnu.org/software/gift/>), which itself borrows ideas from the text retrieval domain. Accordingly, an image is considered as a document containing several *visual words*. Visual features derived from local and global color/texture cues are mapped to these keywords in order to describe the image in terms of a set of words. As such, the system relies on an annotation tool. With such a text-based description, the image retrieval problem turns into one of standard text-based information retrieval. In<sup>15</sup>, the system adopts a content description approach similar to the one considered above. The description scheme relies upon a statistical learning framework, which extracts vocabularies of meaningful medical terms (VisMed terms) associated with the visual appearance of image samples. These terms are segmentation-free image regions described by color and texture features that are meaningful to medical practitioners and that can be learned statistically using a training set of manually cropped image regions. The learned VisMed terms are used to span a new indexing space. A medical image is indexed in terms of a compact spatial distribution (histogram), where each histogram bin corresponds to a VisMed term from the learned vocabulary, and thus has a semantic interpretation.

*Specialized* descriptors exploit the interrelations between feature sets based on domain-specific knowledge. The information that can be represented through the use of these interrelations is complementary to the other descriptors.<sup>16,36,43,61</sup> For example in<sup>16</sup>, the system ([http://cobweb.ecn.purdue.edu/~cbirdev/WEB\\_ASSERT/assert.html](http://cobweb.ecn.purdue.edu/~cbirdev/WEB_ASSERT/assert.html)) addresses the retrieval of pathology bearing regions (PBRs) in lung CT images by a human-in-the-loop approach. Starting from the well-

Table 2. Main Characteristics of the Medical CBIR<sup>a</sup> Systems

Ref.	Descriptors	Similarity measures	Segmentation	Relevance feedback	Notes on modalities and datasets
24	General	Classifier-based	–	Yes	<ul style="list-style-type: none"> <li>• 5,000 images from daily routine (a subset of imageCLEFmed)</li> <li>• 20 categories (different organs, modalities, views)</li> </ul>
21a	General	Classifier-based	Manual	No	<ul style="list-style-type: none"> <li>• 147 ROIs extracted from CT liver images</li> <li>• 76 control vs. 71 pathological</li> </ul>
32	General	Classifier-based	Manual	No	<ul style="list-style-type: none"> <li>• 57 ROIs extracted from mammograms</li> <li>• 37 benign vs. 20 malignant</li> </ul>
16	Specialized	Classifier-based	Interactive	No	<ul style="list-style-type: none"> <li>• 302 lung CT images</li> <li>• 8 lung disease categories</li> </ul>
15	Mixed	Classifier-based	–	No	<ul style="list-style-type: none"> <li>• 8,725 images (CasImage<sup>b</sup> database)</li> <li>• Medical retrieval task in imageCLEF 2004</li> <li>• 26 query topics</li> </ul>
34	Specialized	Procrustes	Manual	No	<ul style="list-style-type: none"> <li>• NHANES II spine X-ray images</li> <li>• 250 vertebra boundary profiles</li> <li>• 10 categories of cervical and lumbar vertebra shapes</li> </ul>
17a	Mixed	Classifier-based	Manual	No	<ul style="list-style-type: none"> <li>• 150 endoscopy images</li> <li>• Several classes from endoscopic findings/diagnoses</li> </ul>
36a	Specialized	Classifier-based	–	No	<ul style="list-style-type: none"> <li>• Brain MR images</li> <li>• Hippocampus in schizophrenia (15 control vs. 15 patients)</li> <li>• Corpus callosum in affective disorder (20 controls vs. 18 patients)</li> </ul>
33	Specialized	Procrustes	Manual	No	<ul style="list-style-type: none"> <li>• NHANES II spine X-ray images</li> <li>• 2,812 vertebra boundary profiles</li> <li>• No classification analysis</li> </ul>
58	General	Classifier-based	Manual	Yes	<ul style="list-style-type: none"> <li>• 76 Mammograms containing clustered microcalcifications</li> <li>• Ground truth similarity obtained from human observer studies</li> </ul>
59a	General	Elastic deformation	–	No	<ul style="list-style-type: none"> <li>• 90 cardiac ultrasound images</li> <li>• View classification</li> </ul>
61	Specialized	Elastic deformation	Interactive	No	<ul style="list-style-type: none"> <li>• 100 intravascular ultrasound images containing calcium plaque structures</li> <li>• Similarity-based retrieval used for improving registration</li> </ul>
14	General	Vector distance	–	No	<ul style="list-style-type: none"> <li>• Abdominal ultrasound images</li> </ul>
57	General	Vector distance Classifier-based	Interactive	No	<ul style="list-style-type: none"> <li>• 70 brain MR images</li> <li>• Hippocampus localization and identification</li> <li>• 10 epileptic patients</li> </ul>
66	General	Classifier	Automated	No	<ul style="list-style-type: none"> <li>• Image categorization and retrieval on 1500 radiological images (IRMA project X-ray library).</li> <li>• 17 radiological X-ray classes</li> </ul>
40a	General	Vector distance	Manual	No	<ul style="list-style-type: none"> <li>• fMRI activation contrast maps used as correlates of Alzheimer's disease</li> <li>• 9 controls vs. 9 patients</li> </ul>
60	Mixed	Classifier-based Vector distance Textual	Automatic	No	<ul style="list-style-type: none"> <li>• 300 VOIs extracted from 13 dynamic PET brain scans</li> <li>• 2 tumor cases, 3 normal, 8 other neurological cases</li> </ul>
43	Specialized	Graph matching	Manual	No	<ul style="list-style-type: none"> <li>• 124 MR images</li> <li>• No classification analysis</li> <li>• Target application: indexing and fast search</li> </ul>
55	General	Vector distance Graph matching Classifier-based	Automatic	Yes	<ul style="list-style-type: none"> <li>• 1,617 radiographs from daily routine</li> <li>• Classification based on image modality, body orientation, anatomic region, biological system</li> </ul>
56	Mixed	Vector distance Textual	–	No	<ul style="list-style-type: none"> <li>• ImageCLEF 2005 medical retrieval tasks</li> </ul>
67	General	Classifier	–	No	<ul style="list-style-type: none"> <li>• Categorization: database of 12,000 radiographs (ImageClef 2007); 11,000 training, 1,000 testing; 116 different categories</li> <li>• Retrieval: database of 66,000 images, 30 query topics (ImageClef 2008). Return ranked set of 1,000 images.</li> </ul>

<sup>a</sup>The citations are *classification* systems with components relevant to CBIR. An exhaustive list of medical image classification is beyond the scope of this survey

<sup>b</sup>[www.casimage.com](http://www.casimage.com)

founded assumption that the extraction of PBRs by automatic image segmentation techniques is difficult, if not impossible, the system lets the user delineate PBRs in lung images. The manually extracted regions are then characterized by their grayscale, texture, and shape attributes, which are organized based on their spatial location with respect to certain anatomical landmarks. The system proposed in<sup>43</sup> describes the images by the so-called attributed relational graphs (ARGs). In an ARG, the image regions are segregated by graph nodes and their spatial relationships are represented by edges between these nodes. As in<sup>16</sup>, the regions of interest are extracted manually but they are not necessarily pathology bearing. Both nodes and edges are labeled by attributes corresponding to the properties of these regions and their relationships, respectively. The attributes can be derived from any type of the visual cues described above, e.g., texture properties, global geometric features of regions (e.g., size, roundness), or features specified in some transform domain (e.g., Fourier coefficients of the extracted region boundary). Typical attributes for the spatial relationships (the edges) can be the distance between two connected regions, their relative orientation, etc. Specialized descriptors are primarily employed with statistical classifiers, however in few cases these descriptors are also used with elastic deformation<sup>61</sup> and graph matching-based<sup>43</sup> similarity analyses.

### Similarity Measures

One of the biggest challenges in any CBIR system is how to define an appropriate measure assessing the similarity to be used for database indexing and/or similarity-based ranking of the retrieved images with respect to the query. A common and rather straightforward method is to employ *vector distances* in a high-dimensional normed vector space, commonly a Euclidean space, in which each image is represented with a point corresponding to its image descriptor/feature vector.<sup>14,40,55–57,60,62</sup>

*Procrustes methods* are used when descriptors consist of landmark points, which usually delineate a shape boundary.<sup>33,34</sup> The shape of object is considered as a member of an equivalence class formed by removing the effects of translation, rotation, and isotropic scaling, collectively denominated as similarity transformations. Thus,

the similarity between two shapes is defined up to a rigid body transformation (translation + rotation) and an isotropic scaling. It is possible to extend this notion of similarity to allow nonrigid deformations of the shapes such as via *elastic matching methods* as in<sup>59,61</sup>. In fact, approaches based on continuous (and even diffeomorphic) mappings are becoming increasingly popular in the medical domain.<sup>49,63</sup> Yet, elastic matching methods are more suitable for cases of abnormalities of overall size and/or shape of an organ, which limits their applicability to some extent.

*Graph matching*, as pointed out earlier, refers to a special class of similarity measurement that is only applicable when the image content is represented by a graph.<sup>43,55</sup> A graph is a natural way of representing features and their interrelations. For instance in<sup>55</sup>, the description (indexing) step uses the selected local features following the Blobworld approach,<sup>64</sup> which outputs a graph-based descriptor of the image. Each node of the graph represents an image region and the corresponding local features of that region are used as the attributes of the node. The similarity between the query and database images is computed by an attributed graph-matching scheme.<sup>48</sup> Basically, a cost function is first evaluated between the graph nodes in terms of the distance between the attributes of the node. The minimum total cost over possible matchings, defined as the sum of the costs between matched nodes, serves as a similarity measure. Combinatorial algorithms and/or continuous relaxation schemes can then be used to find a pairwise matching of the nodes by (approximately) minimizing the total cost.

In contrast to the above methods that directly measure the similarity in terms of image information alone, *classifier-based* similarity measures use the classification of a query image with respect to a fixed set of predetermined labels to assess similarity.<sup>15–17,21,24,32,36,40,55,57,58</sup> Statistical classifiers need to be pretrained. The membership of the query image to each class is usually used as a feature set representing the image content. However, classifiers can also be used as a preprocessing step to narrow the search space in CBIR systems. For example, in<sup>55</sup>, the classifier serves as an automatic medical image annotation tool on its own. It can thus be used to retrieve similar images on a coarse level, e.g., a lung CT would retrieve other lung CT images in the database.

*Textual* similarity measures refer to the common text-based query/retrieval scheme. They can be applied to CBIR through the use of manual/automatic annotations represented as *words*.<sup>56,60</sup> The burden of manual annotations may be an obstacle in their wide-spread use. Accordingly, automatically obtaining medical annotations would thus be of significant importance. In fact, the medical image annotation track in ImageCLEFmed (<http://www.imageclef.org/>; <http://ir.ohsu.edu/image/>), organized as part of the Cross Language Evaluation Forum (CLEF; <http://www.clef-campaign.org/>) since 2004, serves as a benchmarking platform to streamline the research on medical image annotation.

### Segmentation

Segmentation is a key preprocessing step in CBIR systems that describe the image content through regions of interest. The goal is to identify the semantically meaningful regions/objects within an image. Several methods have employed *manual* segmentation to rule out retrieval errors due to wrong segmentation.<sup>17,21,32–34,40,43,58,65</sup> However, despite its advantages, manual segmentation is a tedious task limiting the usability of CBIR systems. This has led to the development of *semi-automatic*<sup>16,57,61</sup> and *automatic*<sup>55,60</sup> segmentation algorithms to extract the regions of interest. In semiautomatic methods, an initial segmentation usually in the form of boundary delineation is provided by the user. The segmentation is iteratively refined and the user can intervene between iterations by correcting the boundaries if they tend to move away from the desired solution. Yet, some systems preferred to design segmentation-free image descriptors to avoid the complications due to imperfect segmentation.<sup>14,15,24,36,38,56,59</sup>

### Relevance Feedback

A promising technique to fill the semantic gap in the medical CBIR systems is to adopt an *expert-in-the-loop* approach. This refers to integrating the physician's high-level expert knowledge into the retrieval process by acquiring his/her relevance judgments regarding a set of initial retrieval results.<sup>24,55,58</sup> A relevance judgment is task-dependent. For instance, an image can be relevant to a certain query in terms of its modality, or the

anatomical region that it belongs to, or the disease that it depicts, etc. These judgments, provided in the form of discrete labels or ordinal/continuous ratings, are used in a statistical learning algorithm to obtain an iteratively refined similarity measure, which is expected, at the end of the search, to become more suitable for a particular search task, e.g., finding all images of a particular organ or retrieving all cases related to a certain disease.

### Performance

A comparative assessment of the performances of the reviewed medical CBIR systems is not possible as not only their application domains (imaging modalities that the systems are built for) differ, but also there is a lack of common database to evaluate different systems. The ImageCLEFmed, as mentioned earlier, is one of the few (if not the single) platform to evaluate and compare different systems. The IRMA,<sup>55</sup> the medGIFT,<sup>56</sup> and the VisMed<sup>15</sup> projects are participants of ImageCLEFmed. The evaluation protocols are largely influenced by the tasks addressed by the IRMA system that mainly targets modality and body part similarity-based retrieval (<http://www.imageclef.org/>; <http://ir.ohsu.edu/image/>; <http://www.clef-campaign.org/>). Objective evaluation is even more challenging for the retrieval systems based on patient/case similarity as there is currently no consensus on how to rate the case similarity even manually. Consequently, task specific evaluation platforms are required rather than a generic approach.

### CHALLENGES AND OPPORTUNITIES FOR CBIR IN RADIOLOGY

In a typical CBIR system (e.g., in multimedia), subtle differences between images are considered as irrelevant for matching, as such, they are often ignored. On the contrary, in medical images such subtle differences can be highly relevant for diagnosis. Thus one of the challenges differentiating radiology CBIR from general purpose multimedia applications is the granularity of classification; this granularity is closely related to the level of invariance that the CBIR system should guarantee. Consequently, CBIR systems in radiology should, at minimum, capture fine details of the image

content, particularly because they are related to diagnostic features used by humans in unassisted diagnosis. In addition, computer-derived features that may not be easily discerned by humans may also be useful.<sup>68</sup> The sheer breadth of imaging findings is a major challenge, as “diseases” are depicted differently by different modalities, they affect a number of organs and tissues, and the findings themselves are highly varied. For example, some tumors are discrete masses while some are poorly defined and infiltrative. Entire organs may be abnormal due to diffuse disease, such as in cirrhosis or leukemia; and some clinically important features only manifest as abnormalities of overall size or shape of an organ or organ part (e.g., the hippocampus in epilepsy), without changing aspects of the image. For human image interpretation, several image features stand out as likely to be the most relevant, though the relative importance of features will vary across modality and disease targets. Some of the key features include lesion shape, boundaries, density or intensity, presence or absence of enhancement with intravenous contrast material, texture, and whether a lesion is solitary or multiple.

The rich and varied information content in medical images and their implicit knowledge about anatomic structures is not leveraged in image-only approaches to CBIR. Many image search methods in fields other than medicine use index terms (metadata) associated with the images, rather than the image data alone.<sup>69,70</sup> The latter type of image search is based on *semantics* (i.e., meaning) as described by annotations. The class of metadata used in retrieval can be extended to include non-image data such as laboratory reports, physiological measurements, etc., which are used routinely by radiologists in providing their interpretations, but are widely ignored in current CBIR systems. This domain-specific metadata largely depends on the radiologists’ observations. The current approach to representing these observations/interpretations is unstructured free text reports. Searches based on these reports are limited because there is no enforcement of controlled terminology, and the linkages between image content and report text are loose or nonexistent. Consequently, efficient and effective annotation tools are needed to generate such image metadata. Nonimage clinical data are stored in electronic medical record (EMR) systems. Such data, when linked to images, can be used to associate

PACS data with corresponding EMR, hence the full potential of compound image and metadata approaches could be exploited. In fact, given the complexity of medical images themselves and the richness of the associated metadata, a combined approach is likely to be advantageous. Including both semantic and pixel content would help identify not only similar images, structures, or lesions, but also find these within similar patients and clinical situations.

Several attempts have been made to establish the connection between the visual content and domain-specific semantic content in medical images. One notable example is the Essence framework,<sup>71</sup> serving as a knowledge repository and exchange platform for medical image databases. In Essence, visual abnormalities and pathologies are extracted and mapped to physician-defined semantic terms using a shared ontology based on the common knowledge from expert radiologists and information from medical references. The system is also capable of refining the shared ontology by adapting the assignment of semantic terms to image features based on individuals’ preferences. Closely related examples include the retrieval strategy adopted in the medGIFT project,<sup>56</sup> employing purely text-based schemes for image retrieval, and the VisMed approach,<sup>15</sup> which extracts vocabularies of meaningful medical terms associated with visual appearance from image samples (see “Challenges and Opportunities for CBIR in Radiology” section), thus providing a mapping from visual feature space to semantic terms. Controlled terminologies such as RadLex<sup>72</sup> have recently been developed, and work is underway to create tools to enable semantic annotation of images using ontologies.<sup>73</sup> The emergence of these technologies provides an opportunity to enhance CBIR systems with richer descriptions of images. A system that attempts at combining nonimage clinical data and image data in a real patient-analysis scenario is the Advanced Analytics for Information Management (AALIM) system.<sup>74</sup> AALIM is a novel multimodal decision support system that seamlessly extracts, analyzes, and correlates information from patient’s EMRs for purposes of decision support. Since health records today have become multimodal, including images, video, text, and charts, AALIM analyzes multiple modalities to identify similar patient records. Current work in AALIM has focused on

the domain of cardiology. Sophisticated feature extraction and search techniques have been developed to extract structured data from a database (demographics, diagnosis codes, vitals, medications), as well as information from free text (cardiology reports), graphical data (EKG waveforms), and imaging studies (ultrasound videos).<sup>74-77</sup> Similar patients are then found based on fusion of information across the multimodal sources. Cohorts of patients with the same condition are extracted from large patient archives, matched on demographics and comorbidities.

We believe that a combination of image-based and metadata-based retrieval will perform much better than either independently, and eventually can help overcome many of the obstacles preventing CBIR from being clinically useful. This is much easier said than done. The image content itself has been extensively studied and ultimately contains a fixed amount of information. Basic metadata is easily linked to the images from clinical information or the DICOM header, such as imaging technique, patient age, weight, and gender. Moving further, there is an enormous amount of (potentially) relevant metadata that can be used to augment the image content itself. This may include structured image descriptors provided by radiologists, laboratory data, pathology, genetic profile, detailed clinical history, family history, and physical exam findings. Such a multidimensional database extends the concept of “image” retrieval to a much broader “case” retrieval, wherein each patient along with their images and associated metadata represents a building block of the database.

If each image-based and metadata approaches are challenging enough in isolation, an integrated system may seem intractable, even grandiose. Content representation becomes complex in dealing with image data itself, mathematical representations of the data, and semantic bridges to visual cues. Nonimage metadata may be text based, alpha-numeric (e.g., lab data), or more complex and multidimensional itself (e.g., pathology results, genetic profiles). Inevitably, the (content) representation in this vast *case space* has to be heterogeneous, yet represented faithfully in a database.

Beyond data collection and storage, a case-based system increases markedly the complexity of making similarity measurements between cases. Here, the major challenge is to combine different types of information in a single and continuous

*case similarity* metric. A useful metric should (a) be continuous as no two cases are the same, thus the similarity measurement cannot be viewed as a classification task, (b) be able to handle the significant differences between types of information, such as scale differences between alpha-numeric features and continuous vs. discrete (e.g. class labels) features, (c) be able to deal with missing data which is expected to be a frequently seen problem due to the high dimensionality of the case space, and (d) should define a convex case space (i.e., case space should be defined as a manifold within the multidimensional feature space).

#### STRATEGIES FOR MOVING FORWARD IN THE NEXT DECADE

CBIR in general has advanced considerably on the heels of the dramatic proliferation of digital media, search technology, and the Internet. Much of what has been learned for nonmedical applications can be ported to the medical domain. Medical images contain varied, rich, and often subtle features that are important clinically and that separate this domain from multimedia applications. We believe that with coordinated efforts and a long-range view, current and future technology will fully enable a system integrating pixel- and metadata-based retrieval of similar images which can then be leveraged for diagnostic decision support. With the appropriate selection of clinically relevant disorders and engagement of front line radiologist users, significant advances in medicine and health can be achieved in the next decades. Progress towards clinically useful CBIR in radiology will require a broad-based and multidisciplinary approach. We propose the following guiding principles as a framework for future work:

- *Focus on clinical importance.* The disorders targeted should be those with a significant impact on human health, such as oncology, wherein improvements in imaging diagnosis and monitoring would be expected to improve treatment outcomes. In that sense, early detection of a disease is a real priority. Furthermore, disorders should be common enough that centers can collect large amounts of data in a reasonable time frame.
- *Focus on areas where current imaging methods are imperfect.* While imaging technol-

ogy itself will continue to improve, areas in which imaging results are inaccurate or ambiguous, such as early detection of hepatocellular carcinoma, should be of high priority.

■ *Create tools for data collection.* Fortunately, computer hardware and storage are no longer barriers to creation of large repositories of data. However, data must be entered systematically and efficiently to achieve collections with millions of images. Automation of segmentation (if needed) or interactive tools requiring minimal user input will be critical. Taking advantage of clinical PACS as the starting point of data collection will also be critical for scaling. Simple text entry or voice-activated tools that map to structured lexicons such as RadLex will be enabling.

■ *Facilitate data sharing.* The creation of large databases will be a key for this effort. Hosting institutions or Federal agencies such as NCI may serve as data centers, assuming all personal health identifiers are scrubbed from case material. Grid or cloud architecture would allow data centers to interconnect. Expertise with networking and database engineering will be necessary.

■ *Leverage image processing expertise.* The vast amounts of knowledge in existence should be leveraged to optimize pixel level analysis and constructions of mathematical representations of the data. Not only single images, but volumetric data and temporal data should be represented.

■ *Leverage bioinformatics developments.* The variety of data types and complex database structures will require specific focus on how to efficiently and effectively search for similarities. This must not just address clinically suspected correlations within the data, but enable discovery of unexpected relationships that may prove to be clinically useful. This expertise is currently being developed in the bioinformatics communities and consists of supervised and unsupervised learning techniques for understanding complex multidimensional data.

■ *Use composite image features and metadata together whenever possible.* As we emphasized in “Challenges and Opportunities for CBIR in Radiology” section, the joint use of visual information in the form of image descriptors and all nonvisual metadata to construct *medical cases*, which are composite entities by themselves, is one of the essential directions to follow for CBIR to be truly beneficial in

medical decision support systems. This integrated data use policy also applies to the choice of image descriptors to employ in a medical CBIR system. One type of image descriptor is extremely unlikely to resolve all the semantic ambiguities inherent to the medical domain. As such, a rich set of image descriptors and metadata should be integrated to augment each other whenever possible.

## CONCLUSIONS

This survey has focused on the CBIR applications in medical domain, with a review of state-of-the-art approaches, discussion of challenges and opportunities in the medical domain, and speculations for future research.

From a technical perspective, we have put forward the diversity of the proposed solutions for image description and similarity assessment techniques that are considered as the two fundamental components of a CBIR system. The majority of the medical CBIR systems have emerged as adaptations of the multimedia CBIR systems, basic technologies of which are discussed in “Current CBIR Technology” section. The medical CBIR approaches reviewed in “State-of-the-Art Medical CBIR Systems” section reflect the continued efforts of researchers working at the crossroads of visual content analysis and medical imaging in order to adapt existing CBIR techniques and to develop dedicated approaches that take into account the special aspects of the radiology domain. However, due to the diversity of medical image content in terms of modalities, human anatomy and diseases as well as the challenges specific to medical domain as compared to the multimedia applications, this diversity is expected to increase even further to meet the requirements of different subdomains within medicine.

Our analysis of state-of-the-art medical CBIR systems has shown that despite the progress in image description and similarity analysis, medical image segmentation remains as a formidable challenge as in other visual domains. Most of the approaches resort to semiautomatic or manual methods. This not only limits the development of more effective solutions for automated and objective analyses of anatomical structures and/or

abnormalities but also adversely affects the scalability of the proposed systems. Another challenge is the lack of large verified datasets for reliable benchmarking, especially for systems having diagnostic purposes. Performance evaluations are usually carried out using a handful of examples, which further delays the adoption of medical CBIR systems in clinical routines.

In “Challenges and Opportunities for CBIR in Radiology” section, we have focused on other challenges that concern the inherent variability in imaging findings. While computer-driven image descriptors and similarity analyses can help reduce the variability, it can be conjectured that image-only approaches have limited resolving power especially for diagnostic purposes. Medical diagnosis is a complex process that should also take into account the rich and varied metadata, patient history, and other relevant (not necessarily visual) information. Accordingly, a medical retrieval system should extend its scope in the type of information it uses and be able to seamlessly incorporate such information. We have conceptualized this data integration problem as finding mixed descriptors and similarity measures on a conceptual *case space* rather than an image-only space. Should appropriate information processing technologies be developed, the information diversity in medicine is not necessarily a curse but may be an opportunity.

Finally, in “Strategies for Moving Forward in the Next Decade” section, we have drawn attention on operational issues in medical CBIR and proposed strategies to tackle them. We believe and hope these guiding principles would help researchers, in this fast growing field, build more useful and reliable medical search/retrieval systems.

#### ACKNOWLEDGMENTS

This work has partly been supported by NIH CA72023 and TÜBİTAK KARIYER-DRESS (104E035).

#### REFERENCES

1. Rubin GD: Data explosion: the challenge of multidetector-row CT. *Eur J Radiol* 36(2):74–80, 2000
2. Siegle RL, et al: Rates of disagreement in imaging interpretation in a group of community hospitals. *Acad Radiol* 5(3):148–154, 1998
3. Barlow WE, et al: Accuracy of screening mammography interpretation by characteristics of radiologists. *J Natl Cancer Inst* 96(24):1840–1850, 2004
4. McDonald CJ: Medical heuristics: the silent adjudicators of clinical practice. *Ann Intern Med* 124:56–62, 1996
5. Doi K: Computer-aided diagnosis in medical imaging: historical review, current status and future potential. *Comput Med Imaging Graph* 31(4–5):198–211, 2007
6. Burnside ES, et al: Bayesian network to predict breast cancer risk of mammographic microcalcifications and reduce number of benign biopsy results: initial experience. *Radiology* 240(3):666–673, 2006
7. Rubin DL, Burnside ES, Shachter R: A Bayesian Network to assist mammography interpretation. In: M.L. Brandeau, F.S. F, and W.P. Pierskalla, Eds. *Operations Research and Health Care*. Kluwer Academic Publishers: Boston, 2004, pp 695–720
8. Kahn CE: Artificial intelligence in radiology: decision support systems. *Radiographics* 14:849–861, 1994
9. Datta R, et al: Image retrieval: ideas, influences, and trends of the new age. *ACM Comput Surv* 40(2), 2008
10. Muller H, et al: A review of content-based image retrieval systems in medical applications—clinical benefits and future directions. *Int J Med Informatics* 73(1):1–23, 2004
11. Smeulders AWM, et al: Content-based image retrieval at the end of the early years. *IEEE Trans Patt Anal Mach Intell* 22(12):1349–1380, 2000
12. Swain MJ, Ballard DH: Color indexing. *Int J Comput Vis* 7(1):11–32, 1991
13. Comaniciu D, Meer P, Foran DJ: Image-guided decision support system for pathology. *Mach Vis Appl* 11(4):213–224, 1999
14. Kwak DM, et al: Content-based ultrasound image retrieval using a coarse to fine approach. *Ann N Y Acad Sci* 980:212–224, 2002
15. Lim J, Chevallet J-P: Vismed: A visual vocabulary approach for medical image indexing and retrieval. in *Second Asia Information Retrieval Symposium*. 2005. Jeju Island, Korea
16. Shyu CR, et al: ASSERT: a physician-in-the-loop content-based image retrieval system for HRCT image databases. *Comput Vis Image Underst* 75(1/2):111–132, 1999
17. Cauvin JM, et al: Computer-assisted diagnosis system in digestive endoscopy. *IEEE Trans Inf Technol Biomed* 7(4):256–262, 2003
18. Güld MO, et al: Content-Based Retrieval of Medical Images by Combining Global Features. *Accessing Multilingual Information Repositories*. in *Accessing Multilingual Information Repositories*. 2005: Springer LNCS 4022
19. Lubbers K, et al: A Probabilistic Approach to Medical Image Retrieval, in *Multilingual Information Access for Text, Speech and Images*. Springer Berlin, 2005, pp 761–772
20. Doyle S, et al: Using manifold learning for content-based image retrieval of prostate histopathology. in *MICCAI 2007 Workshop on Content-Based Image Retrieval for Biomedical Image Archives*. Brisbane, Australia, 2007
21. Gletsos M, et al: A computer-aided diagnostic system to characterize CT focal liver lesions: design and optimization of a neural network classifier. *IEEE Trans Inf Technol Biomed* 7(3):153–162, 2003
22. Zhu H, et al: A new local multiscale Fourier analysis for medical imaging. *Med Phys* 30:1134–1141, 2003
23. Oliveira MC, Cirne W, Marques PDA: Towards applying content-based image retrieval in the clinical routine. *Future Gener Comput Syst* 23(3):466–474, 2007
24. Rahman M, Bhattacharya P, Desai BC: A framework for medical image retrieval using machine learning and statistical

similarity matching techniques with relevance feedback. *IEEE Trans Inf Technol Biomed* 11(1):58–69, 2007

25. Mao J, Jain AK: Texture classification and segmentation using multiresolution simultaneous autoregressive models. *Pattern Recogn* 25(2):173–188, 1992

26. Iqbal Q, Aggarwal JK: Combining structure, color and texture for image retrieval: a performance evaluation. in *International Conference on Pattern Recognition (ICPR)*. Quebec City, Canada, 2002

27. Belongie S, Malik J, Puzicha J: Shape matching and object recognition using shape contexts. *IEEE Trans Pattern Anal Machine Intelligence* 24(4):509–522, 2002

28. Akgul CB, et al: 3D model retrieval using probability density-based shape descriptors. *IEEE Trans Pattern Anal Machine Intelligence* 31(6):1117–1133, 2009

29. Gokturk SB, et al: A statistical 3-D pattern processing method for computer-aided detection of polyps in CT colonography. *IEEE Trans Med Imag* 20(12):1251–1260, 2001

30. Yoshida H, Nappi J: Three-dimensional computer-aided diagnosis scheme for detection of colonic polyps. *IEEE Trans Med Imag* 20(12):1261–1274, 2001

31. Rubin GD, et al: Pulmonary nodules on multi-detector row CT scans: performance comparison of radiologists and computer-aided detection. *Radiology* 234(1):274–283, 2005

32. Alto H, Rangayyan RM, Desautels JEL: Content-based retrieval and analysis of mammographic masses. *Journal of Electronic Imaging* 14(2):1–17, 2005

33. Qian XN, Tagare HD: Optimal embedding for shape indexing in medical image databases. in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2005

34. Antani S, et al: Evaluation of shape similarity measurement methods for spine X-ray images. *J Vis Commun Image Represent* 15(3):285–302, 2004

35. Balmachnova E, et al: Content-based image retrieval by means of scale-space top-points and differential invariants. in *MICCAI 2007 Workshop on Content-based Image Retrieval for Biomedical Image Archives*. Brisbane, Australia, 2007

36. Golland P, et al: Detection and analysis of statistical differences in anatomical shape. *Med Image Anal* 9(1):69–86, 2005

37. Bansal R, et al: Statistical analyses of brain surfaces using Gaussian random fields on 2-D manifolds. *IEEE Trans Med Imaging* 26(1):46–57, 2007

38. Toews M, Arbel T: A statistical parts-based model of anatomical variability. *IEEE Trans Med Imag* 26(4):497–508, 2007

39. Wang JZ: Pathfinder: multiresolution region-based searching of pathology images using IRM. in *AMIA Symp*. 2000

40. Pokrajac D, et al: Applying spatial distribution analysis techniques to classification of 3D medical images. *Artif Intell Med* 33(3):261–280, 2005

41. Nielsen J, Nelson M, Liu L: Image-matching as a medical diagnostic support tool (DST) for brain diseases in children. *Comput Med Imaging Graph* 29(2/3):195–202, 2005

42. Sasso G, et al: A visual query-by-example image database for chest CT images: potential role as a decision and educational support tool for radiologists. *J Digit Imaging* 18(1):78–84, 2005

43. Petrakis EGM, Faloutsos C, Lin KI: ImageMap: an image indexing method based on spatial similarity. *IEEE Trans Knowl Data Eng* 14(5):979–987, 2002

44. Shantanu HJ, Washington M: Statistical shape analysis: clustering, learning, and testing. *IEEE Trans Pattern Anal Mach Intell* 27(4):590–602, 2005

45. Tong L, Hongbin Z: Riemannian manifold learning. *IEEE Trans Pattern Anal Mach Intell* 30(5):796–809, 2008

46. Zhang J, et al: Object representation and recognition in shape spaces. *Pattern Recogn* 36(5):1143–1154, 2003

47. Rubner Y, Tomasi C, Guibas LJ: The earth mover's distance as a metric for image retrieval. *Int J Comput Vis* 40(2):99–121, 2000

48. Bunke H, Irmiger C, Neuhaus M: *Graph Matching—Challenges and Potential Solutions*, in *Image Analysis and Processing—ICIAP 2005*, 2005, pp 1–10

49. Glaunès J, et al: Large deformation diffeomorphic metric curve mapping. *Int J Comput Vis* 80(3):317–336, 2008

50. Veltkamp RC: Shape matching: similarity measures and algorithms. in *Shape Modeling and Applications*, SMI 2001 International Conference on, 2001

51. Akgül CB, et al: Similarity learning for 3D object retrieval using relevance feedback and risk minimization. To appear in *Int. Journal of Computer Vision*, 2010

52. Rahmani R, et al: Localized content-based image retrieval. *IEEE Trans Pattern Anal Mach Intell* 30(11):1902–1912, 2008

53. Rui Y, et al: Relevance feedback: a power tool for interactive content-based image retrieval. *IEEE Trans Circuits Syst Video Technol* 8(5):644–655, 1998

54. Tao D, et al: Asymmetric bagging and random subspace for support vector machines-based relevance feedback in image retrieval. *IEEE Trans Pattern Anal Mach Intell* 28(7):1088–1099, 2006

55. Keysers D, et al: Statistical framework for model-based image retrieval in medical applications. *J Electron Imaging* 12(1):59–68, 2003

56. Müller H, et al: The Use of MedGIFT and EasyIR for ImageCLEF 2005. in *Accessing Multilingual Information Repositories*. 2005: Springer LNCS 4022

57. Mohammad-Reza S, et al: Content-based image database system for epilepsy. *Comput Methods Programs Biomed* 79(3):209–226, 2005

58. El-Naqa I, et al: A similarity learning approach to content-based image retrieval: application to digital mammography. *IEEE Transactions On Medical Imaging* 23(10):1233–1244, 2004

59. Aschkenasy VS, et al: Unsupervised image classification of medical ultrasound data by multiresolution elastic registration. *Ultrasound Med Biol* 32(7):1047–1054, 2006

60. 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 Trans Inf Technol Biomed* 10(3):598–607, 2006

61. Amores J, Radeva P: Registration and retrieval of highly elastic bodies using contextual information. *Pattern Recognit Lett* 26(11):1720–1731, 2005

62. Chin Y, et al: An automatic liver segmentation initialization information retrieval strategy for a CBIR followed by a new liver volume segmentation method for CT and MRI image datasets. in *MICCAI 2007 Workshop on Content-based Image Retrieval for Biomedical Image Archives*. 2007. Brisbane, Australia

63. Miller MI, Younes L: Group actions, homeomorphisms, and matching: a general framework. *Int J Comput Vis* 41(1):61–84, 2001
64. Carson C, et al: Blobworld: image segmentation using expectation-maximization and its application to image querying. *IEEE Trans Pattern Anal Mach Intell* 24(8):1026–1038, 2002
65. Korn P, et al: Fast and effective retrieval of medical tumor shapes. *IEEE Trans Knowl Data Eng* 10(6):889–904, 1998
66. Greenspan H, Pinhas AT: Medical image categorization and retrieval for PACS using the GMM-KL framework. *IEEE Trans Inf Technol Biomed* 11(2):190–202, 2007
67. Avni U, et al., X-ray image categorization and retrieval using patch-based visual words representation, in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI)*. 2009
68. Brown R, et al: The use of magnetic resonance imaging to noninvasively detect genetic signatures in oligodendroglioma. *Clin Cancer Res* 14(8):2357–2362, 2008
69. Badr Y, Chbeir R: Automatic Image Description Based on Textual Data, in *Journal on Data Semantics VII*. 2006. pp. 196–218
70. Besançon R, et al: Cross-Media Feedback Strategies: Merging Text and Image Information to Improve Image Retrieval, in *Multilingual Information Access for Text, Speech and Images*. 2005. pp. 709–717
71. Barb AS, Chi-Ren S, Sethi YP: Knowledge representation and sharing using visual semantic modeling for diagnostic medical image databases. *IEEE Trans Inf Technol Biomed* 9(4):538–553, 2005
72. Langlotz CP: RadLex: a new method for indexing online educational materials. *RadioGraphics* 26(6):1595–1597, 2006
73. Rubin DL, et al: Medical Imaging on the Semantic Web: Annotation and Image Markup. in *AAAI Spring Symposium Series, Semantic Scientific Knowledge Integration*. 2008. Stanford University
74. Syeda-Mahmood T, et al: AALIM: Multimodal Mining for Cardiac Decision Support. *Comput Cardiol* 1 and 2:209–212, 2007
75. Syeda-Mahmood T, Beymer D, Wang F: Shape-based matching of ECG recordings. *2007 Annual International Conference of the Ieee Engineering in Medicine and Biology Society, Vols 1–16:2012–2018*, 2007
76. Syeda-Mahmood T, et al: Characterizing spatio-temporal patterns for disease discrimination in cardiac echo videos. *Medical Image Computing and Computer-Assisted Intervention–MICCAI, Pt 1. Proceedings* 4791:261–269, 2007
77. Syeda-Mahmood T, Beymer D, Amir A: Disease-specific extraction of text from cardiac echo videos for decision support, in *Intl. Conf on Document Analysis and Recognition (ICDAR)*. 2009