
Similarity-Based Retrieval for Biomedical Applications

L. G. Shapiro¹, I. Atmosukarto¹, H. Cho², H. J. Lin¹, S. Ruiz-Correa³, and J. Yuen¹

¹ University of Washington, Seattle, WA 98195, U.S.A.

(shapiro@cs,indria@cs,hjlin@u,jenny@cs).washington.edu

² Svision, Inc., Bellevue, WA 98006, U.S.A. hansang@svisionllc.com

³ Children's National Medical Center, Washington DC 20010, U.S.A.
srcorrea@cnmc.org

1 Introduction

Similarity-based image retrieval is part of the case-based reasoning scenario. It allows for the retrieval of images from a database that are similar in some way to a given query image. It has been used in case-based reasoning systems for both image segmentation and image interpretation. Whereas case-based reasoning focuses on incremental learning of prototypical case-classes and index structures and on case mining, computer vision has been studying similarity measures, image indexing and image features.

Similarity-based retrieval has become an important area of computer vision research. It has reached a state of maturity in which it is starting to be used in real commercial and medical applications. In the commercial domain, similarity measures are being used to organize both personal and institutional databases and to retrieve visually similar images for advertising and marketing. In the medical domain, similarity-based retrieval is used by physicians who want to compare imaging studies from a new patient to those from a database of prior patients to help them determine the diagnosis and potential treatment options. Furthermore, research scientists in the biomedical domain are starting to use similarity-measures to organize and retrieve data from large-scale experiments involving multiple types of image and signal data. We have collaborated with three different biomedical research groups who needed similarity-based classification and retrieval systems to aid in their scientific research. In this paper, we discuss each of these applications, describe the retrieval systems we have developed for them, and suggest the need for a unified query formulation for a general multimedia information retrieval system for biomedical applications.

Section 2 provides an account of previous work on content-based retrieval systems. Section 3 describes a brain data retrieval system that retrieves neuronal signals and fMRI data from a database of patients who have had surgery for epileptic tumor resection. Section 4 describes a skull retrieval system that retrieves skull images from a database of CT images. Section 5 describes a mouse eye retrieval system that retrieves images in order to study the cataract development in the eyes. Section 6 describes our current work in developing a unified query framework for multimedia biomedical data retrieval. Finally, section 7 provides a summary of the chapter.

2 Related Work

Content-based image retrieval (CBIR) has been a heavily studied area of computer vision for the past 10 years including both general [48] [49] and medical [31] retrieval systems. The majority of these systems retrieve according to the ‘query by example’ paradigm; that is, the user provides or selects a *query image* and chooses a distance measure (or combination of such measures) that will be used to compare the query image to the images stored in the database. The system retrieves those database images that are judged ‘similar to’ the query image by the distance measure. Usually the retrieved images are returned in order of similarity to the query, and the user has the capability to browse through them in this order. Much of the research work in this area has related to the design of distance measures for accurate retrieval in various application domains. A smaller amount of work has considered the efficiency of retrieval systems, and some indexing methodologies have been developed [6] [46] [47] [9] [3] [8] as well as incremental learning methods for the index structure and the prototypical representation of case-classes [36] [17].

In recent years image databases have gone from theory to reality. There are commercial systems, such as IBM’s QBIC [14] and Virage’s image search engine [4] and research systems including PhotoBook [34], Chabot [32], ImageRover [44], VisualSEEk [50], WebSEEk [51], MARS [28], and our own FIDS [6]. The most common features used for retrieval in these systems are color histograms, texture measures, and simple shape measures. They do not recognize specific objects and are not intended for biological applications. Region-based systems such as Blobword [7], NETRA [24], and the Composite Region Template (CRT) System [52] perform segmentation based on color-texture properties to identify regions of interest, which can be used in spatial-relationship queries. Some systems employ relevance feedback where the system refines its concept of the user’s query according to the user’s feedback [29, 33, 40]. Both regions of interest and relevance feedback are useful in a biological multimedia database system.

There have been a number of systems developed for medical image retrieval [31] [30] [1] [37] [20] [56] [57] [46] [2] [13] as well as for case-based reasoning [25] [15] [18] [35]. Most of these systems were developed for a particular type of retrieval and did not attempt to work with multiple data types. Kelly and Cannon [19] used a global texture signature to characterize images of diseased lungs and a signature distance function to compare them. Chu and Cardenas [12] developed the KMeD (Knowledge-Based Multimedia Medical Distributed Database) System. This large and ambitious project allows querying of medical multimedia data by both image and alphanumeric content and allows for the modeling of both spatial and temporal content of medical objects. Liu *et al.* [23] have worked on retrieval of brain images with abnormalities, while Perner [35] worked on retrieval of CT brain images based on the image characteristics as well as on patient data for the automatic determination of degenerative brain diseases. Tagare *et al.* [55, 54, 38] have studied content-based medical image retrieval for a number of years. Their current research is on high-dimensional indexing in medical image databases applied to cervical and lumbar spine x-ray images, where retrievals are based on shape features. Tang *et al.* developed a system for histological image retrieval [56], while Zheng *et al.* designed a system for pathology image retrieval [57]. Shyu *et al.* [46] developed a structure called the statistical k-d tree to allow for rapid searches in content-based retrieval. More recently they have developed a real-time protein structure retrieval system with a multidimensional index [47][9].

3 Brain Data Retrieval for the Study of Language Sites in the Brain

During surgery for epileptic tumor resection, a technique called cortical stimulation mapping (CSM) is used to avoid areas on the cortical surface that are essential for language. The technique involves bringing the patient to an awakened state during surgery and presenting text, pictorial, or audio cues of a familiar object for the patient to name. During this process an electrical current is applied to selected *cortical surface sites*, marked by placing small numbered tags on the surface. The electrical current results in a short-term localized disruption of neural function. If this stimulation results in a naming error, then the cortical surface site at which it occurs is marked as essential for language function and is avoided during surgery.

In addition to their clinical use, the CSM language sites are a valuable source of data for understanding language organization in the brain, since their location is highly variable from one patient to the next. Studies have shown that the distribution of sites is correlated with such factors as sex and verbal IQ. The hope is that further insight into language organization can be obtained by correlating these sites, not only with additional demographic

data, but also with anatomical features and other measures of language such as functional magnetic resonance imaging (fMRI) and single unit recording (SUR) of individual neurons.

3.1 Experimental Procedure

Prior to surgery, each patient undergoes an fMRI scan while being shown the same stimuli as in the later surgical CSM studies (text, audio, pictures). Following each stimulus the scanner generates a time-series of 3-D volumetric data, where each voxel represents the time course of cerebral blood flow at that voxel. Subsequent statistical processing generates additional 3-D volumes, in which each voxel represents the probability that a statistically significant change in blood flow occurred between a stimulus and corresponding control. Single Unit Recording (SUR) studies are performed during surgery. As in the CSM studies the locations of the electrodes are marked by numbered tags associated with cortical surface sites. The same stimuli are presented as for fMRI and CSM, but in this case the electrodes record the firing patterns of individual neurons in response to the stimuli. Later processing looks for significant changes in firing rates between stimuli and controls, as well as changes in the firing patterns.

In order to correlate these diverse data sources for a patient, they are registered to a 3-D model of the patient's brain, using a structural MRI image volume acquired at the same time as the fMRI volumes. Population data are correlated by warping individual patient brain anatomy to a common "canonical" brain, in the process carrying along the registered functional data. All of the data are stored in several databases that are beginning to be used to help understand the relationships among various factors related to language.

3.2 Preprocessing and Feature Extraction

Both the signal data and fMRI data are preprocessed before they can be used in our retrieval system. The raw neuron spike data often contains spikes from two or more separate neurons. This signal data goes through a process called *spike sorting* in which the multiple neuron signals are separated. We have developed a template-extracting method to discriminate the neuron spikes using a rotated principal component analysis algorithm [10]. The raw fMRI data is also noisy. We have developed a method for dynamically detecting the activations (areas of high blood flow) using a maximal overlap wavelet transform to extract hemodynamic responses with minimum shape distortion and a dynamic time-warping algorithm to classify the different types of dynamic waveforms [11].

Figure 1 shows the multimedia data that must be stored and retrieved for this application. In addition to standard textual patient information there

is SUR data and fMRI data for each patient. The SUR data includes, for each stimulated electron, its ID, the stimuli shown to the patient, the firing pattern in the time and frequency domains, and several numeric statistics that we compute. The fMRI data also has stimuli and statistics and includes the 4-dimensional activation data (3D image over time). The activation level at each voxel over time is the important characteristic of this data.

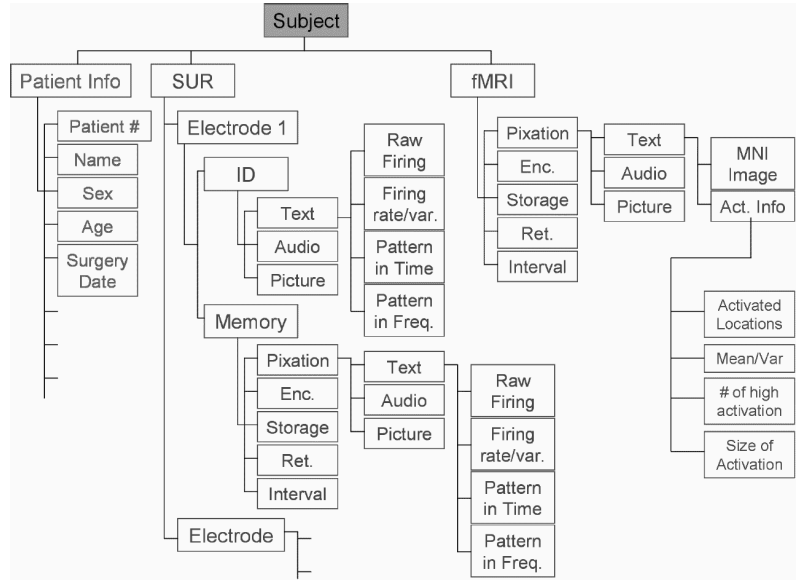


Fig. 1. Data structure

3.3 Similarity-Based Retrieval

Retrieval of signals is based on prior matching of the raw signal data to eight signal templates. The templates come from application of a rotated principal component analysis method [10]. They are shown in Figure 2 for the time domain only. A query signal is decomposed and the two best-matching templates are returned. From these, all signals in the database that match these two templates can be examined.

Figure 3 illustrates the results of a query for retrieving firing patterns. The top box shows the patient, microelectrode, and trial protocol information, and the second box gives the firing rate. The third box shows the best and second best template matches to the temporal firing pattern, and the fourth box shows the best and second best template matches to the frequency firing pattern. Users can extend the query by asking for similar results to those

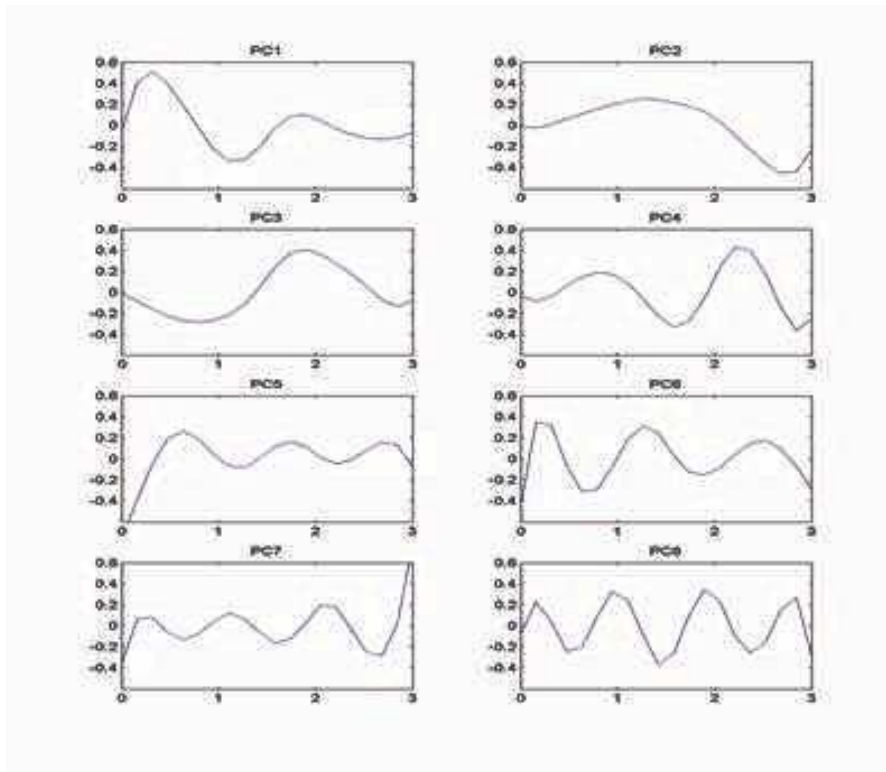


Fig. 2. Eight templates in the temporal domain.

retrieved by selecting one of the two options at the bottom of the screen: a) similar firing patterns from a single unit recording or 2) related fMRI results if they exist.

Figure 4 shows the results of a query to display both firing patterns from neurons and the related hemodynamic response from the fMRI image. The top box gives patient information, the closest two activated brain areas in the UW parcellation scheme, and the Brodman area in which the results lie, and the second box gives the average firing rate. The third box shows best template matches to the temporal and frequency firing patterns. The bottom box shows the hemodynamic response that occurred in the most highly activated voxel on the brain surface.

Figure 5 shows the results of an fMRI activation query. The top box shows the fMRI number, the MNI coordinates of the point with highest activation, and the closest two activated brain areas in the UW parcellaton scheme. The

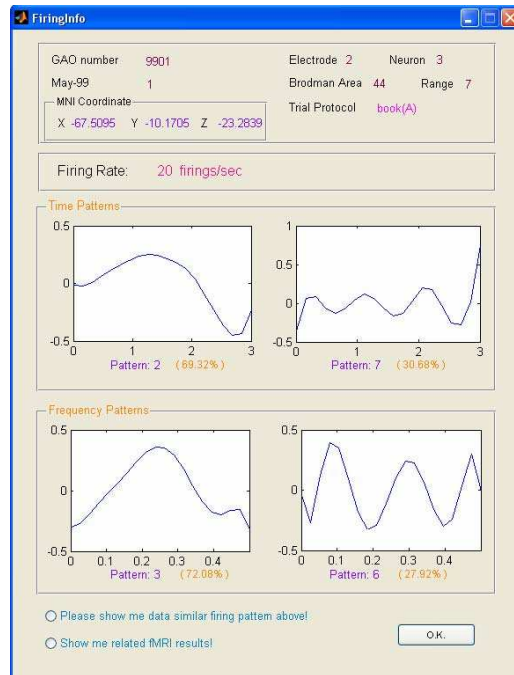


Fig. 3. Firing pattern retrieval

images show the highly activated areas in red.

The query system was custom-built for this application. There are three main routes that a query can take: 1) query by patient information, 2) query by trial protocols, and 3) query by firing patterns and/or fMRI activations. The third route is content-based, while the other two are text-based. The system allows the user to extend the basic query and move between the different query types as needed. A number of different queries have been used to illustrate the system, as summarized below.

1. query by patient information with SUR characteristics
2. query by patient information with fMRI characteristics
3. query by trial protocols of SUR data
4. query by trial protocols of fMRI data
5. query by common trial protocols of SUR and fMRI data
6. query by firing patterns of SUR data
7. query by fMRI activations
8. query by common features in both SUR and fMRI data

The prototype query system was developed for the scientists studying this particular application. It will become part of a more general system we are developing under National Science Foundation support.

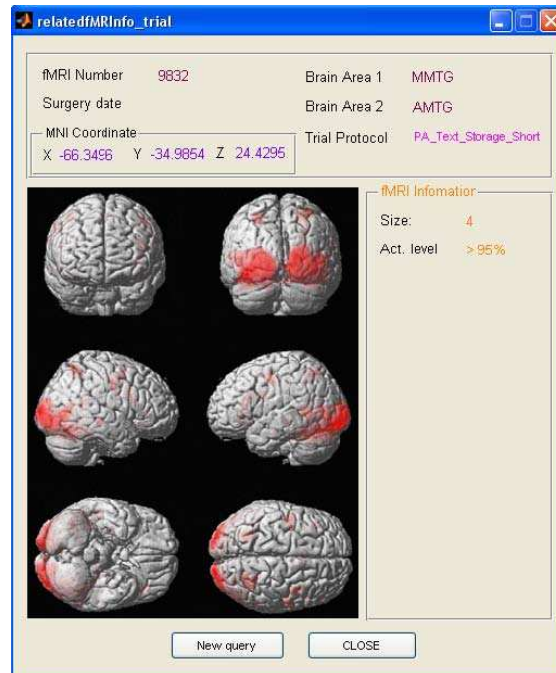


Fig. 4. FMRI retrieval

4 Skull Data Retrieval for Studying the Effects of Craniosynostosis

Researchers at the Pediatric Imaging Research Laboratory of the Children's Hospital and Regional Medical Center in Seattle study craniofacial disorders in children. In particular, they develop new computational techniques to represent, quantify, and analyze variants of biological morphology from imaging sources such as stereo cameras, CT scans and MRI scans. The focus of the research is the introduction of principled algorithms to reveal genotype-phenotype disease associations. Current projects include the development of quantitative descriptors of syndromic and non syndromic craniosynostotic head shape to a) investigate genotype-phenotype correlations, b) predict neurobehavioral and surgical outcomes, and c) develop a shape-based information retrieval system for craniofacial information.

Craniosynostosis is the pathological condition of premature fusion of one or more calvarial (skull) sutures or fibrous skull joints in childhood, affecting 1 in 2,500 individuals. Normally, an infant is born with open sutures, allowing for the development and expansion of the brain. However, in children with craniosynostosis, one or more of these sutures close prematurely. The early closure of these sutures results in abnormalities in calvarial shapes due to

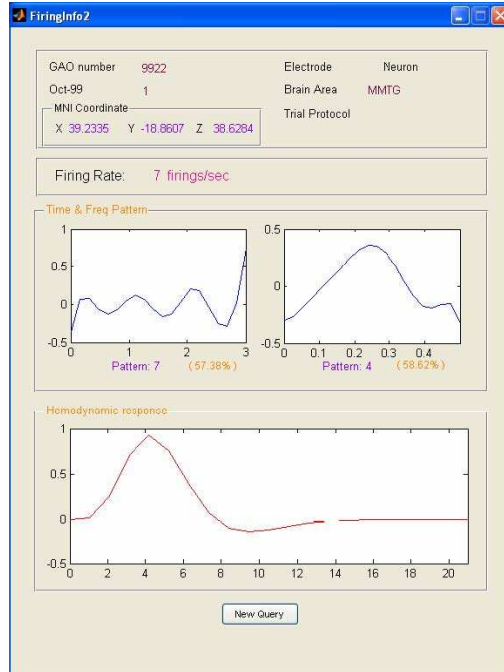


Fig. 5. Firing pattern and hemodynamic responses retrieval

the combination of restriction of osseous growth perpendicular to the fused suture and compensatory growth in unfused calvarial bone plates. Different types of craniosynostosis are classified according to the calvarial suture(s) involved with fusion detected in 3-D computed tomography (CT) images.

Isolated craniosynostosis, or single-suture synostosis, includes isolated fusion of the sagittal, metopic, and left or right unicoronal or lambdoid sutures (Fig. 6). The incidence of an isolated suture fusion is about 1 in 2,000 live births [45]. Sagittal synostosis is the most common form of isolated suture synostosis with an incidence of approximately 1 in 5,000, accounting for 40-60% of single-suture synostosis [21]. Early closure of the sagittal suture results in scaphocephaly, denoting a long narrow skull often associated with prominent ridges along the prematurely ossified sagittal suture (Fig. 6a). Unilateral coronal synostosis is the next most common sutural fusion, with incidence rates of about 1 in 11,000 [21]. It is manifest at birth as an asymmetrically skewed head with retrusion of the forehead and brow on the same side as the fused suture and with compensatory bulging of the forehead on the side opposite the fused suture (Fig. 6b). Metopic synostosis is less common and affects 1 in 15,000 individuals [21]. The premature fusion of the metopic suture produces trigonocephaly, denoting a triangular shaped head with prominent

frontal crest (Fig. 6c). The degree of skull shape deformity and changes in frontal and occipital bulging and biparietal narrowing can vary significantly between individuals even in the same disease class. In most cases of single-suture synostoses, a single surgery, i.e. cranioplasty, is required to release the fused suture and reshape the deformed calvaria. This surgery is preferentially performed within the first year to capitalize on the malleability of the infant’s skull and to minimize secondary facial deformation [27] [26].




Fig. 6. Lateral, frontal and top views of a) a patient with sagittal synostosis, b) a patient affected with unicoronal synostosis, and c) a patient affected with metopic synostosis.

Whether or not different calvarial shapes directly affect the severity of the cases, neuropsychological development, complications during cranioplasty, and post-surgical long-term outcomes for cases of single-suture synostoses remain largely unknown. In addition, the assessment of surgical “success” has been subjective, and no standardized method has been established for the post-surgery evaluation. However, surgeons and craniofacial experts often use cases of similar skull shapes from their past experience as guidelines in the preparation and evaluation of the reconstruction of the skull. This “case-based” study of reasoning makes it possible to reuse prototype images and diagnoses of previously resolved cases to assess the possible surgical complications and the post-surgery outcomes of the new cases. Therefore, the case-based clinical decision support technique produces a need to retrieve similar images of

shapes in patients with single-suture synostosis objectively and reproducibly. In addition, medical images, especially the diagnostic images, are produced in ever-increasing quantities. Searching through a large collection of digital images by keyword indexing, or simply by browsing, may be time-consuming and may not satisfy physicians' needs for retrieving cases of similar shapes. Therefore, it is critical to design a shape-based image retrieval system to not only ease the management of clinical data but to aid the radiologists and surgeons in the decision making analysis of the reconstruction of the skull.

In order to design an image-based clinical decision support system for surgeons and craniofacial experts, it is critical to first develop shape descriptors that will enable objective and reproducible detection and quantification of similarities and differences in the three-dimensional skull shapes. We have developed several different quantitative shape descriptors for our retrieval system.

4.1 Shape Descriptors

Our shape descriptors were computed from CT image slices obtained with skull imaging. To standardize our computations, we used a calibrated lateral view of a 3-D reconstruction of the skull to select three CT planar slices defined by internal brain landmarks (Fig. 7) These planes are parallel to the skull base plane, which is determined by connecting the frontal nasal suture anteriorly and the opisthion posteriorly. The A, F, and M planes we use are shown in Figure 7. The A-plane is at the top of the lateral ventricle, the F-plane is at the Foramina of Muntro, and the M-plane is at the level of the maximal dimension of the fourth ventricle. Oriented outlines were extracted from the CT image at each of these planes.




Fig. 7. All shape descriptors are constructed using the CT image slices extracted from skull imaging. These bone slides are defined by selecting the base plane on the 3-D image and finding internal anatomical landmarks on cerebral ventricles.

We have developed four different quantitative shape descriptors, three of which are numeric [43][42] [41] and one of which is symbolic [22]. The numeric

shape descriptors range from a single number per planar slice to a large matrix of numbers. The symbolic shape descriptor is a vector of probabilities obtained from a bag-of-words approach to shape description. The shape descriptors we have developed to characterize skull morphology in craniosynostosis are summarized below.

1. **Scaphocephaly Severity Index (SSI)**. SSI is similar to the cephalic index used as the current clinical standard to measure the variation in head shapes in patients with isolated craniosynostosis. However, instead of measuring the ratio of the head width to the length by the measurer’s subjective judgment, SSI takes the ratio, β/α , computed at an outline extracted from CT image slices from skull imaging (Fig. 8) [43]. The SSI assumes that a skull shape can be approximated as an ellipse that has eccentricity

$$e = \sqrt{1 - \frac{\beta^2}{\alpha^2}} \quad (1)$$

Note that if $e = 0$, $\beta/\alpha = 1$ and the outline shape would be a perfect circle. An eccentricity value close to 1 would suggest a very narrow outline shape. To increase the sensitivity of the shape description using SSI, users can combine the SSI measurements from multiple outlines extracted from skull images to represent the 3D data. Depending on the performance evaluation metric, either individual or combined SSI values can be used as shape descriptors. This measurement is very simple and efficient to compute, but it disregards a broad range of shape variations that may be of importance in capturing skull morphology.




Fig. 8. The scaphocephaly severity index is computed as the head width to length ratio β/α as measured on a CT plane extracted from skull shape imaging.

2. **Cranial Spectrum (CS)**. CS is a Fourier-based shape representation that describes a skull shape with the magnitude of the Fourier series coefficients of a periodic function [42]. The planes extracted from CT images are oriented outlines that have directions defined by their corresponding

tangent (T) vectors (Fig. 9b). An oriented outline can be represented as a periodic function by using polar coordinates with the center of mass as the origin of the coordinate system (Fig. 9c). This periodic function is then decomposed into a weighted sum of basis outlines by Fourier series. The coefficients from this decomposition then constitute the resulting shape descriptor. This representation encompasses shape information that cannot be captured by the SSI ratios. It is also closely related to traditional DFT-based descriptors [39]. The aim of Fourier analysis as applied here is to decompose an outline shape into a weighted sum of basis outlines, where each basis stratifies particular geometric features of a shape.




Fig. 9. a) Bone CT slice at the level of the A-plane; b) oriented outline counter clockwise direction; c) same outline represented in polar coordinates (ρ, θ) ; and d) 21 components of the corresponding cranial spectrum. Key: α (maximum outline length), T (tangent vector), N (normal vector), and (CM) center of mass.

3. **Cranial Image (CI).** The CI descriptor is a matrix representation of pairwise normalized square distances computed for all the vertices of an oriented outline that has been discretized into N evenly spaced vertices [41]. Let D be a symmetric matrix with elements $D_{ij} = d_{ij}/\alpha$, for $i, j = 1 \dots, N$, where d_{ij} is the Euclidean distance between vertices i and j , α is the maximum length of the contour (Fig 10), and N is an arbitrary number defined by the user. Since the outline is oriented, the vertices can be sequentially ordered up to the selection of the first vertex. As a consequence, the matrix D is defined up to a periodic shift along the main diagonal. The definition of CI can be extended to incorporate an arbitrary number of oriented outlines by computing inter and intra-oriented outline distances for each of the vertices of all of the outlines representing a skull.
4. **Symbolic Shape Descriptor (SSD)** The SSD algorithm was developed to model morphological variations that characterize different synostotic skull shapes[22]. It uses a *bag of words* (BOW) approach to capture the local shape variations. Even though this BOW approach has the advantage of producing a simple data representation, it creates a high-dimensional feature space and subsequently hinders efficient statistical analysis and image retrieval. In order to overcome these challenges, we use *probabilistic*




Fig. 10. a) Oriented contour represented as a sequence of N evenly spaced points. b) Cranial image. c) Top view and normalized distance scale.

latent semantic analysis (PLSA) to capture the co-occurrence information and relationship between elements in the bag of words in order to reduce the high computational complexity from the BOW representation.

The construction of the symbolic shape descriptors from skull imaging involves several steps, as illustrated in Figure 11. In brief, the cranial image distance matrices are computed for all skull outlines in the training set. K -means clustering of the rows of the distance matrices are applied and yields a set of clusters of the outline points. A symbolic cluster label is assigned to each outline vertex. A co-occurrence matrix is then computed for all training data using the strings of symbols constructed from the cluster labels of the vertices. Finally, PLSA is employed to reduce the dimension of the co-occurrence matrix and to construct the SSD.




Fig. 11. Construction of the SSD representation of an extracted skull outline. 1 = cranial image representation, 2 = k -means clustering, 3 = co-occurrence matrix construction, 4 = PLSA analysis

The input of the SSD algorithm is a set of skull shapes $\mathcal{S} = \{S_1, \dots, S_M\}$. Each skull shape is represented by L oriented outlines, and each outline is discretized into N evenly spaced vertices. For the sake of simplicity and without loss of generality, we assume that $L = 1$. The feature generation algorithm is as follows:

- a) For each shape S_j in \mathcal{S} and each vertex v_i of S_j , compute the vector of distances from all other vertices of S_j to v_i . This vector is the same

as the i -th row of the cranial image matrix descriptor (Fig. 10).

- b) Cluster all vectors in S by the k -means clustering algorithm with user-selected k and assign each cluster a symbolic label. Each vertex receives the label of its cluster.
- c) Compute a *bag of words (BOW)* representation of the skull outlines in \mathcal{S} . More specifically, the symbols associated with the vertices of an oriented outline are used to construct strings of symbols or *words*. The string size is fixed at some integer $1 \leq W \leq N$ and is specified by the user. For instance, when $W = 3$, each word contains a string of three symbols. A *BOW* representation for the outline in Figure 12a is the unordered set $s = \{ 'CAA', 'AAB', 'ABB', 'BBC', 'BCD', 'CDB', 'DBC', 'BCA' \}$. These strings represent the *local* geometric properties in the skull shapes.




Fig. 12. Symbolic labels are assigned to the vertices of the oriented outlines after applying k -means clustering to their numeric attributes. Local geometric aspects are represented by forming strings of these symbols. Oriented outlines of a) sagittal, b) metopic and c) normal head shapes, computed at the level of the A-plane.

- d) Compute a $M \times V$ co-occurrence matrix of counts $n(s_i, w_j)$, denoting the number of times the word w_j occurs in the *BOW* s_i associated with the skull outline S_i in the training sets.
- e) Apply *probabilistic latent semantic analysis (PLSA)* to the co-occurrence matrix \mathcal{S} [16]. PLSA is a latent variable model which associates an unobserved class variable $z_k \in z_1, \dots, z_P$ with each observation, an observation being the occurrence of a string of symbols from a particular *BOW* s_j . PLSA was originally applied to document retrieval for which a particular object is a document and the string of symbols is a word in that document.

- f) Use the class-conditional probabilities $P(s|z)$ estimated in the previous step to construct the symbolic shape descriptors for the outlines in \mathcal{S} . More specifically, for each outline S_i in \mathcal{S} , form its corresponding *symbolic shape descriptor* as the P -dimensional vector $[P(s_i|z_1), \dots, P(s_i|z_P)]$.

4.2 Shape-Based Retrieval

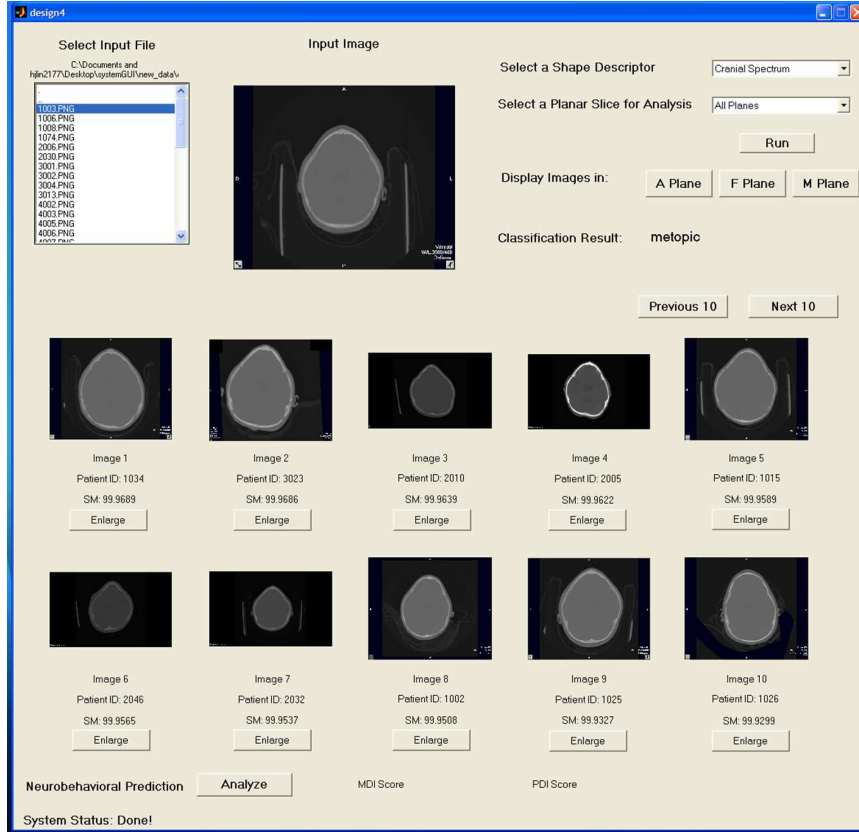


Fig. 13. An example of the system retrieving database images using the cranial spectrum representation on all 3 planar slices.

We have implemented a prototype system for shape-based image retrieval of skull CT imaging for craniofacial medicine. The system of database images supports retrieval based on their shape similarity to a query image. The shapes are represented by either numeric or symbolic features extracted using the four shape descriptors described above. A cosine similarity measure is

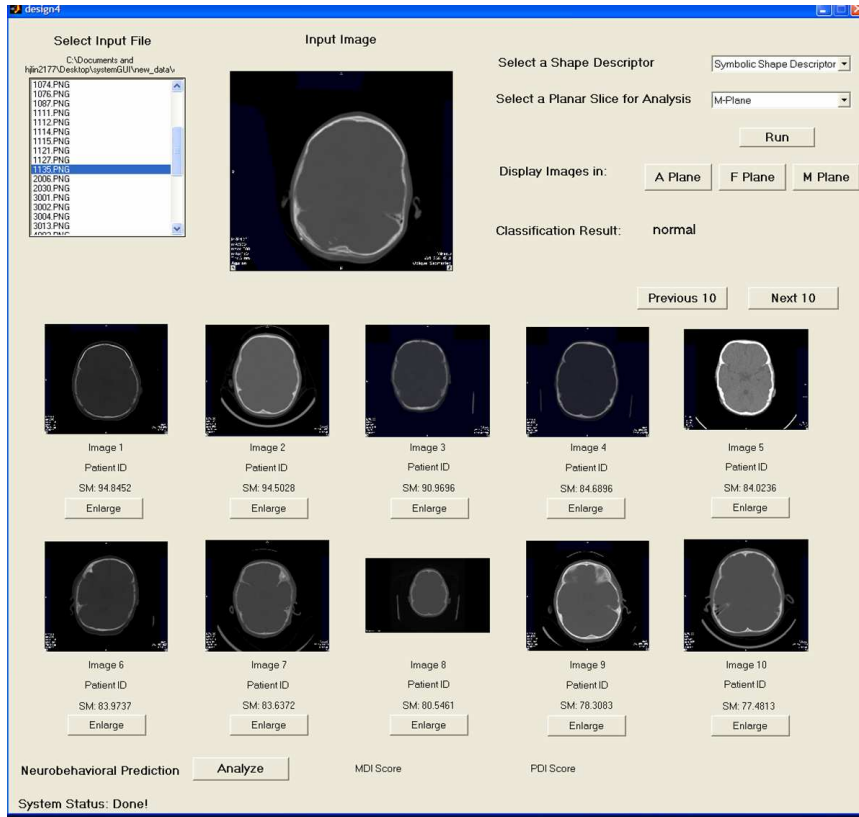


Fig. 14. An example of the system retrieving database images using the SSD representation on the M-plane.

incorporated for the retrieval of similar cases. The proposed system can help physicians and craniofacial experts with such tasks as diagnosis, intervention, and neurodevelopmental prediction.

The user is given a graphical user interface to the retrieval system. All the feature extraction and computation, feature classification, and feature vector comparison functions are done in the background. The queries to the image data can be specified using a query image to retrieve images that share similar features in the system. Figure 13 shows the system retrieving database images that are most similar to a query image using the cranial spectrum representation on all 3 planar slices. Figure 14 shows an example of the system retrieving database images that are most similar to a query image using the symbolic shape descriptor (SSD) on all 3 planar slices. The system is currently undergoing a rigorous evaluation that compares the results it retrieves to those selected by a craniofacial expert.

5 Mouse Eye Image Retrieval for the Study of Cataract Development

Researchers in the Eye Lab at the University of Washington are studying cataract formation in the eye, using mice as the subjects of their experiments. The mice are of several different strains with different genetic factors including a control group of normal mice. The eyes of the mice are photographed at regular intervals using a slit-lens technique. Thus each mouse has associated subject data and genotype that need to be correlated with the progression of opacification observed in a sequence of slit lamp images. A current study will relate genotype (the mutation in a particular strain of mouse) to phenotype (the presentation of the opacity with respect to pattern and intensity of light scattering) using these images. The experiment involves 8 different mutations in 2 different strains that are imaged approximately once per week. The number of images is enormous. Organizing the images through classification and clustering procedures as well as through the particular mice they came from is essential for the success of the project.

The set of cataract classes used in our work contains three classes: (1) WT, which stands for wild type and has no laboratory-induced cataract, (2) Secreted Protein Acidic and Rich in Cysteine (SPARC) knockout (a matricellular protein), and (3) the Synaptic Cleft (SC1) knockout. A knockout means the gene coding for a specific protein has been truncated or replaced so that the functional protein is no longer expressed. Figure 15 shows typical images for each class.

The original high-resolution images are resized to 300 by 300 pixel versions. Some of the properties of the images in the datasets are:

- The eyes in the images are approximately the same size and are approximately centered at the same location in each image.
- The illumination during the image capture process varies among the different experiments (Figure 16).
- There are artifacts caused by the illumination that are independent from the cataracts (Figure 17).
- The pattern in the center of the eye is directly related to the cataract in the lens.
- A ring pattern can be observed in the WT class. Depending on the cataract, partial or total occlusion of this ring pattern is observed.
- Because of the way the images are taken, the ring pattern is not circular but elliptical. This makes the task of detecting rings more difficult (Figure 16).
- The images contain, besides the eye, adjacent parts of the mouse such as eyelashes, which are not of interest.




Fig. 15. Classes of our study: Wild Type (WT), Synaptic Cleft 1 protein knockout (SC1) [53], and Secreted Acydic Rich Cysteine (SPARC) knockout [5]




Fig. 16. The shape of the eye in an image varies depending on both the mouse and the angle of incidence of the slit-lamp ray with the surface of the eye. In addition to the shape, the mean intensity of the image varies due to the illumination variations not associated with the type of cataract.

The occlusion or cloudiness of the lens caused by the presence of certain cataracts changes the perception of the cell layers in the lens. Therefore, the pattern of the lens rings such as the relative colors between two rings will change, making this an important feature in the characterization of the cataract class.

The angle of incidence of the slit-lamp light on the cornea is a reason for the elliptical shape of the ring pattern shown in the medical image. In order to encapsulate the largest amount of information, images with ring patterns as close to a circle as possible are preferred. However, not all the given images have this property, since they are selected manually and the image of the

mouse lens may only approximate a sphere. Fitting circular arcs to the rings in the images will not necessarily give a good approximation.

The numerous noise factors in the image make this problem non-trivial. Because the conditions in which the image is taken are not completely controlled, we cannot assume constant illumination or even constant shape (figure 17).



Fig. 17. Some artifacts caused by the camera are marked by yellow ellipses.

5.1 Feature Extraction

We have developed three different features that can discriminate between the different known classes. The features are (1) ring pattern, (2) intensity profile, and (3) histogram features. The features we have developed are summarized below.

Ring Pattern

Rings, or elliptical layers of cells, are present in both the normal and cataractous lens, unless the lens cells are disrupted. Visualization of the layers of lens cells depends on the magnification and contrast. In the normal lens, the contrast between adjacent cells or layers of cells is small (observation of the layers of cells in a normal lens requires high magnification in a microscope). The contrast increases when an increase in opacity occurs in some layers and not in adjacent layers. The increased contrast allows the rings to be observed at the magnification of the slit-lamp.

The pattern of rings can be correlated with the formation of cataracts. The proposed approach in identifying and quantifying these characteristics of the rings consists of a five-step process:

1. Ring enhancement.

2. Isolation of an elliptical sector of the lens.
3. Transformation of the elliptical sector into a rectangular image containing only pixels corresponding to the lens.
4. Compression of the rectangular image representation into a 1-dimensional array of mean intensities.
5. Extraction of feature vector values from the 1-dimensional array produced in the previous step.




Fig. 18. The original image (a), the original image after a histogram equalization (b), and the original image after a local equalization (c).

First the ring pattern is enhanced by a local histogram equalization transform as shown in Figure 18. Once the rings have been enhanced, the parameters of the ring pattern have to be extracted. Let the coordinates of the center of the eye be (c_x, c_y) . We will consider the region that corresponds to the elliptical sector centered at (c_x, c_y) and between the angles α and $-\alpha$, where $0 < \alpha < \frac{\pi}{2}$.

A ring is modeled as an elliptical arc with axes of length a and b parallel to the x and y axes respectively and centered at (c_x, c_y) . Every point (x, y) that lies in this elliptical arc is mapped to a point on a vertical line by the transformation $Rect(x, y) = (c_x + a, y)$ as shown in Figure 19.

The *Rect* transformation is applied to every ellipse centered at (c_x, c_y) with a fixed $c = \frac{a}{b}$ and $0 < a < w - c_x$, where w and h are the width and height of the image, respectively, to form a rectangular version of the elliptical-form eye. The resulting image of the *Rect* mapping is cropped to remove the cornea and the area outside the lens. The elliptical-to-rectangular transformation has three degrees of freedom: c_x , c_y , and c . The value of c_y is restricted to $\frac{h}{2}$, working with two degrees of freedom for each image.

The *Rect* transformation converts the elliptical rings into vertical lines for easier and more accurate analysis. For a given center (c_x, c_y) and length-width ratio c , it produces a vector $M(c_x, c)$ of the mean intensities for each of the

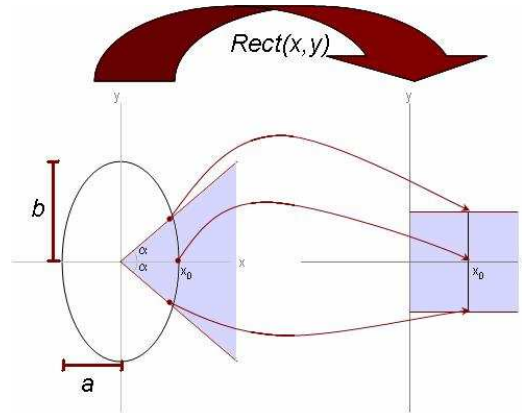


Fig. 19. The region inside an elliptical sector is mapped to a rectangular sector using the $Rect(x, y) = (c_x + a, y)$ transformation.

columns in $Rect(c_x, c_y)$ as shown in Figure 20.

The (c_x, c_y) and c that produce the best-fitting ellipse are used to generate the mean vector M that is analyzed to provide a feature vector of numeric attributes that can be used for pattern recognition. The vector we use contains the following ring attributes:

- number, maximum, minimum, mean, and variance of maxima of the function;
- number, maximum, minimum, mean, and variance of minima of the function;
- maximum, minimum, mean, and variance of the distance between pairs of consecutive maxima;
- maximum, minimum, mean, and variance of the distance between pairs of consecutive minima;
- maximum, minimum, mean, and variance of the difference between each pair of consecutive maxima and a minima;
- number of consecutive pairs of maxima and minima (regions formed by peaks and valleys).

These attributes capture several characteristics of the rings: globally, the attributes include the number of rings and the distribution of their properties such as mean intensity, width, and others; locally, the attributes describe the width and opacity of each ring and provide a comparison of these characteristics with the characteristics of other rings in the same image.




Fig. 20. An iterative process selects a center point, isolates an elliptical sector of the lens and maps it to a rectangular version. The variance at each column of the rectangular mapping is computed, and the case where the sum of the column variances is the lowest is selected as the best ellipse sector isolation; its center corresponds to the best center of the lens.

Intensity Profile

In order to capture the level of cloudiness of the lens, the intensity profile of the row of pixels in the middle of the image is also considered. Due to the different opacities in each layer of cells of the lens, a characteristic distribution of intensities exists for each class. With this information, a plot of intensity vs column (with the row being constant) is created. Call this function $I(x)$, where the domain is the columns in the image, and the range is $\{0, 255\}$ representing the possible values for intensity. One way of quantifying $I(x)$ is to fit a polynomial to it and use the values of its coefficients as features.

However, the large number of peaks and valleys creates too much noise to produce an accurate fit as shown in figure 21b. Instead, a Fast Fourier Transform $Fourier_I(x)$ as observed in figure 21c is applied. A polynomial of degree 5 is fit to this function, resulting in a complex function. The 12 coefficients of this polynomial (6 real and 6 imaginary) are concatenated to form a feature vector.

Histogram of Western Quadrant of the Lens

The difference in opacity, particularly in the western quadrant (as shown in figure 22) of the lens layers is characteristic for each cataract class. A histogram of the western quadrant of the lens is created and is fitted to a 1-dimensional Gaussian using Maximum Likelihood. The variance of the Gaussian is included in the feature vector. Because of the variations in illumination caused by external sources in the image, only the variance, which is related to the





Fig. 21. The white line in the left image(a) indicates the row of pixels that is considered for the intensity profile feature. The image in the middle (b) shows a polynomial fit of degree 5 on the intensity profile. The image on the right (c) shows a polynomial fit on the output of the Fast Fourier Transform of the intensity profile.

distribution of intensities, and not the mean, which is directly associated with the change in external illuminations, is considered. Figure 23 shows the histogram of the images corresponding to each of the three classes.



western quadrant

Fig. 22. Western quadrant of the lens.

Segmented Least Squares Fitting on the Intensity Profile

While the intensity profile feature can characterize some classes, the function is assumed to be continuous. Therefore, non-continuous fluctuations are approximated as smooth curves. Some classes such as SPARC and WT have similar polynomial fits but differ in the number of non-continuous changes. To encapsulate the amount of “continuity,” we fit a piecewise linear function on $I(x)$.

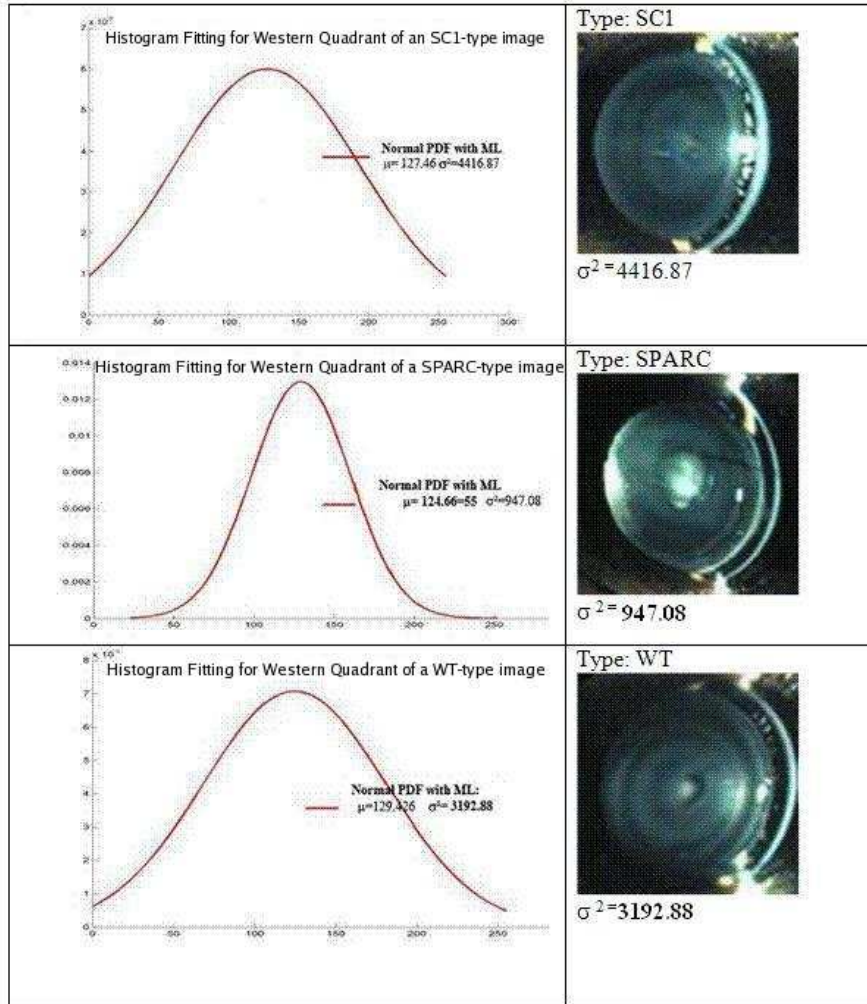


Fig. 23. The histogram of the western quadrant of the eye in three different classes. The variance of pixel intensities is smaller in SPARC type images and larger in SC1 type images, while the WT type images tend to have a variance value in between.

The standard least squares linear fitting algorithm fits a line to a set of n 2-dimensional points $P = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$. The line with the minimum error is $y = ax + b$, where

$$a = \frac{n \sum_i x_i y_i - (\sum_i x_i)(\sum_i y_i)}{n \sum_i x_i^2 - (\sum_i x_i)^2}$$

$$b = \frac{\sum_i y_i - a \sum_i x_i}{n}$$

The error of the fit is:

$$Error(L, P) = \sum_{i=1}^n (y - i - ax_i - b)^2$$

The objective of the segmented least squares approximation is to find the partition of consecutive points that minimizes the error of the fit, which corresponds to the sum of the errors for each linear segment used in the fit with an additional cost C for each segment used.

$$E_{fit} = \sum_{p_i \in P} Error(l_i, p_i) + C$$

where p_i is a set of consecutive points in $I(x)$, l_i is the least squares linear fit on this set, and C is the cost for each extra segment. We used values of C on the order of 500.

After fitting this piecewise linear function to $I(x)$, we only consider the left side of the eye for this feature extraction. This is done by considering the segments up to the middle of the image and discarding the first segments that have a very small value for their slope and intersection at the origin. These segments correspond to the dark section of the image, not a part of the mouse lens. The features extracted are:

- number of segments;
- means of the slope and the intersection at the origin (a, b) of each of the segments;
- variances of the slope and the intersection at the origin (a, b) of each of the segments.

5.2 Similarity-Based Retrieval

We have implemented a prototype system for eye image retrieval. The system retrieves eye images based on the similarity of their rings and intensities. The features are represented by a 44 dimension feature vector. A Euclidean distance measure is used to calculate the distance between each pair of images.

The user is presented with a graphical user interface of the retrieval system. Feature extraction and classification are done in the background. Figure 25 shows the system retrieving database images that are most similar to the query image. This system will allow the scientists doing the basic research to test their theories of genotype-phenotype relationships.




Fig. 24. The non-continuous changes in the intensity profile are not well encapsulated when fitting the intensity profile to a continuous function. The graph on the top is the intensity profile for the SPARC type lens on the top right; the bottom graph corresponds to the intensity profile for the WT type, shown on the bottom right. The non-continuous fluctuations of intensities in the image are observed as sharp edges in the intensity profile that are characteristic in SPARC images, while their presence is not as prominent in the WT class.

6 A Unified Query Framework for Multimedia Data

We wish to develop a unified query framework for multimedia biomedical data. The unified query framework must allow the specification of queries on alphanumeric data, text data, and multiple kinds of image and signal data. Our query framework will be an extended SQL that can handle probabilistic queries through *similarity measures*, which compare a query object to a database object of the same or comparable type and return a value between 0 (no match) and 1 (perfect match), which can be loosely interpreted as a

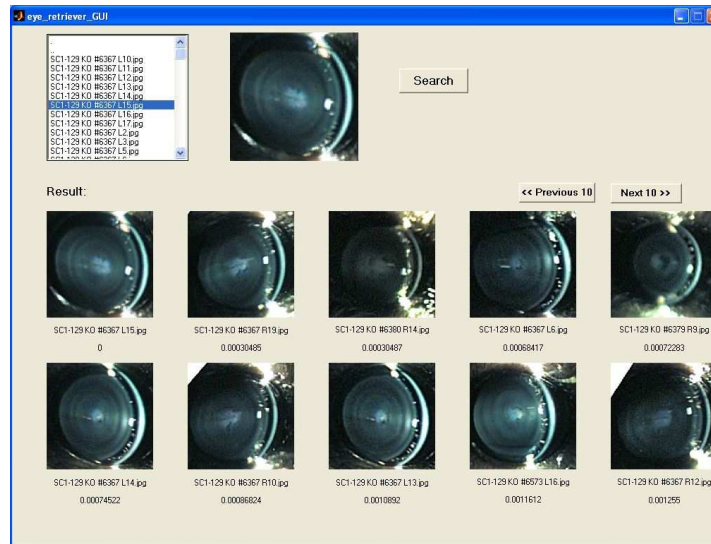


Fig. 25. An example of the system retrieving database images that are similar to the query image from class SPARC.

probability.

A query will have three integral parts, which are common to all database systems and can be expressed in SQL or other query languages: 1) the query specifications, 2) the result specifications, and 3) the matching requirements. The query specifications will include the objects that are provided to the query. These can include many different data types, such as text, numeric, date, matrix, table, signal, 2D image, 3D image, 3D mesh, and 3D image over time (video). The result specifications will include the objects that are to be returned by the query, which can include the same variety of data types. The matching requirements will specify constraints (predicates) on the returned objects, including both hard constraints that must be satisfied and soft constraints whose satisfaction is probabilistic in nature and involves the execution of similarity measures. For soft constraints, the data will be retrieved in ranked order, according to the probability of a match. The main focus of our work is on these similarity measures and probabilistic retrievals.

7 Summary

We have described three separate content-based retrieval systems for biomedical applications. The first retrieves neuronal signals and fMRI data from a database of patients who have undergone epileptic tumor resection. The sec-

ond retrieves skull images from a database of CT images from patients with both normal and abnormal craniofacial structure. The third retrieves slit-lens mouse eye images from a database of images from normal and DNA-modified mice. All three use very specialized image and signal features in their similarity measures. Our goal in our future work is to develop a set of feature extractors that can be used to construct useful similarity measures as part of a unified system for multimedia biomedical data retrieval.

References

1. A. M. Aisen, L. S. Broderick, H. Winer-Muram, C. E. Brodley, A. C. Kak, C. Pavlopoulou, J. Dy, C. R. Shyu, and A. Marchiori. Automated storage and retrieval of thin section CT images to assist diagnosis: System description and preliminary assessment. *Radiology*, 228:265–270, 2003.
2. S. Atnafu, R. Chbeir, and L. Brunie. Content-based and metadata retrieval in medical image database. In *Proceedings of the IEEE Symposium on Computer-Based Medical Systems*, pages 327–332, 2002.
3. Z. Aung and K.-L. Tan. Rapid 3d protein structure database searching using information retrieval techniques. *Bioinformatics*, 20(7):1045–1052, 2004.
4. J. R. Bach, C. Fuller, A. Gupta, A. Hampapur, B. Horowitz, R. Humphrey, R. Jain, and C. F. Shu. The virage search engine: An open framework for image management. In *Proceedings of SPIE Storage and Retrieval of Image and Video Databases*, 1996.
5. J. A. Bassuk, T. Birkebak, J. D. Rothmier, J. M. Clark, A. Bradshaw, P. J. Muchowski, C. C. Howe, J. I. Clark, and E. H. Sage. Disruption of the sparc locus in mice alters the differentiation of lenticular epithelial cells and leads to cataract formation. *Experimental Eye Research*, 68(3):321–331, 1999.
6. A. Berman and L. G. Shapiro. A flexible image database system for content-based retrieval. *Computer Vision and Image Understanding*, 75:175–195, 1999.
7. C. Carson, S. Belongie, H. Greenspan, and J. Malik. Region-based image querying. In *Proceedings of IEEE Workshop on Content-Based Access of Image and Video Libraries*, pages 42–49, June 1997.
8. K. Chakrabarti and S. Mehrotra. The hybrid tree: An index structure for high dimensional feature space. In *International Conference on Data Engineering*, pages 440–447, 1999.
9. P.-H. Chi, G. Scott, and C.-R. Shyu. A fast protein structure retrieval system using image-based distance matrices and multidimensional index. *International Journal of Software Engineering and Knowledge Engineering*, 15(4), 2005.
10. H. Cho, D. Corina, G. A. Ojemann, J. Schoenfield, L. Zamora, and L. Shapiro. A new neuron spike sorting method using maximal overlap discrete wavelet transform and rotated principal component analysis. In *Proc. IEEE Engineering in Medicine and Biology Society Annual International Conference*, volume 3, pages 2921–2924, 2003.
11. H. Cho, G.A. Ojemann, D. Corina, and L. G. Shapiro. Detection of neural activity in event-related fmri using wavelet and dynamic time warping. In *Proc. IEEE Applications of Digital Image Processing XXVI*, volume 5203, pages 638–647, 2003.

12. W. W. Chu, C. C. Hsu, A. F. Cardenas, and R. K. Taira. Knowledge-based image retrieval with spatial and temporal constructs. *IEEE Transactions on Knowledge and Data Engineering*, 10(6):872–888, 1998.
13. T. Deselaers, D. Keysers, and H. Ney. Fire - flexible image retrieval engine: Image clef 2004 evaluation. In *Proceedings of the CLEF 2004 Workshop*, pages 535–544, 2004.
14. M. Flickner, H. Sawhney, W. Niblack, J. Ashely, Q. Hyang, B. Dom, M. Gorkani, J. Hafner, D. Lee, D. Petkovic, D. Steele, and P. Yanker. The QBIC project: Querying images by content using color, texture, and shape. In *Proceedings of SPIE Storage and Retrieval of Image and Video Databases*, pages 173–181, 1993.
15. M. Haddad, K-P. Adlassnig, and G. Porenta. Feasibility analysis of a case-based reasoning system for automated detection of coronary heart disease from myocardial scintigrams. *Artificial Intelligence in Medicine*, 9:61–78, 1997.
16. T. Hofmann. Unsupervised learning by probabilistic latent semantic analysis. *Machine Learning*, 42:177–196, 2001.
17. S. Janichen and P. Perner. Conceptual clustering and case generalization of two-dimensional forms. *Computational Intelligence*, 22(3/4), 2006.
18. MC Jaulent, C. Le Bozec, E. Zapletal, and P. Degoulet. Case-based diagnosis in histopathology of breast tumors. *Medinfo*, 9(1), 1998.
19. P. Kelly, M. Cannon, and D. Bush. Query by image example: the CANDID approach. In *Proceedings of the SPIE Conference on Storage and Retrieval for Image and Video Databases III*, pages 238–248, 1995.
20. P. Korn, N. Sidiropoulos, C. Faloutsos, E. Siegal, and Z. Protopapas. Fast and effective retrieval of medical tumor shapes. *IEEE Transactions on Knowledge Data Engineering*, 10(6):889–904, 1998.
21. E. Lajeunie, M. Le Merrer, C. Marchac, and D. Renier. Genetic study of scaphocephaly. *American Journal of Medical Genetics*, 62(3):282–285, 1996.
22. H.J. Lin, S. Ruiz-Correa, L.G. Shapiro, A.V. Hing, M.L. Cunningham, M.L. Speltz, and R.W. Sze. Symbolic shape descriptors for classifying craniosynostosis deformations from skull imaging. In *Engineering in Medicine and Biology Society 2005. IEEE EMBS 2005. 27th Annual International Conference of the*, pages 6325–6331, 2005.
23. Y. Liu, N. Lazar, and W. Rothfus. Semantic-based biomedical image indexing and retrieval. In *International Conference on Diagnostic Imaging and Analysis*, 2002.
24. W. Y. Ma and B. S. Manjunath. NETRA: A toolbox for navigating large image databases. In *Proceedings of the IEEE International Conference on Image Processing*, pages 568–571, 1997.
25. R. Macura and K. Macura. Macrad: Radiology image resource with a case-based retrieval system. In M. Veloso and A. Aamodt, editors, *Case-Based Reasoning: Research and Development*, Springer, pages 43–54, 1995.
26. J.L. Marsh, A. Jenny, Galic M, Picker S, and M.W. Vannier. Surgical management of sagittal synostosis. a quantitative evaluation of two techniques. *Neurosurgery Clinics of North America*, 2(3):629–640, 1991.
27. J.L. Marsh and M.W. Vannier. Cranial base changes following surgical treatment of craniosynostosis. *The Cleft Palate Journal*, 23:9–19, 1986.
28. S. Mehrotra, Y. Rui, M. Ortega, and T. S. Huang. Supporting content-based queries over images in MARS. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*, 1997.

29. T. P. Minka and R. W. Picard. Interactive learning with a society of models. In *Proceedings of CVPR-96*, pages 447–452, 1996.
30. A. Mojsilovic and J. Gomes. Semantic based image categorization, browsing and retrieval in medical image databases. In *IEEE. International Conference on Image Processing*, 2000.
31. H. Muller, N. Michous, D. Bandon, and A. Geissbuhler. A review of content-based image retrieval systems in medical applications—clinical benefits and future directions. *International Journal of Medical Informatics*, 73:1–23, 2004.
32. V. E. Ogle and M. Stonebraker. Chabot: Retrieval from a relational database of images. *IEEE Computer*, 28:40–48, 1995.
33. J. Peng, B. Bhanu, and S. Qing. Probabilistic feature relevance learning for content-based image retrieval. *Computer Vision and Image Understanding: Special Issue on Content-Based Access of Image and Video Libraries*, 75:150–164, 1999.
34. A. Pentland, R. W. Picard, and S. Sclaroff. Photobook: Content-based manipulation of image databases. In *Proceedings of SPIE Storage and Retrieval of Image and Video Databases II*, pages 34–47, 1994.
35. P. Perner. An architecture for a cbr image segmentation system. *Journal on Engineering Application in Artificial Intelligence*, 12(6):749–759, 1999.
36. P. Perner. Case-base maintenance by conceptual clustering of graphs. *Engineering Applications of Artificial Intelligence*, 19(4):381–295, 2006.
37. E. G. M. Petrakis. Content-based retrieval of medical images. *International Journal of Computing Research*, 11(2):171–182, 2002.
38. X. Qian and H. D. Tagare. Optimally adapted indexing trees for medical image databases. In *IEEE International Symposium on Biomedical Imaging*, pages 321–324, 2002.
39. C. Rao. Geometry of circular vectors and pattern recognition of shape of a boundary. *Proc. Nat. Acad. Sci.*, 95:12783, 2002.
40. Y. Rui, T. S. Huang, M. Ortega, and S. Mehrotra. Relevane feedback: A power tool for interactive content-based image retrieval. *IEEE Transactions on Circuits and Systems for Video Technology: Special Issue on Segmentation, Description, and Retrieval of Video Content*, 8(5):644–655, 1998.
41. S. Ruiz-Correa, R.W. Sze, H.J. Lin, L.G. Shapiro, M.L. Speltz, and M.L. Cunningham. Classifying craniosynostosis deformations by skull shape imaging. In *Proceedings of the 18th IEEE International Symposium on Computer-Based Medical Systems*, pages 335–340, 2005.
42. S. Ruiz-Correa, R.W. Sze, J.R. Starr, A.V. Hing, H.J. Lin, and M.L. Cunningham. A fourier-based approach for quantifying sagittal synostosis head shape. In *American Cleft Palate Craniofacial Association*, 2005.
43. S. Ruiz-Correa, R.W. Sze, J.R. Starr, H.T. Lin, M.L. Speltz, M.L. Cunningham, and A.V. Hing. New scaphocephaly severity indices of sagittal craniosynostosis: A comparative study with cranial index quantifications. *Cleft Palate Craniofacial Journal*, 43(2):211–221, 2006.
44. S. Sclaroff, L. Taycher, and M. L. Cascia. Imagerover: A content-based image browser for the world wide web. In *Proceedings of the IEEE Workshop on Content- Based Access of Image and Video Libraries*, pages 2–9, June 1997.
45. A. Shuper, P. Merlob, M. Grunebaum, and S.H. Reisner. The incidence of isolated craniosynostosis in the newborn infants. *American Journal of Diseases of Children*, 139(1):85–86, 1985.

46. C.-R. Shyu, C. E. Brodley, A. C. Kak, A. Kosaka, A. M. Aisen, and L. S. Broderick. Assert: A physician-in-the-loop content-based image retrieval system for hrct image databases. In *Computer Vision and Image Understanding {Special Issue on Content-Based Retrieval from Image Databases}*, pages 111–131, 1999.
47. C.-R. Shyu, P.-H. Chi, G. Scott, and D. Xu. Proteindbs: A real-time retrieval system for protein structure comparison. *Nucleic Acids Research*, 32:W572–W575, 2004.
48. A. W. M. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain. Content-based image retrieval at the end of the early years. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(12):1349–1380, 2000.
49. A.W. M. Smeulders, T. S. Huang, and T. Gevers. Special issue on content-based image retrieval. *International Journal of Computer Vision*, 56:5–6, 2000.
50. J. R. Smith and S. F. Chang. Visually searching the web for content. *IEEE Multimedia Magazine*, 4(3):12–20, 1997.
51. J. R. Smith and S. F. Chang. Visually searching the web for content. *IEEE Multimedia Magazine*, 4(3):12–20, 1997.
52. J. R. Smith and C. S. Li. Image classification and querying using composite region templates. *Computer Vision and Image Understanding: Special Issue on Content-Based Access of Image and Video Libraries*, 75(1–2):165–174, 1999.
53. M. M. Sullivan and E. H. Sage. Hevin/sc1, a matricellular glycoprotein and potential tumor suppressor of the sparc/bm-40/osteonectin family. *International Journal of Biochemistry and Cell Biology*, 36(6):482–490, 2004.
54. H. D. Tagare. Increasing retrieval efficiency by index tree adaptation. In *IEEE Workshop on Content-Based Access of Image and Video Libraries*, pages 28–35, 1997.
55. H. D. Tagare, C. C. Jaffe, and J. Duncan. Medical image databases: a content-based retrieval approach. *Journal of the American Medical Informatics Association*, 4(3):184–198, 1997.
56. H. Tang, R. Hanka, and H. Ip. Histological image retrieval based on semantic content. *IEEE Transaction on Information Technology in Biomedicine*, 7(1):26–36, 2003.
57. L. Zheng, A. W. Wetzel, J. Gilbertson, and M. Becich. Design and analysis of a content based pathology image retrieval system. *IEEE Transaction on Information Technology in Biomedicine*, 7(4):249–255, 2003.