Content-based 3D Neuroradiologic Image Indexing and Retrieval: Preliminary Results

Yanxi Liu, William E. Rothfus (M.D.)[a], Takeo Kanade

CMU-RI-TR-98-04

[a]Associate Professor of Radiologic Sciences, Diagnostic Radiology, Allegheny University of the Health Sciences (Allegheny General Hospital)

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

This research is sponsored by the Allegheny-Singer Research Institute under prime contract through the Advanced Technology Program of the National Institute of Standards and Technology (NIST#70NANB5H1183). Views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of the National Institute of Standards and Technology or the United States Government.

This paper appears in the proceedings of 1998 IEEE International Workshop on Content-Based Access of Image and Video Databases in conjuncation with ICCV'98, Bombay, India, January, 1998.

ABSTRACT

A content-based 3D neuroradiologic image retrieval system is being developed at the Robotics Institute of CMU. The special characteristics of this system include: 1) directly dealing with multimodal 3D images (MR/CT); 2) image similarity based on anatomical structures of the human brain; 3) combining both visual and collateral information for indexing and retrieval. A testbed has been implemented for using detected salient visual features for indexing and retrieving 3D images.

1 Introduction

On-line image data is expanding rapidly in quantity, content and dimension. Existing "content-based" image retrieval systems, for example [7, 9, 23, 25, 22, 8], depend on general visual properties such as color and texture to classify diverse, two-dimensional (2D) images, thus are not truly content-based nor suitable for handling three-dimensional (3D) volumetric images. Furthermore, the difference of image sets taken within a single domain where the images all appear similar, differing only through subtle, domain-specific cues, cannot be captured using global statistical color and texture measures. It is the general consensus of the content-based image retrieval community[1] that future research should be more focused on domain specific images instead of general images due to the lack of evaluation standards.

The two main **objectives** in this research are to

1. develop a set of novel 3D image indexing and retrieval techniques, and
2. demonstrate a true content-based image retrieval system within a specific image domain, with a built-in evaluation scheme.

The most prevalent source of volumetric image data is medical and biological imaging. Other applications include the non-destructive diagnosis of malfunctioning mechanical objects such as engines, and quality inspection of manufactured objects (to see interior defects). Medical images form an essential and inseparable component through out the diagnosis and treatment process, which is especially true in neurology. Likewise, neuroradiologic images become a natural candidate for front-end indexing to retrieve relevant medical cases in a multimedia medical case database[2]. Using normal and pathological (bleed, stroke, tumor) neuroradiologic (MR/CT) images as an application domain, we propose a framework and a set of methodologies for content-based, 3D volumetric grey-level image retrieval with an emphasis on image indexing and evaluation of retrieved results.

A 3D neuroradiology image can be expressed as a stack of 2D images. An *ideal head coordinate system* can be centered in the brain with positive X_0, Y_0 and Z_0 axes pointing in the right, anterior and superior directions respectively (Figure 1, white coordinate axes). Ideally, a set of *axial (coronal,sagittal)* 2D slices is cut perpendicular to the $Z_0(Y_0, X_0)$ axis. Figure 2 shows a set of axial CT images placed on a plane and Figure 3 shows the same set of CT images interpolated in 3D (thresholded). In clinical practice, due to various positioning errors, we are presented not with the ideal coordinate system, but rather a *working coordinate system* XYZ in which X and Y are oriented along the rows and columns of each image slice, and Z is the actual axis of the scan

[1] IEEE Content-based video and image retrieval workshop associated with CVPR97, June, 1997

[2] Such a database "National Medical Practice Knowledge Bank" is currently under construction, which is funded by NIST

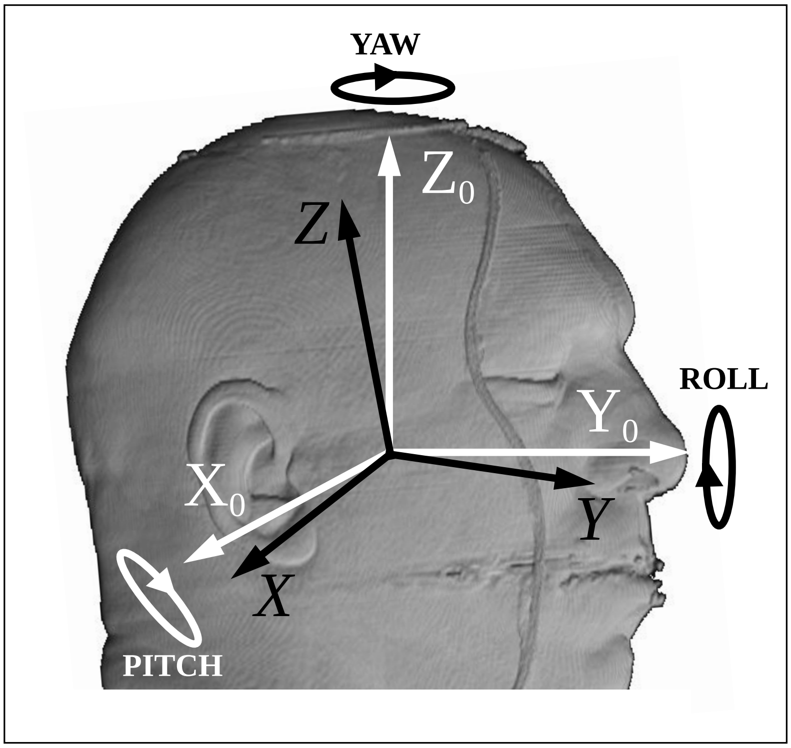

Figure 1: Ideal head coordinate system $X_0Y_0Z_0$ vs. the working coordinate system XYZ. Rendered head courtesy of the Visible Human Project.

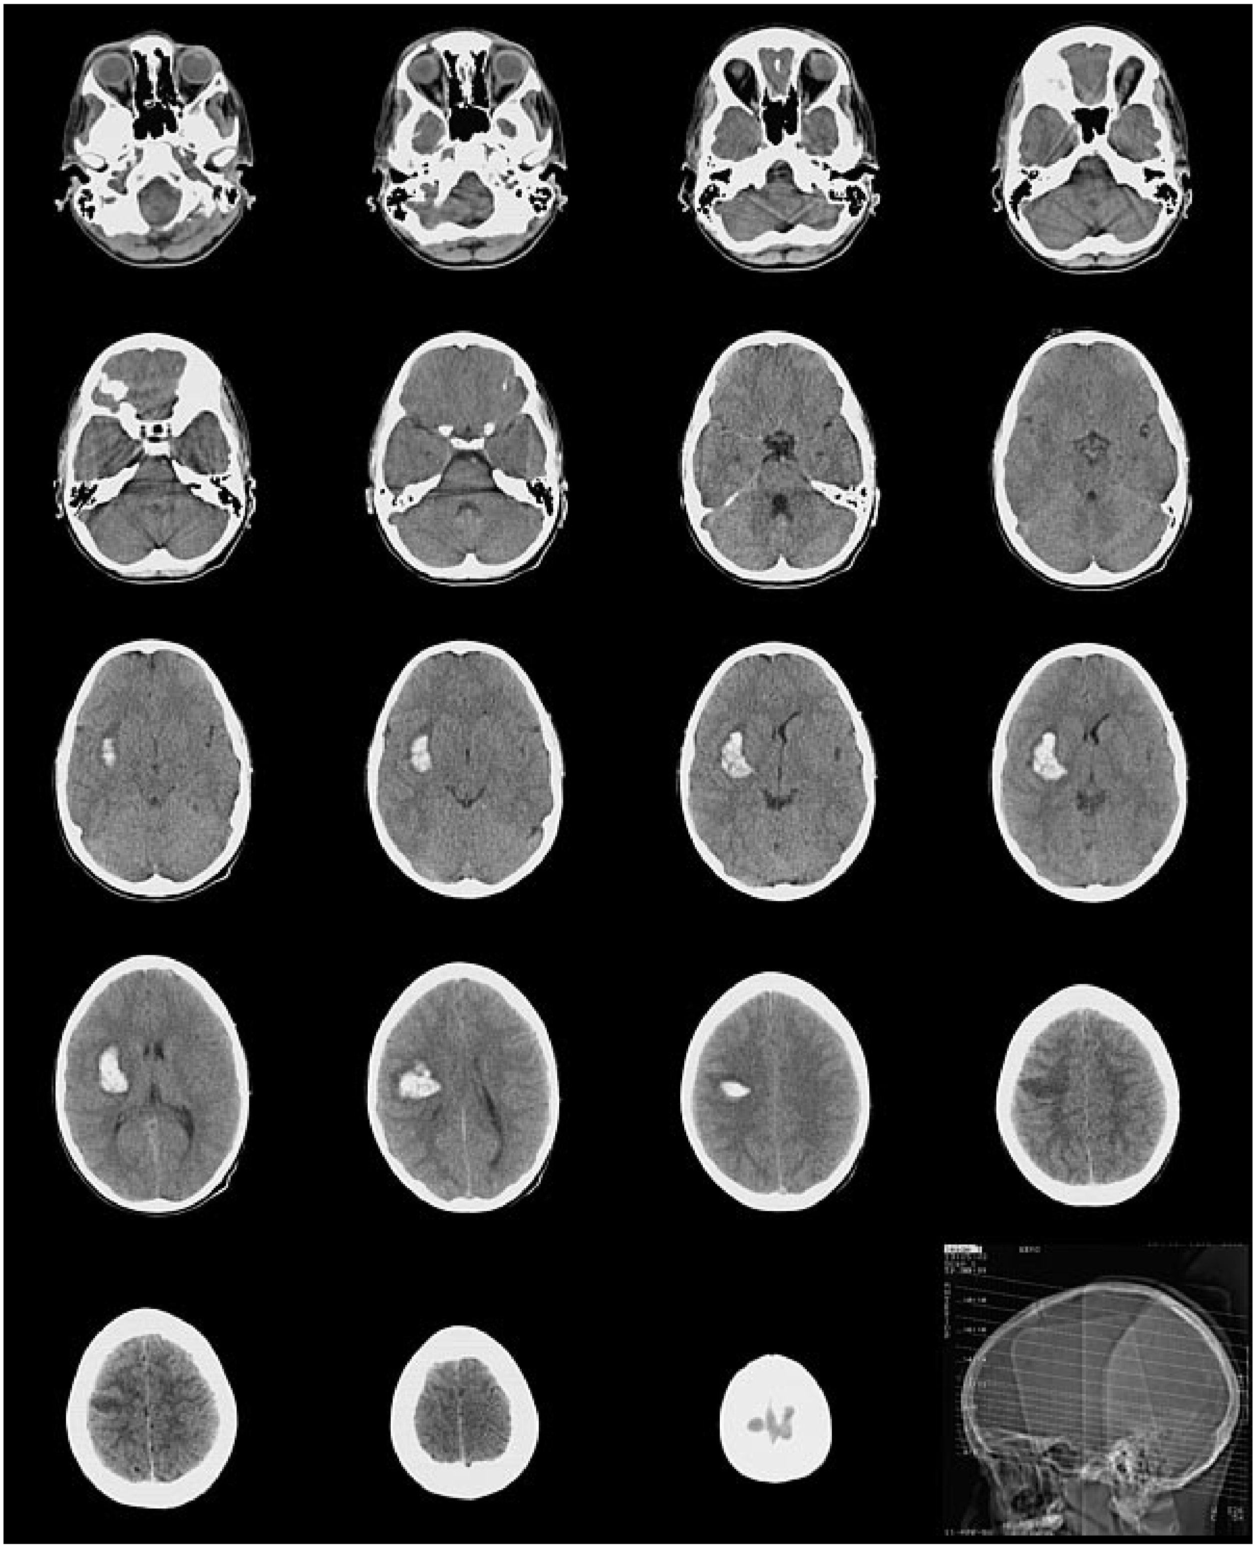

Figure 2: A 3D image which is composed of a set of clinical CT scans (axial), only a portion of a patient's head is captured as shown in a side view on the lower right corner. This is a case of acute right basal ganglion bleed.

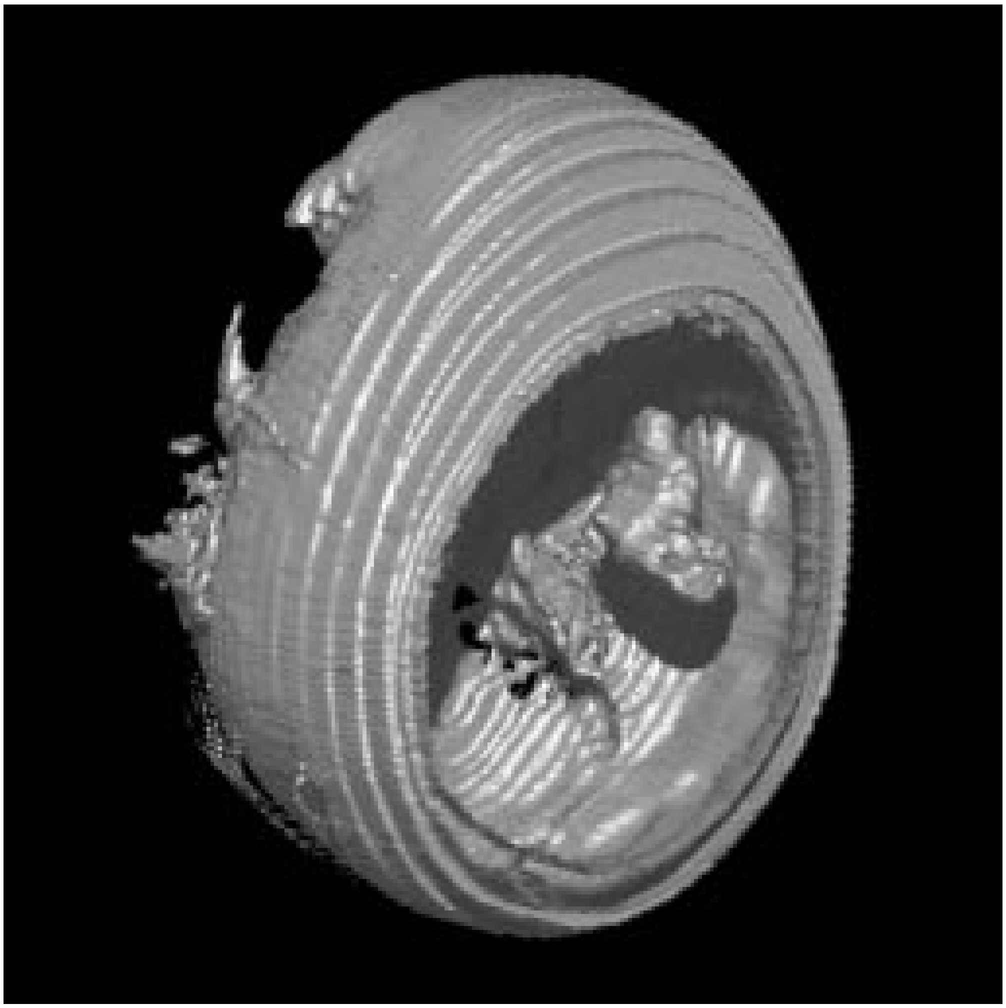

Figure 3: The 3D visualization of the lesion (shaded image) with respect to the skull

(Figure 1, black coordinate axes). The orientation of the working coordinate system differs from the ideal coordinate system by three rotation angles, *pitch*, *roll* and *yaw*, about the X_0, Y_0 and Z_0 axes, respectively. Several challenging issues that are unique in 3D medical image retrieval:

3D v. 2D images: Since the images are collected from different scanners in different hospitals, each set of images may start and end at different portions of the brain (not necessarily the whole brain is scanned, especially for CT images), and may be scanned in different angles or along different axes. Until two 3D images are properly registered and segmented, existing techniques for content-based retrieval using color, texture, shape and position can not be applied directly for meaningful results. This reveals one of the fundamental differences in 2D and 3D image retrieval: matching before comparison is necessary in 3D.

Rigid v. Deformable Registration: Since we are dealing with the similarity of 3D anatomical structures of human brains, which vary from person to person, a system enabling interpretation of deformable brain anatomy is required. 3D image intra-subject multimodality (same person's head, different modality) rigid registration has achieved excellent results [27, 28, 29],but rigid registration is not sufficient for matching two 3D images of different human brains (inter-subject). Normal brain deformable registration across patients [1, 4, 21, 26, 19], on the other hand, does not address problems in pathological clinical brain images. The topology of normal brains remains the same, while the shape may vary from person to person. Pathological brains violate this basic principle. This is one of the unsolved problems in medical image understanding.

Visual Content v. Conceptual Content: The ultimate goal for content-based medical image retrieval is to find *medically meaningful similar* cases. These cases may or may not present medical images that are similar in the usual visual sense. Similar visual features on images may not imply similar diagnosis or symptoms, and vice versa. Thus sometimes it is necessary to combine visual features in the image with background knowledge such as the anatomical location of the lesion or the symptoms of the patient to reach the right subset of relevant cases in the database.

Our retrieval task is targeted at a multimedia database (medical knowledge bank) composed of medical cases. Each case in the database contains at least one 3D image plus collateral information, including: the patient's age, sex, symptom, test results, diagnosis, operations, treatments, and outcomes in text, speech, or video formats. The idea is to use 3D images as the front-end key index to access the database and the goal is to retrieve a set of relevant medical cases from the database. The input to the retrieval system is a 3D image with a set of user determined weights on various visual features. Table 1 shows the parameters of a few sample 3D images in the database. The output is a set of retrieved cases ranked by their similarity.

From the application point of view, the result of this research will enhance the quality of medical specialist consultations and medical education. The usefulness of such an image retrieval system

Table 1: A Sample of 3D Image Data

Set	Modality	#Slices	Form and Size	Voxel (mm^3)	Pathology
1	CT	18	axial 664x534	0.5x0.5x2 (1-10) 0.5x0.5x10 (11-18)	Right occipital/parietal acute bleed
2	CT	19	axial 524x518	0.5x0.5x5 (1-11) 0.5x0.5x10 (12-19)	Basal ganglion acute bleed
3	CT enhanced	33	axial 686x550	0.5x0.5x5 (1-15) 0.5x0.5x10 (16 - 33)	right parietal/occipital meningioma
4	MR	187	axial 176x236	0.98x0.98x1.2	Normal
5	MR (T1)	123	coronal 256x256	0.9375x0.9375x1.5	Atlas, Normal
6	CT	16	axial 686x550	0.5x0.5x10	Basal ganglion acute bleed
7	CT	9	axial 686x550	0.5x0.5x10	Right thalamic acute bleed
8	CT	17	axial 678x542	0.5x0.5x5 (1-9) 0.5x0.5x10 (10-17)	Frontal astrocytoma high grade glial

can be observed in the following scenarios: (1) A general practitioner medical doctor may confirm her diagnosis of a specific patient and explore possible treatment plans through a consultation to the medical knowledge bank over the web; (2) A medical student may have a set of images and would like to explore possible diagnoses. In both cases, the system will return a set of content relevant images ranked according to their similarity scores to the target image. Each retrieved image is associated with its collateral information including a diagnosis which may or may not be shared by the rest of the retrieved images. Based on the diversity of the diagnosis in this set, either a consensus can be reached so that the doctor can confirm her diagnosis and follow the treatment plans accordingly or further information is needed to confirm a diagnosis. Through this process the medical student will learn about the subtlety of differential diagnosis in radiology by varying weights on different visual features.

Although this exploratory research is carried out using medical images, several important general research issues in computer vision and database retrieval are addressed, including:

- exploring multimodality 3D deformable registration for images of varied topology,
- developing similarity analysis of irregular 3D shapes,
- incorporating both image and text data in indexing to achieve improved retrieval performance,
- discovering a good balance between automated and interactive operations during the process of image retrieval, and
- establishing a quantitative evaluation scheme for the retrieval system.

More generally, this research is developing a prototype interactive 3D image retrieval system for exploring the potential of combining cutting edge image understanding techniques, machine learning, and neurology domain knowledge. This work extends image retrieval research into 3D images, and initiates a new generation of true content-based (semantic) image retrieval systems.

2 Related Work

In the medical imaging domain, most of the work on automatic 3D image retrieval has been dealing with dense (typically with $256 \times 256 \times 125$ voxels each with size $1 \times 1 \times 1.5mm^3$) MR images on normal subjects. The retrieval process is turned into a deformable registration process regardless of feature-based [5] or density-based [11]. In [11] a database of 10 MR images of healthy subjects are used to retrieve corresponding anatomical structure defined as "volume of interest" (VOI) given by a user in a reference image. The basic approach is a density-based registration algorithm using global to local affine followed by free-form transformations for the best match between the VOI and images in the database. Correlation is used as a measure for morphological differences. The CANDID project [15] carried out at Los Alamos National Lab deals with retrieving similar CT slices with lung diseases using texture analysis.

3 Image Indexing

The inherent categorization of pathologies in medicine provides a way to classify medical images (cases). This content-based organization of medical images naturally forms a hierarchical structure with its bottom leaves corresponding to a set of specific images (cases), and higher nodes correspond to a subcategory of pathological cases. Figure 4 is an example of how relevant cases are grouped in the database. Using this structure, cases are first sorted by their pathological nature (tumor,bleed,stroke, etc.), and then sorted by their proximate anatomical locations (intra-axial,

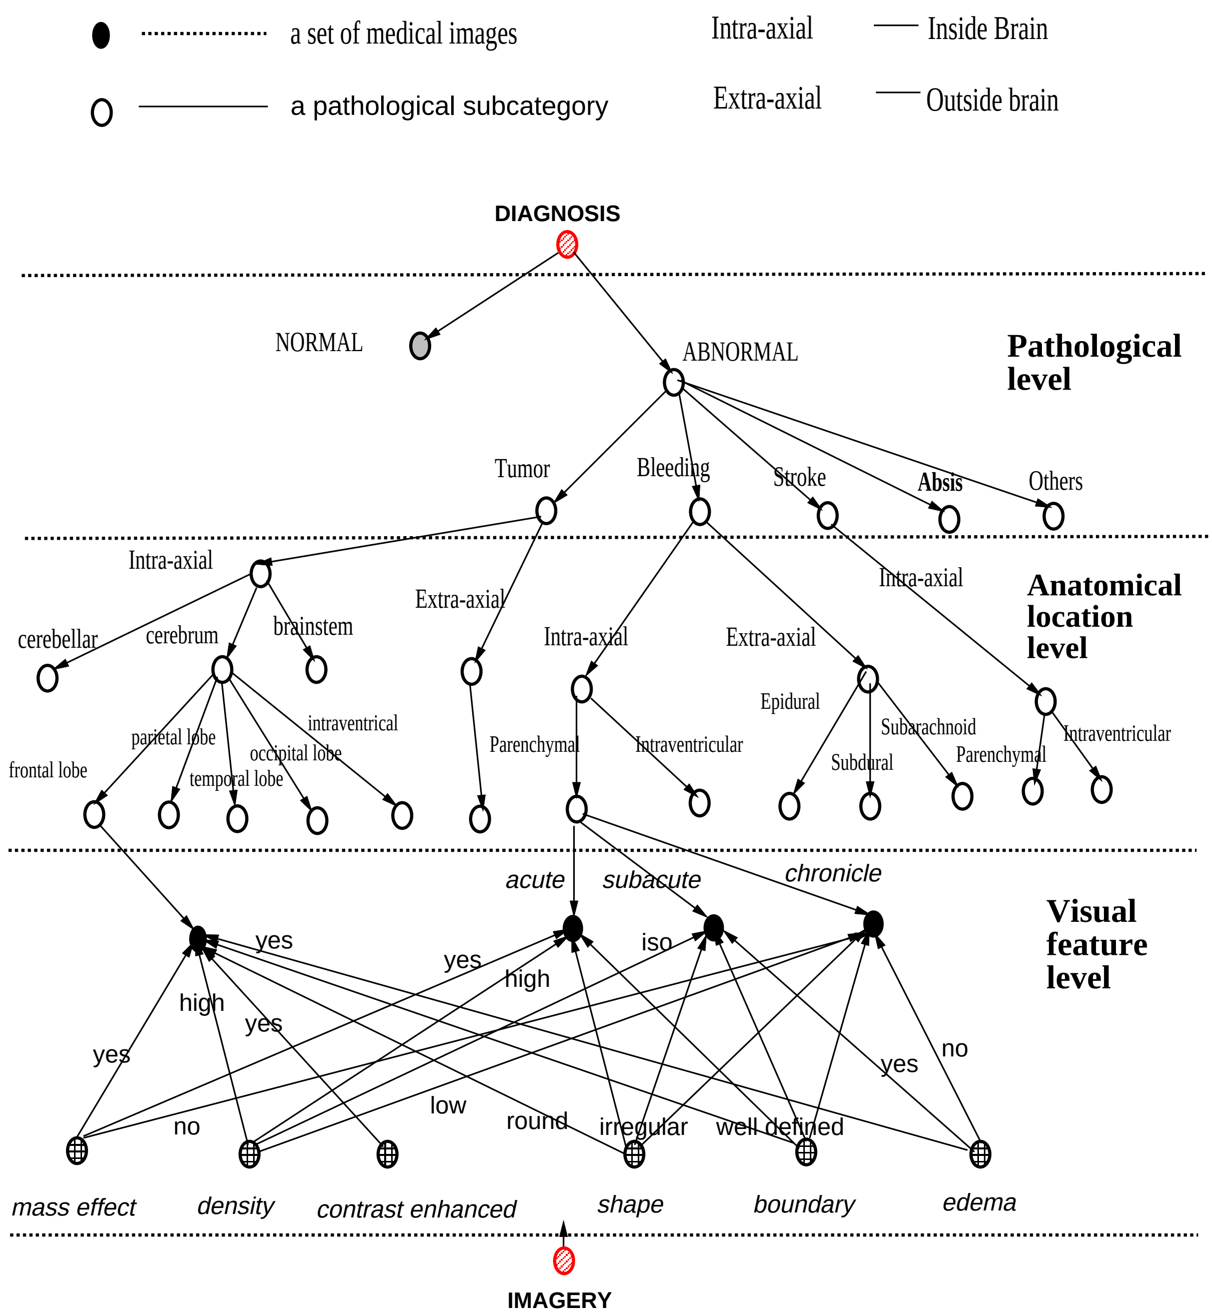

Figure 4: Pathological cases are classified by causes (bleed, stroke, tumor ...), anatomical locations (intra-zial, extra-axial, ...) and then visual features (mass effect, lesion density, shape, boundary, ...). For clarity, only a couple of the bottom nodes are drawn.

extra-axial), and lastly by their visual features. This data hierarchy serves to structure a massive set of entries in a database into smaller subsets based on their content. Such a data structure can only be effective if we have a discriminating indexing module that can lead us to the relevant cases quickly. The key question is how we can reach the right branch at the pathological level by starting from the visual feature level (Figure 4).

Each image can be viewed as a data point in some multidimensional feature space. The feature vector, the coordinates of a 3D image, functions in turn as the image index. With the help of neuroradiologists, the following salient visual features, which have significant semantic meanings and medical implications in interpreting brain images, have been identified:

- *mass effect*: asymmetry with respect to the ideal center line, due to structural/density imbalance
- *anatomical location*: where the lesion resides in terms of the brain's 3D anatomical structure
- *density*: relative brightness and darkness of the lesion
- *contrast enhancement*: lesion sensitivity to contrast enhancement
- *boundary*: the region between the lesion and its surroundings
- *shape*: a characterization of the 3D volume the lesion occupies
- *edema*: lucent area around a lesion, usually caused by excessive liquid
- *texture*: the texture of the lesion
- *size*: the dimension of the lesion
- *age:* some visual features vary with patient's age.

Figure 5 is a sketch of a flow chart for constructing an image index. The basic idea is to build a set of feature detectors or demons [10] which can capture the values of relevant salient visual features in the input image. Some initial design and experiments on building some of the important feature detectors are described below.

3D Symmetry Plane Detector: Since normal human brains present an approximate bilateral symmetry, which is often absent in pathological brains, a *symmetry detector* is constructed to automatically extract the ideal midsagittal plane (the plane $X_0 = 0$, see Figure 1) from a 3D, probably pathologic, brain image [17]. The basic idea is to find the transformation that brings a given 3D brain scan and its bilateral reflection (also a 3D image) into the best matching position. This can be achieved by the following two methods: 1) compute the yaw (roll) angle and offset of the bilateral symmetry axis for each 2D axial (coronal) slice via cross-correlation of the slice

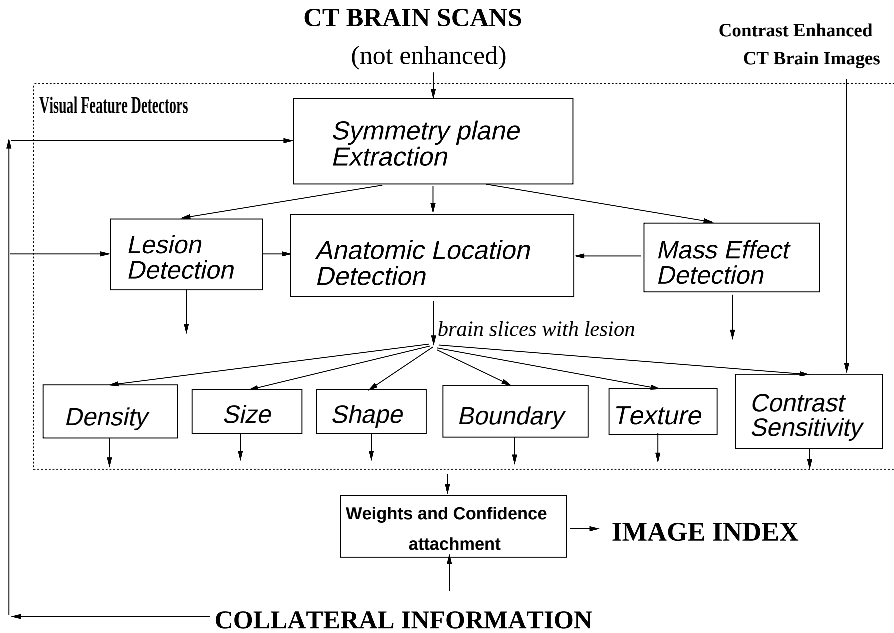

Figure 5: Image indexing by finding the salient visual feature values using visual feature detectors. Both images and collateral information are used for this purpose. Note, some brain slices can be screened out from further examinations due to the absence of lesion.

and its bilaterally reflected image, fit a 3D symmetry plane to all the 2D symmetry axes in 3D space thus determine the roll (yaw) angle of the symmetry plane. 2) directly register the given 3D brain with its bilateral reflection (about the Y-Z plane in Figure 1) using maximization of mutual information theory under affine transformations [18]. This symmetry detector has been tested on both CT and MR images of normal and pathological images. The absolute rotational accuracy of the extracted symmetry plane is less than one degree when compared against human experts (neuroradiologists), the relative accuracy is less than one degree as well when reslicing a 3D brain image with varied yaw and roll angles. One of the results of the symmetry plane detection is to locate an ideal bilateral symmetry axis on each 2D slice as shown on Figure 6. Another result of this process is to find the yaw and roll angle errors, the difference between the working coordinate system and the ideal head coordinate system, in the input 3D image data. An image resampling process is followed to minimize these errors.

Lesion Detector: Using the results from the symmetry detector we have further developed the *lesion detector* which aims at automatically locating possible lesions (bleeds, stroke, tumors) by detecting asymmetrical regions with respect to the extracted central symmetry plane (Figure 7). Figure 8 shows the outlines of lesion hypotheses found in one set of CT brain scans. The goal is to make this process adaptive and robust to different image densities, for example, acute blood

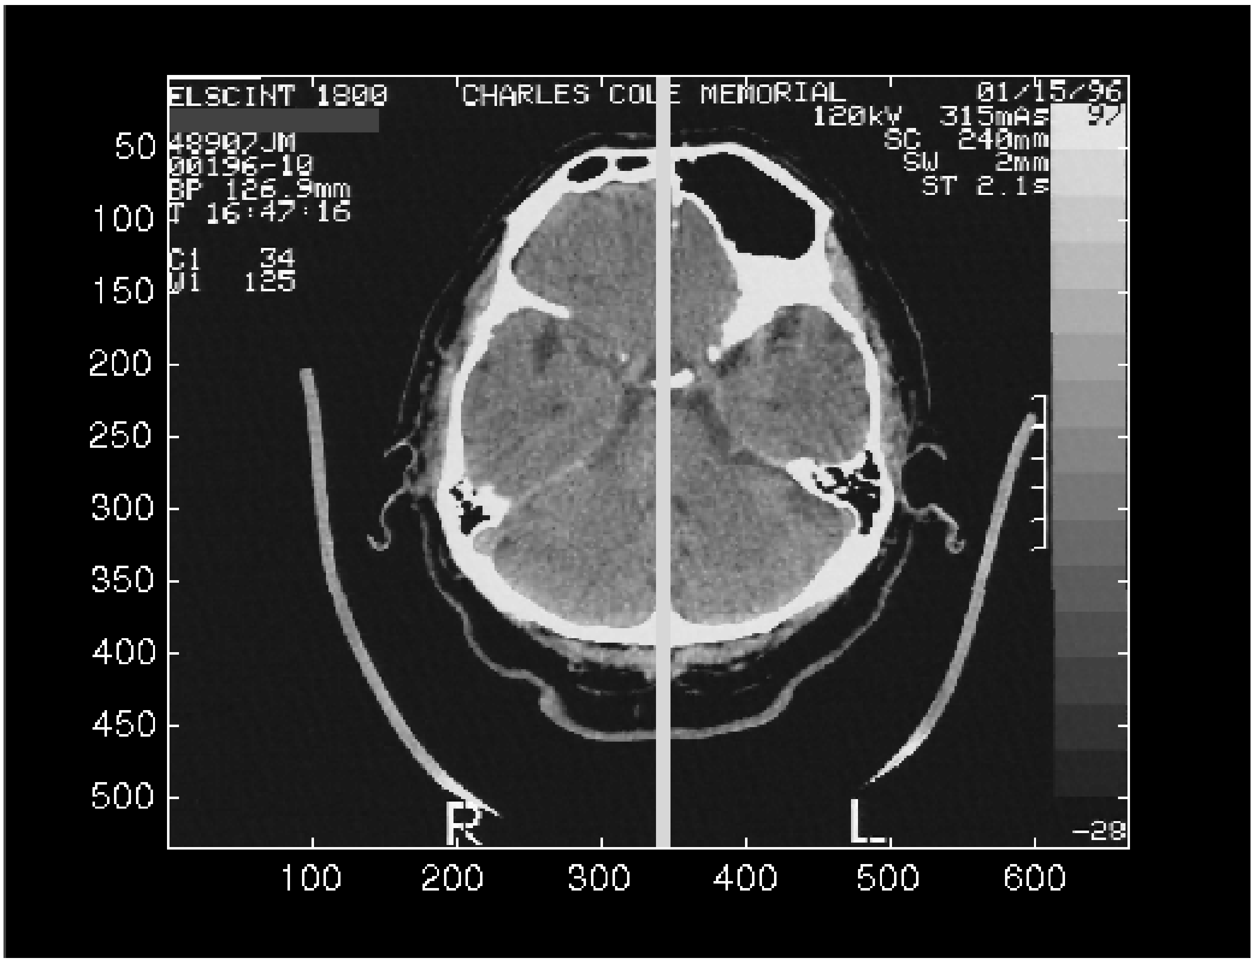

Figure 6: The ideal symmetry axis is extracted as the intersection of a 3D brain image and its midsagittal plane. This is a CT scan with obvious asymmetry due to a sinus near the frontal lobe.

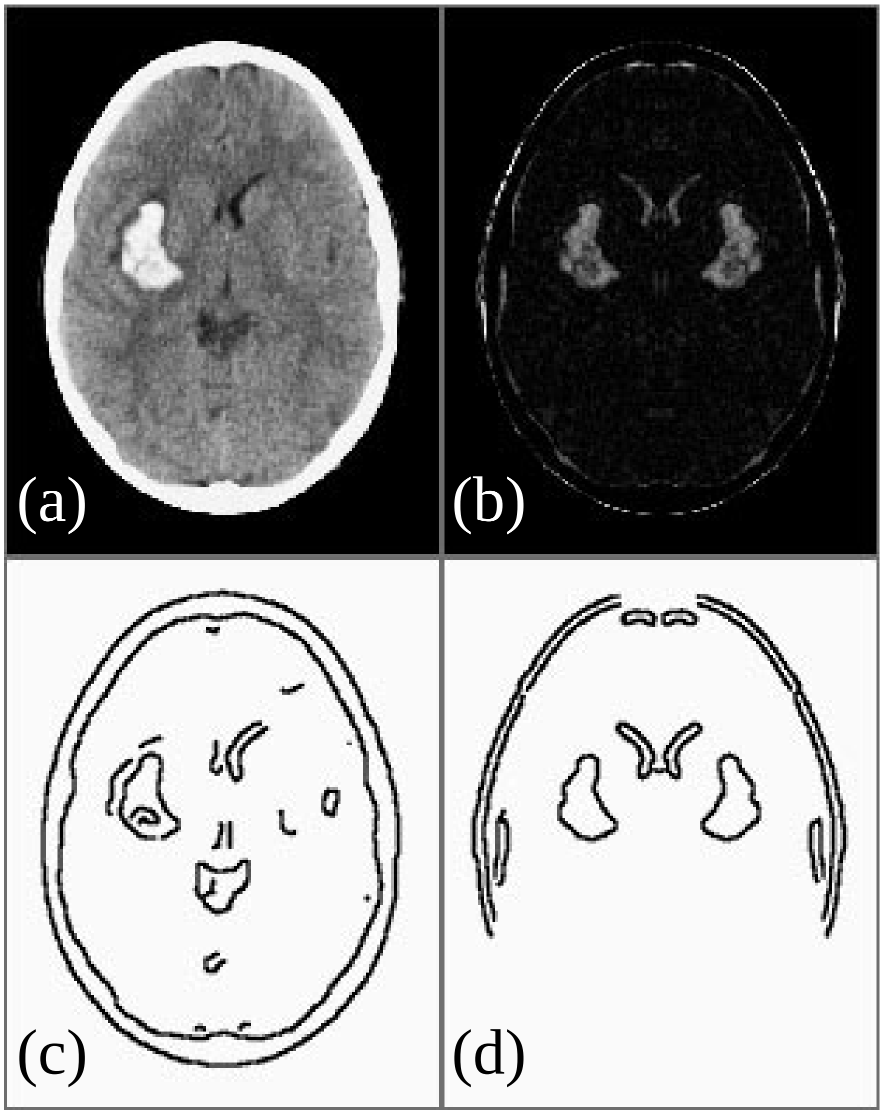

Figure 7: (a) an axial slice S with yaw and roll angle errors corrected such that the symmetry axis is centered and vertical in the image; (b) Absolute value of the difference image between S and its vertical reflection (c) and (d) are the chained canny edges from (a) and (b). Lesions and mass effect are detected by matching curves between (c) and (d). The result of the matching is shown in Figure 8, the left most image on the second row.

appears white on an CT image while acute infarct (stroke) appears dark.

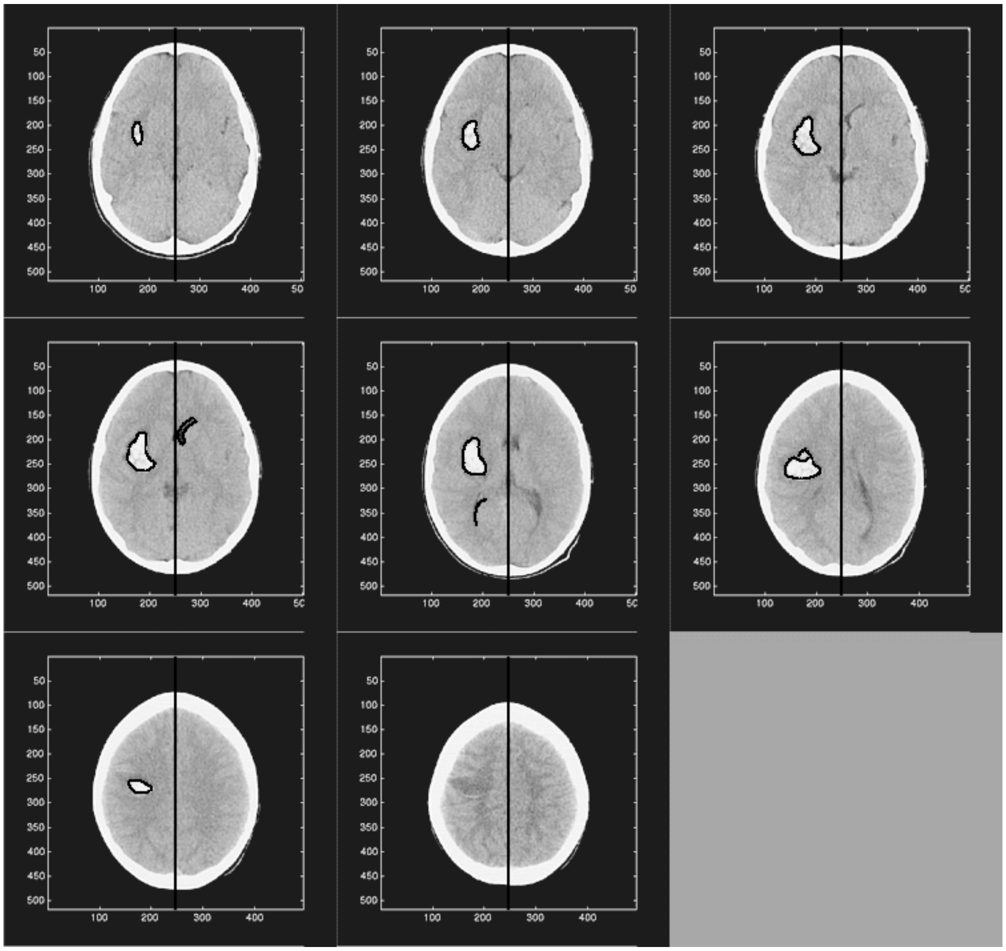

Figure 8: Using the center lines automatically obtained by our midsagittal plane extraction algorithm, the contours of potential lesions are found by reflecting the image about the center lines. Notice that nothing is found in the last 2D slice.

Mass effect detector: The extracted symmetry axis is used as the initial position of an open snake [14]. The final resting position of the snake indicates how much the brain has shifted from its ideal centered position due to a tumor (Figure 9). The difference between the deformed midline of a pathological brain and ideal midline, the ratio of the maximum distance of the two curves over their vertical length, is used as a quantified measurement of mass effect.

Anatomical location of the lesion detector: A 3D digital brain atlas has been obtained from the Harvard Medical School [16]. This atlas is composed of a 3D MR T1 image (the 5th image in Table 1) and an equivalent 3D image with each voxel labeled to a specific anatomical region. In order to determine the anatomical location of a detected lesion in a 3D image, the atlas is

Figure 9: Using an active contour (open snake) to detect mass effect as deformation of the ideal midline, which is automatically extracted using our midsagittal plane extraction algorithm.

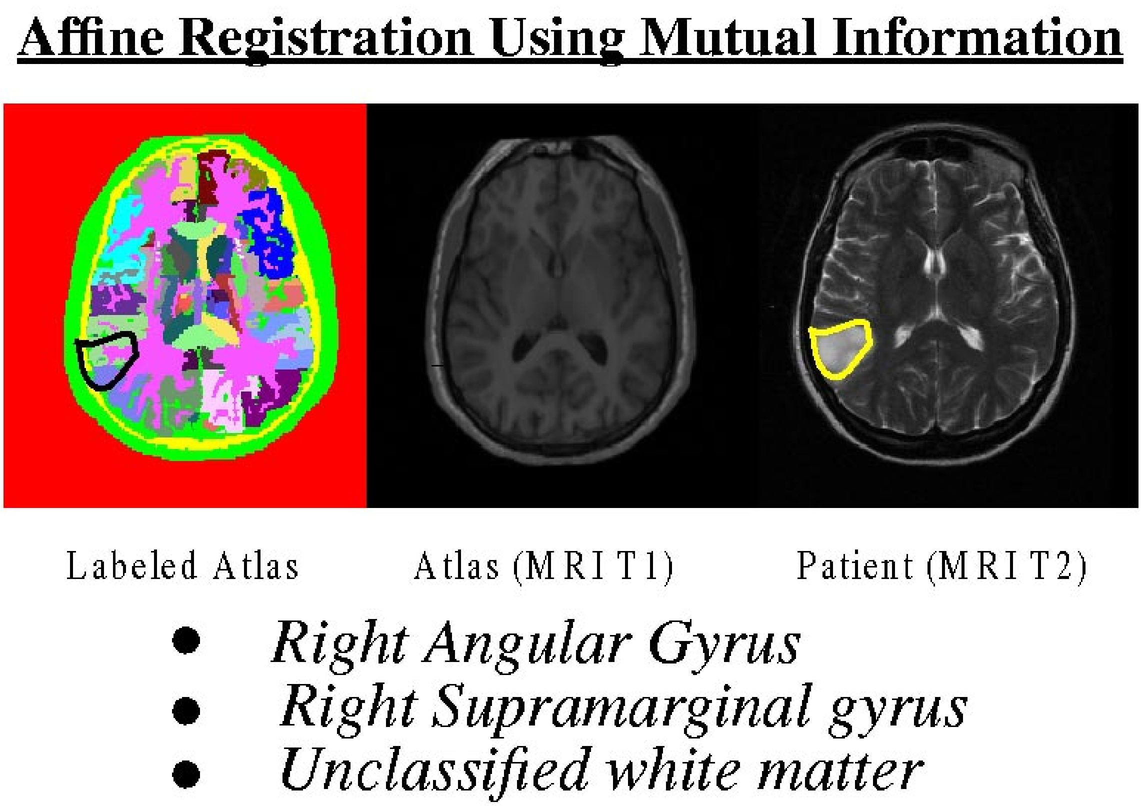

Figure 10: Using maximization of mutual information method to register a 3D digital brain atlas (T1) to a patient 3D image (T2), such that the anatomical location of a detected lesion can be identified.

deformably registered onto the pathological brain. Figure 10 shows the result of a deformable registration (affine warping) from a 3D digital atlas (MRI T1 image) to a brain with lesion (MRI T2 image). Since the atlas is completely labeled, the general anatomical location of the lesion can be identified from the labeled voxels. While we are working on more sophisticated automatic deformable registration of a brain atlas to a pathological 3D brain image for exact matching, an interactive tool has been used that allows the user to identify corresponding anatomical points on the images, and a simple warping algorithm [2] warps one image to the other. Figure 11 gives one glimpse of the process: given the midsagittal planes (automatically extracted using our symmetry plane detector) of both the patient and the atlas, the user draws a few corresponding line segments and clicks a button to instantly creat a deformed image.

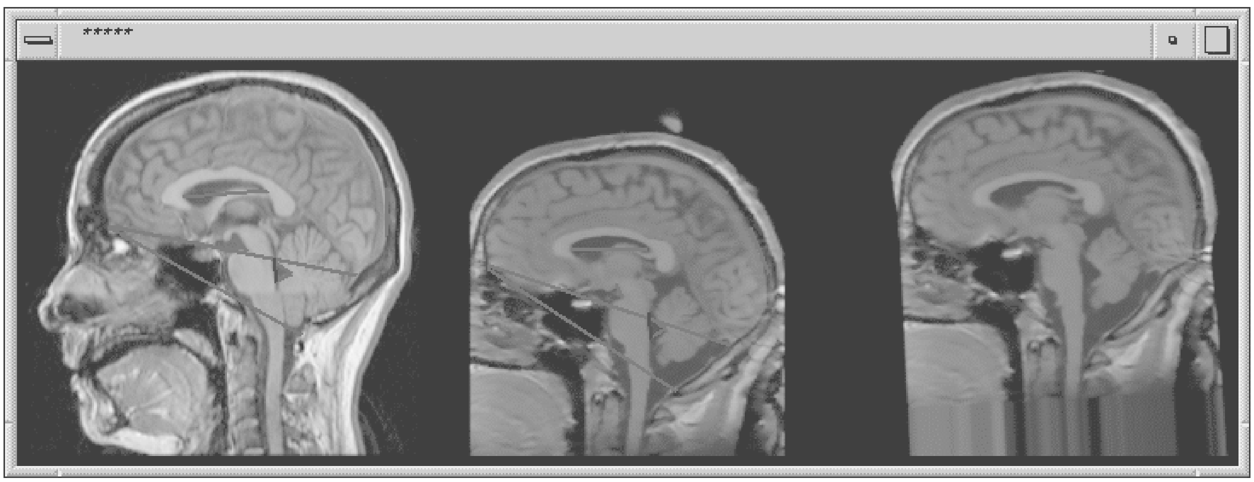

Figure 11: Left: midsagittal of the patient with corresponding lines drawn by user, Middle: midsagittal of the atlas with corresponding lines, Right: the atlas midsagittal deformed into the patient's midsagittal.

3D shape of a lesion: The unit Gaussian sphere has been chosen by several researchers as the space for a quantitative representation of the 3D shapes. Experiment on 3D shape similarities of lesions is carried out on a meshed Gaussian sphere using a simplex attribute image (SAI) [12, 13] to characterize both the convex and concave surfaces. It has been shown that the comparisons of 3D shapes can be carried out using local curvature on each mesh point of the sphere for a 3D surface effectively[24]. However, it remains to be seen whether relatively sparse and unevenly distributed surface data extracted from a lesion in CT or MR scans has sufficient information to carry out such a comparison.

4 Image Retrieval

Assume the image index is an N dimensional vector, each dimension corresponds to an identified visual feature discussed above. The overall similarity function S of two 3D images a, b is a simple metric: $S(a, b) = \Sigma w_i S_i(a, b)$ where w_i is the weight for feature i, $\Sigma w_i = 1$, $i \in [1, ..., N]$, $S_i(a, b)$ is the similarity of images a and b in feature space i:

$$S_i(a,b) = \begin{cases} 100 & \text{if } D_i(a,b) = 0 \\ 0 & \text{if } D_i(a,b) > H_i \\ 100 * (1 - D_i(a,b)/H_i) & otherwise \end{cases}$$

Where D_i is the distance function for feature i and H_i is a cutoff threshold for feature i. The interesting question is how the distance function in each feature space is defined. For example, the distance function for lesion density can use the difference of the relative average density values of the corresponding lesions directly, while the distance function for anatomical locations of the lesions is quite complicated. The latter requires a good understanding of brain anatomy and functionality to construct a predetermined "anatomical location distance" look-up table.

For our initial interactive testbed, exact match of the feature values and a weighted sum of the difference between two feature vectors are used. The weights of features can be adjusted by the user interactively. A simple JAVA user interface panel has been constructed in order to visualize the retrieval results. Figure 12 shows the panel displaying the visual feature values (in text format) of the target and the highest ranked retrieved image case. Figure 13 shows the corresponding target (left) and the top ranked retrieved images. Note, although the bleeds in the target and the retrieved images appear on opposite sides of the brain, they belong to the same brain anatomical structure, while the second ranked retrieved 3D image shown on the right of Figure 14 has an acute bleed on the same side of the brain as the target case but in a different anatomical location. The resulting ranking for the retrieved 3D images concurs with an expert's opinion, albeit on a small database of only ten cases. Our long term goal is to build a sufficiently sized database ($>$ 1000 cases), with which a statistical-based adaptive learning scheme [3, 20, 6] will be used to capture the most discriminating visual feature patterns and their respective weights for effective and efficient image retrieval. The hit rate and similarity rankings of the retrieval system will be statistically monitored and evaluated against human experts periodically (see next section).

5 Evaluation

Our preliminary experiments on constructing 3D medical image indexing and retrieval demo showed promising result, but the true value of a retrieval system partly relies on how it can be

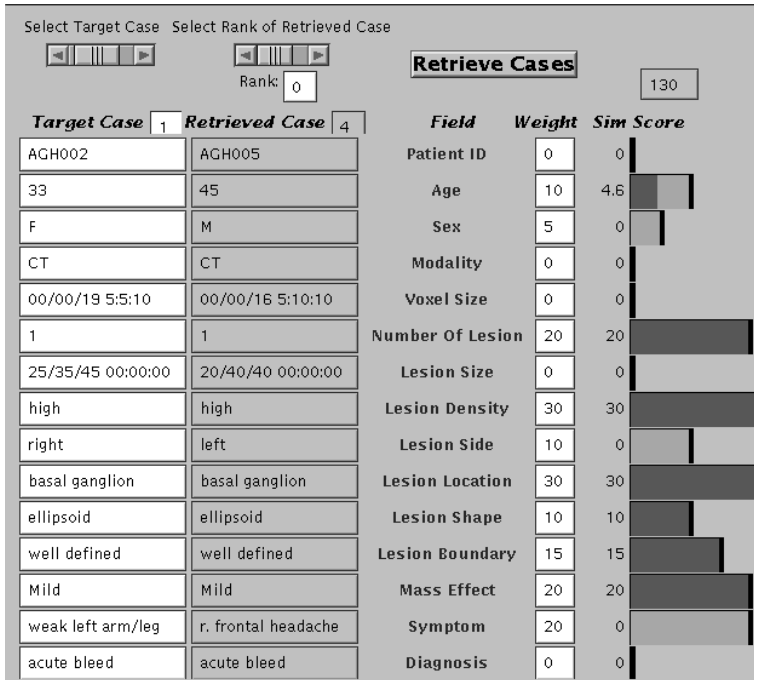

Figure 12: The Medical Image Understanding System testbed. Left: target case number 1, Right: the top ranked (rank 0) retrieved case number 4

evaluated and improved quantitatively. The chosen image domain provides easy measures for evaluation, i.e. the human expert's medical diagnosis associated with each 3D image. There are at least two quantitative measures we use. One is the top ten *hit rate* $= C/10$, here C is the number of the coincidence between the classifications of the top ten retrieved images and the diagnosis given by a human expert for the input image. The other is the *ranking correlation* meaning the cross correlation between the similar images produced by the retrieval system and the ones identified by a human expert. Here the ranking can be a partial ranking, i.e. ties are allowed. Note, sometimes the diagnosis is not a clear-cut thus these evaluation measurements have to be computed based on the probability in each possible diagnosis class.

6 Summary and Conclusions

We have learned that image understanding of approximately symmetrical, 3D biological objects can be guided by a simple principle: the possibility of functional abnormality is implied by a departure of the object from its normal symmetrical structure. The positive results from some of the feature detectors, especially the lesion detector described in Section 3, is benefited from this

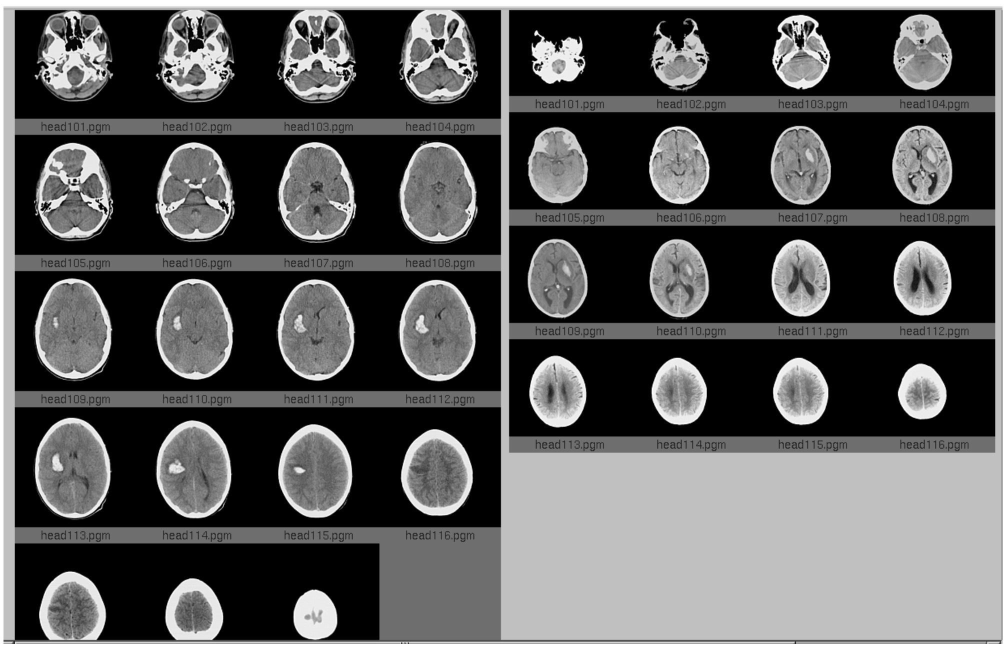

Figure 13: The corresponding images in the previous Figure. Left: target case, which is a right basal ganglion acute bleed. Right: top ranked retrieved case, which is a left basal ganglion acute bleed. Although the bleeds appear on opposite sides of the brain, they belong to the same brain anatomical structure.

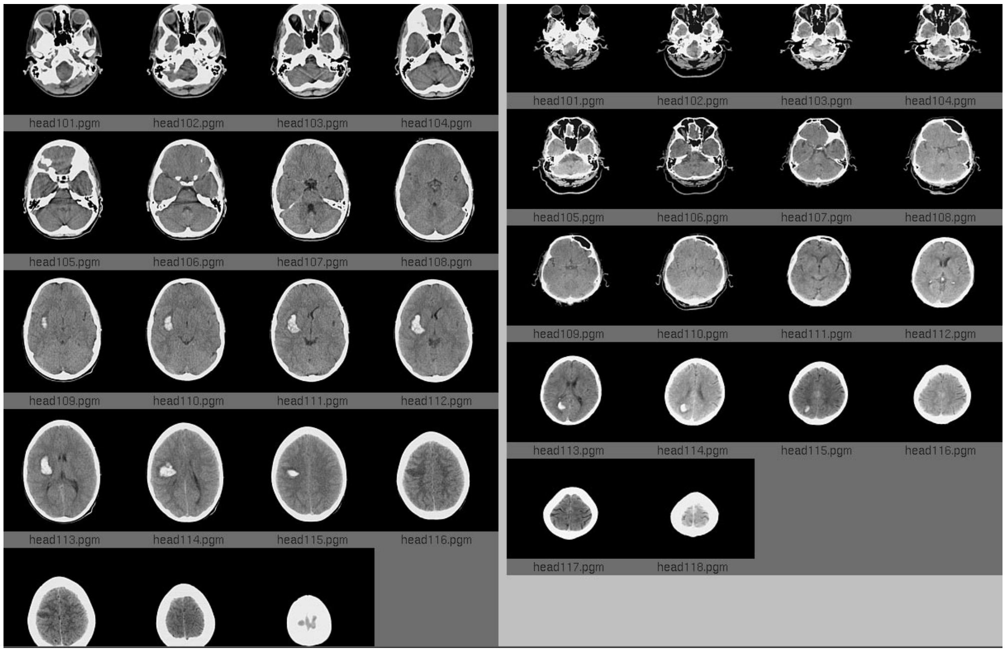

Figure 14: Left: same target case as in the previous Figure. Right: 2nd ranked retrieved case image indicating a right occipital/parietal acute bleed.

realization. However, many open problems remain, for example:

- treatment of sparse and incomplete image data

  The sparseness and incompleteness of CT images have caused problems in image resampling to correct yaw and roll angle errors, 3D shape representation/comparison, and deformable volumetric registration. Experiment results have shown that CT scans which contain more lower cuts with complicated bony structures give better registration result. A dense CT atlas may help to enhance the sparse and incomplete CT image.

- multi-modality deformable registration between a brain atlas made of normal subject and a pathological brain

  This is useful in two folds in the context of medical image retrieval: one is to determine the exact anatomical location of a given lesion in a patient's image to fill in the value of the anatomical location feature; and two is to use the result of the registration directly in the detection of the existence of possible lesion, i.e. identify the portion in the patient's image which cannot be registered with the atlas, unusual density for example. This requires the system to have sufficient knowledge on how each anatomical structure deforms (rigidity, range etc.) and their variations across populations.

Neverthless the initial result from this 3D medical image retrieval project is encouraging. In the near future, we foresee to continue our research under the basic framework stated in this paper, but with an emphasis on probability-based learning to enhance the indexing (classification) process as well as to increase the accuracy of retrieval (similarity analysis) on a much larger 3D neuroradiologic image database (500+).

Acknowledgments

This research is sponsored by the Allegheny-Singer Research Institute under prime contract through the National Institute of Standards and Technology (NIST#70NANB5H1183). Views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of the National Institute of Standards and Technology or the United States Government.

The authors would like to thank Dr. Ron Kikinis from the Brigham and Woman's hospital, Harvard Medical School, Boston, for providing the brain atlas; and thank the Laboratory for Medical Imaging Research, ESAT-Radiologie, K.U. Leuven, Belgium, for sharing their multimodality image registration software.

References

[1] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Computer Vision, Graphics, and Image Procession*, 46:1–21, 1989.

[2] T. Beier and S. Neely. Feature-based image metamorphosis. In *SIGGRAPH*, pages 35–42, 1992.

[3] C. M. Bishop. *Neural Networks for Pattern Recognition*. Clarendon Press, 1995. ISBN:0198538499.

[4] C. Davatzikos. Spatial transformation and registration of brain images using elastically deformable models. *Comp. Vis. and Image Understanding, Spefial Issue on Medical Imaging*, May 1997.

[5] J. Declerck, G. Subsol, J-P Thirion, and N. Ayache. Automatic retrieval of anatomical structures in 3d medical images. Technical Report 2485, INRIA, Sophia-Antipolis, France, 1995.

[6] R.O. Duda and P.E. Hart. *Pattern Classification and Scene Analysis*. John Wiley & Sons, New York, 1973.

[7] D. Faloutsos, R. Barber, M. Flickner, J. Hafner, W. Niblack, D. Petkovic, and W. Equitz. Efficient and effective querying by image content. *Journal of Intelligent Information Systems*, 1994.

[8] Y. Gong, H. Zhang, H.C. Chua, and M Sakauchi. An image database system with content capturing and fast image indexing abilities. *IEEE International Conference on Multimedia Computing and Systems*, May 1994.

[9] V.N. Gudivada and V.V. Raghavan. Content-based image retrieval systems. *Computer*, pages 18–22, September 1995.

[10] L.J. Guibas and C. Tomasi. Image retrieval and robot vision research at stanford. In *The Proceedings of the Image Understanding Workshop*, pages 101–108. ARPA, Morgan Kaufmann Publishers, Inc., 1996.

[11] A. Guimond and G. Subsol. Automatic mri database exploration and applications. *Pattern Recognition and Artificial Intelligence*, 11(8), December 1997.

[12] K. Ikeuchi. Recognition of 3d objects using the extended gaussian images. In *IJCAI Conference*, pages 595–600, 1981.

[13] K. Ikeuchi and M. Hebert. From EGI to SAI. Technical Report CMU-CS-95-197, The Robotics Institute, CMU, 1995.

[14] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, pages 321–331, 1988.

[15] P.M. Kelly and T.M. Cannon. CANDID: Comparison algorithm for navigating digital image databases. In *Proceedings of Seventh International Working Conference on Scientific and Statistical Database Management*, pages 252–258, September 1994.

[16] R. Kikinis, M.E. Shenton, D.V. Iosifescu, R.W McCarley, P. Saiviroonporn, H.H. Hokama, A. Robatino, D. Metcalf, C.G. Wible, C.M. Portas, R.M Donnino, and F.A. Jolesz. A digital brain atlas for surfical planning, model-driben segmentation, and teaching. *IEEE Transactions on visualization and computer graphics*, 2(3):232–240, September 1996.

[17] Y. Liu, R.T. Collins, and W.E. Rothfus. Automatic extraction of the central symmetry (mid-sagittal) plane from neuroradiology images. *Carnegie Mellon University, The Robotics Institute Technical Report*, 96(40):1:29, November 1996.

[18] F. Maes, A. Collignon, D. Vandermeulun, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, 16(2):187,198, 1997.

[19] T. McInerney and D. Terzopoulos. Deformable models in medical images analysis: a survey. *Medical Image Analysis*, 1(2):91–108, 1996.

[20] T.M. Mitchell. *Machine Learning*. McGraw-Hill, New York, 1997.

[21] M Moshfeghi. Elastic matching of multimodality medical images. *CVGIP:Graphical Models and Image Processing*, 53:271–282, 1991.

[22] A. Pentland, R.W. Picard, and S. Sclaroff. Photobook: Content-based manipulation of image databases. *IJCV*, 18(3):233–254, June 1996.

[23] R.W. Picard. A societry of models for video and image libraries. *IBM Systems Journal*, 1997.

[24] H. Shum, M. Hebert, and K. Ikeuchi. On 3d shape similarity. In *IEEE CVPR*, pages 526,531, June 1996.

[25] M.J. Swain. Interactive indexing into image databases. *SPIE*, 1908, 1993.

[26] J.P. Thirion. Fast intensity-based non-rigid matching. In *Proc. of 2nd Intl. Symp. on Med. Robotics and Comp. Ass. Surgery*, pages 47–54, 1995.

[27] P. Viola. Entropy, information, computer vision and image processing. In *Ph.D., MIT,*, 1995.

[28] P. Viola and W. Wells III. Alignment by maximization of mutual information. In *ICCV95*, pages 16–23, 1995.

[29] W.M. Wells III, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis. Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1(1):35–51, March 1996.