Development of Computer Vision Techniques to Support the Clinical Study of Ischemic Stroke Treatment

by
Y. Chitti, K. Muhammad, G. Whitten,
B. Stein, D. Geman & J. Horowitz

CMPSCI Technical Report 98-45
October, 1998

Progress Report for the Collaborative Biomedical Research Program between the Baystate Medical Center and the Biotechnology Program at the University of Massachusetts on

Development of Computer Vision Techniques to Support the Clinical Study of Ischemic Stroke Treatment

Yasmina CHITTI, Kamran MUHAMMAD, Gary WHITTEN, Edward M. RISEMAN
Computer Science Department

Benjamin STEIN, Don GEMAN and Joseph HOROWITZ
Department of Mathematics and Statistics

Abstract

The long-term goal of this collaboration between the Baystate Medical Center and the University of Massachusetts is to establish a working understanding of the relationship between blood pressure treatment and ischemic stroke patient outcome. An intermediate goal is the determination of the relationship, if any, between changes in stroke lesion size and patient outcome. To this end, the University of Massachusetts is developing a system to automatically determine stroke lesion volume using a custom graphical user interface for a technician to supply an initial focus of attention.

This report describes progress and results for the first six months of the project.

Introduction

This project is a research collaboration between the Baystate Medical Center, led by Dr. Pleet, Head of the Neurology Department, and Dr. Hicks of in the Radiology Department, and two departments of the University of Massachusetts: the Computer Science Department and Mathematics and Statistics Department. The long-term goal of this project is to develop and study treatment protocols for ischemic stroke patients. A stroke occurs when blood does not properly supply a region of the brain. The result may be some degree of brain damage that is due to damaged or dead cells. The damaged region is known as infarct, an abnormal brain region, or a lesion.

The general goal of the clinical study is to establish a working understanding of the relationship between blood pressure treatment and patient outcome and our specific goal is to develop semi-automated methods to quantitatively determine the volume of abnormal brain regions and relate it to patient condition.

Our subgoals are to provide a simple user interface to facilitate interactive specification of initial lesion location, to allow the user to reshape the identified lesion, and to automatically identify lesion extent in adjacent slices. From this, a lesion will be automatically detected on every slice and its volume will be computed.

Many techniques have already been applied to medical imaging [1] but they have not yet been integrated into a system for automatic volume determination.

One approach to lesion detection is based on image segmentation by using active contour models [6, 11]. An active contour model, or "snake", is defined by an energy functional which is a weighted combination of internal and external energies. Snakes dynamically alter their shape and position in order to seek a minimal energy state. The initial snake contour will be defined by a human operator via a human interface and the final shape of the snake will represent the edge of the lesion.

Another approach in lesion detection is statistical. We examine samples of intensities in an image, and use them to determine a way to segment it into its "lesion" and "non-lesion" components. The results of the statistical segmentation can be naturally integrated with the snake lesion representation by expressing the results of the segmentation as a force applied to the snake. We will refine our segmentation results by exploiting edge orientation information and generalizing our statistical models, as described below. We will also integrate lesion segmentation results from single slices into a full three dimensional representation by adding an orthogonal set of snakes that extend two dimensional smoothness constraints to three dimensions. These future efforts are also described below.

1 Motivation

1.1 What is a stroke lesion?

The human brain consists of more than 10 billion cells. These cells use a surprising amount of energy in the form of oxygen and glucose. While the brain represents only 2% of total body weight, it consumes up to 25% of the oxygen used by the body.

A stroke occurs when an artery to the brain bursts or becomes clogged, thereby stopping or interrupting the blood supply to the brain. Deprived of the blood's essential oxygen and nutrients, even for only a few minutes, the brain cells begin to die. Dead cell regions are known as infarcts or lesions. Since dead brain cells do not regenerate, damage from a stroke is often permanent. Depending on what portion of the brain is affected by the stroke, the result may be some degree of brain damage and loss of function affecting speech, vision, or memory. A stroke may also result in paralysis, coma or death. Cells that are damaged but potentially viable are known as penumbra.

If a stroke occurs, blood pressure is elevated at the time of admission to the hospital (blood pressure greater than 140 mm Hg systolic and/or greater than 90 mm Hg diastolic).

1.2 Study of treatment

The long-term goal of this project is to provide an effective treatment to people suffering acute strokes. Studying possible treatments involves the following questions:

- Is there improvement or worsening in outcome when blood pressure is treated (scoring on neurological and functional scales) ?
- Should blood pressure treatment be vigorous, or minimal?
- Is there a change in the ischemic penumbra with treatment,and can doctors predict the effectiveness of the treatment?

For the study, affected individuals will be divided into three "treatment" arms:

- No treatment for blood pressure.
- Treatment of blood pressure until it is normal.
- Treatment of blood pressure until it is halfway to normal.

Each patient will have a quantitative neurological examination every day for the five days they are in the hospital, and again 30 days following discharge. At the time of admission another scaling measurement will be used to assess their physical abilities (as opposed to neurological abnormalities). This measurement will be repeated 30 days after discharge from the hospital. A scan of the brain will be obtained just prior to discharge from the hospital (approximately day four or five) and will be repeated on the 30 day follow–up visit.

Both the use of a tool that allows the physician to compute the change in volume of the lesion between two scans and the study of the relationships between the change in volume and the treatment can readily change treatment modalities. The physicians plan to study the consequences of treating blood pressure as well as the effect of "cerebral protective" agents in acute stroke, including calcium channel blocking drugs, NMDA channel blocking drugs, throbolytic agents ("clot busters"), and a host of others.

1.3 Imaging the lesion

Magnetic Resonance (MR) Imaging is a technique that produces images of the body using Nuclear Magnetic Resonance properties of tissues [10]. MR imaging does not use ionizing radiation and can produce contrast between different tissues as white matter, grey matter, infarct, etc, based on a variety of physical parameters as the repetition time of a sequence of radio-frequency pulses (TR), spin-echo delay time (TE) or the inversion time (TI). Therefore a great variety of information can be obtained by appropriate selection of imaging criteria. The variations of these time parameters change the type of MR images. For example, `pilot` images are obtained with small time settings for fast acquisition of pilot scans, `T1` images correspond to a short TR and TE times and tend to contain little noise, `T2` are acquired using long TE and usually also long TR times. With these typical settings, water is imaged light. Other usual time settings include `PD` that combines a relatively long TR with a relatively short TE time and `flair` for which TE and TR are set as for `T2`, but a nonzero TI is used. This can be used to selectively null out-tissue types with certain properties. For instance, TI $\sim$ 2000 ms will cause water (and thus CSF) to be imaged dark [4, 7].

The data for our primary study include:

To cover the entire brain, we have three views for each patient: axial (from bottom to top, the right half of the image corresponds to the left half of the patient's head), sagittal (from side to side) and coronal (from front to back). We have between 10 and 23 images taken in each direction that cover almost the entire brain. These images are called slices because they look like the results of cutting the brain into

many slices, which offer a view into the brain's interior. On our images each slice usually represents a thickness of the brain of 5 millimeters with a gap of 2 millimeters between two slices. These thicknesses are adjustable, and we have some data where the gap is 0 and the thickness of the slice is 3 millimeters. Most of the brain's scans have been taken in the three modalities: `T1`, `T2` and `Flair`. An example of each type is shown Fig. 1 as an example of axial, sagittal and coronal views is shown in Fig. 2.

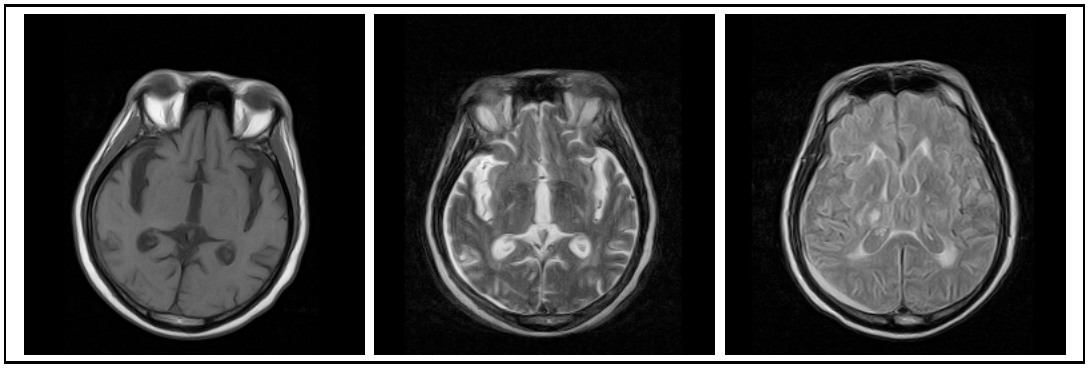

Figure 1: Different protocols of MR images showing the same axial view of the brain. From left to right: `T1`, `T2` and `Flair` type.

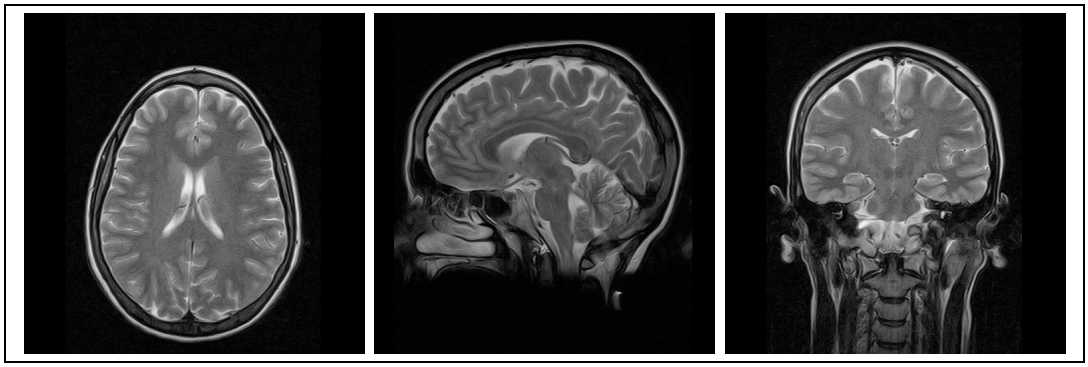

Figure 2: Different views of MR images showing the same brain. From left to right: axial, sagittal and coronal views. MR type is flair in all case.

From these images of the patient's brain our goal is to provide to the physician a tool that allows him to compute the change in volume of the lesion between two sequences of MR images. Relevant questions include:

- Can we define the central core of brain death and separate it from surrounding penumbra cells, which are areas that have a decreased blood supply but might recover?
- Can we see a change in the ischemic penumbra with treatment that allows the physician to predict the effectiveness of the treatment?
- Which kind of MR images can we use to detect the region that represents the "core" of cerebral death or infarction and the region that represents the ischemic penumbra?
- Can we provide a three-dimensional representation of the lesion to help the physician see where the changes occur?

2 Determination of lesion's boundaries

The primary goal in our analysis of brain imagery is to segment the lesion, i.e. classify the pixel of the brain image as lesion or non-lesion pixels. The shape of the lesion is seldom spherical or elliptical. In fact, a lesion can have any shape, size, or location in the brain, (however, the shape is usually relatively smooth) and the edges of the lesion are often indistinct. Thus, finding a way to determine the pixels representing an abnormality on a scan, and then computing the volume of this abnormality, poses a significant problem. An example of a lesion in an axial view of the brain is shown in Fig. 3 and the magnified lesion is shown in Fig. 4.

One of our first tasks was to study image examples of normal and lesion brain regions. This required the manual indication of abnormal brain areas, which were provided to us by Dr. Pleet. Our studies have indicated that contrast is most pronounced in `Flair` and `T2` biased imagery.

There are candidate approaches to the segmentation problem under consideration. At this point we have tested two of them. One approach is based on adaptative contour models and the other one is based on local statistical analysis around a lesion. Both of these methods and the results we have obtained for each will be described further. Since the results we have obtained are encouraging, we are trying to merge these two methods to improve the detection in some of the most difficult cases.

2.1 Statistical study of the lesion

Our statistical segmentation approach involves studying statistical properties of the imagery. We focus on `T2` images since they exhibit the most contrast and `Flair` se-

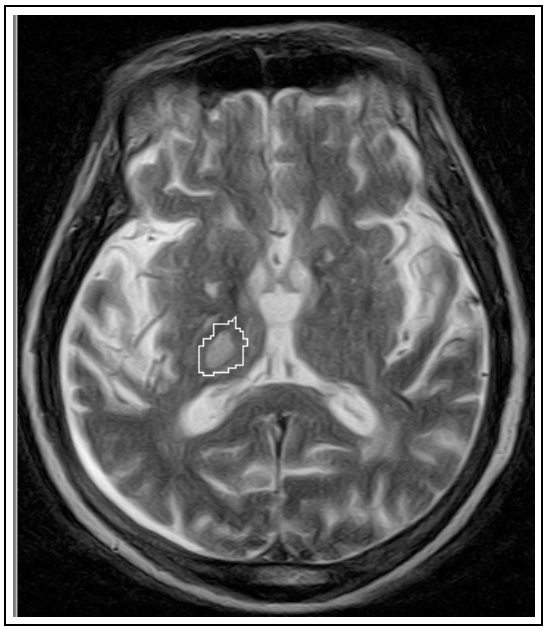

Figure 3: This image shows an axial view of the brain. The lesion is circled in the middle. The lesion is unusually clear.

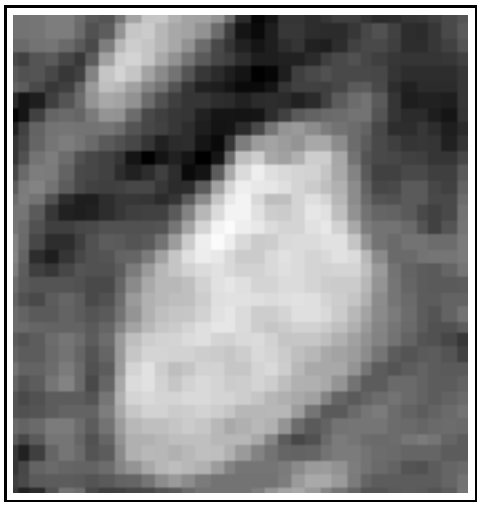

Figure 4: Magnified image of the lesion circled in Fig. 3.

quences are not available for every patient.

If we look at the distribution of the intensity in the T2 axial view of the brain in Fig. 3) containing an infarct as is shown on Fig. 5, we see that there are at least two distinct peaks in the distribution. Unfortunately, neither of these two regions represents pixels belonging to the lesion alone, and in fact, pixels corresponding to the lesion are lost in the lower intensity peak. This implies that lesion detection from a whole brain slice will be problematic.

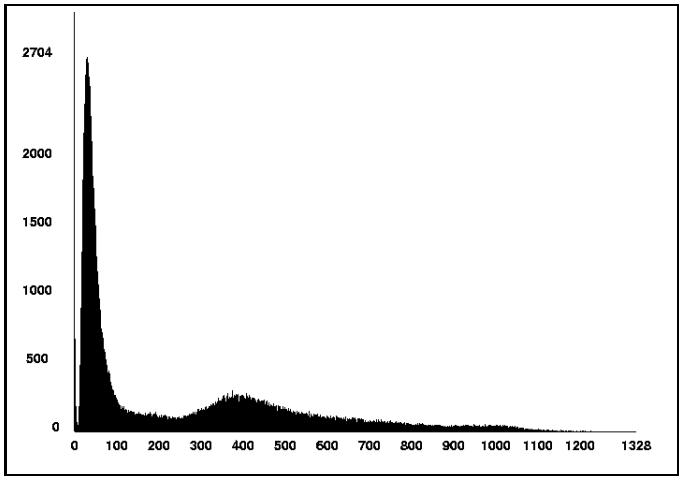

Figure 5: This figure shows the distribution of the intensity of the full brain image slice in Fig. 3. There are two distinct regions but neither represents the intensity of the pixels belonging to the lesion.

The graph in Fig. 6 represents the intensity distribution of two samples of pixels in the magnified view of the lesion in Fig. 4. The left peak represents intensities of a sample of pixels that are outside of but close to the lesion and the right peak represents intensities of a sample of pixels that belong to the lesion. Notice that the grey levels for the lesion data are much higher than for the non-lesion data.

The two peaks of this histogram can be smoothed into continuous density functions using the theory of density estimation [3, 9]. Once these functions are obtained, a maximum likelihood classification scheme is applied to each pixel: if the pixel's grey level is more likely to come from the "lesion" density than the "non-lesion," classify the pixel as "lesion," and otherwise, classify it as "non-lesion."

Beyond this straightforward classical approach, a more complex, and accurate, statistical model can be developed, by creating an appropriate simultaneous autoregressive model [2]. In order to do this, we must first determine a neighborhood

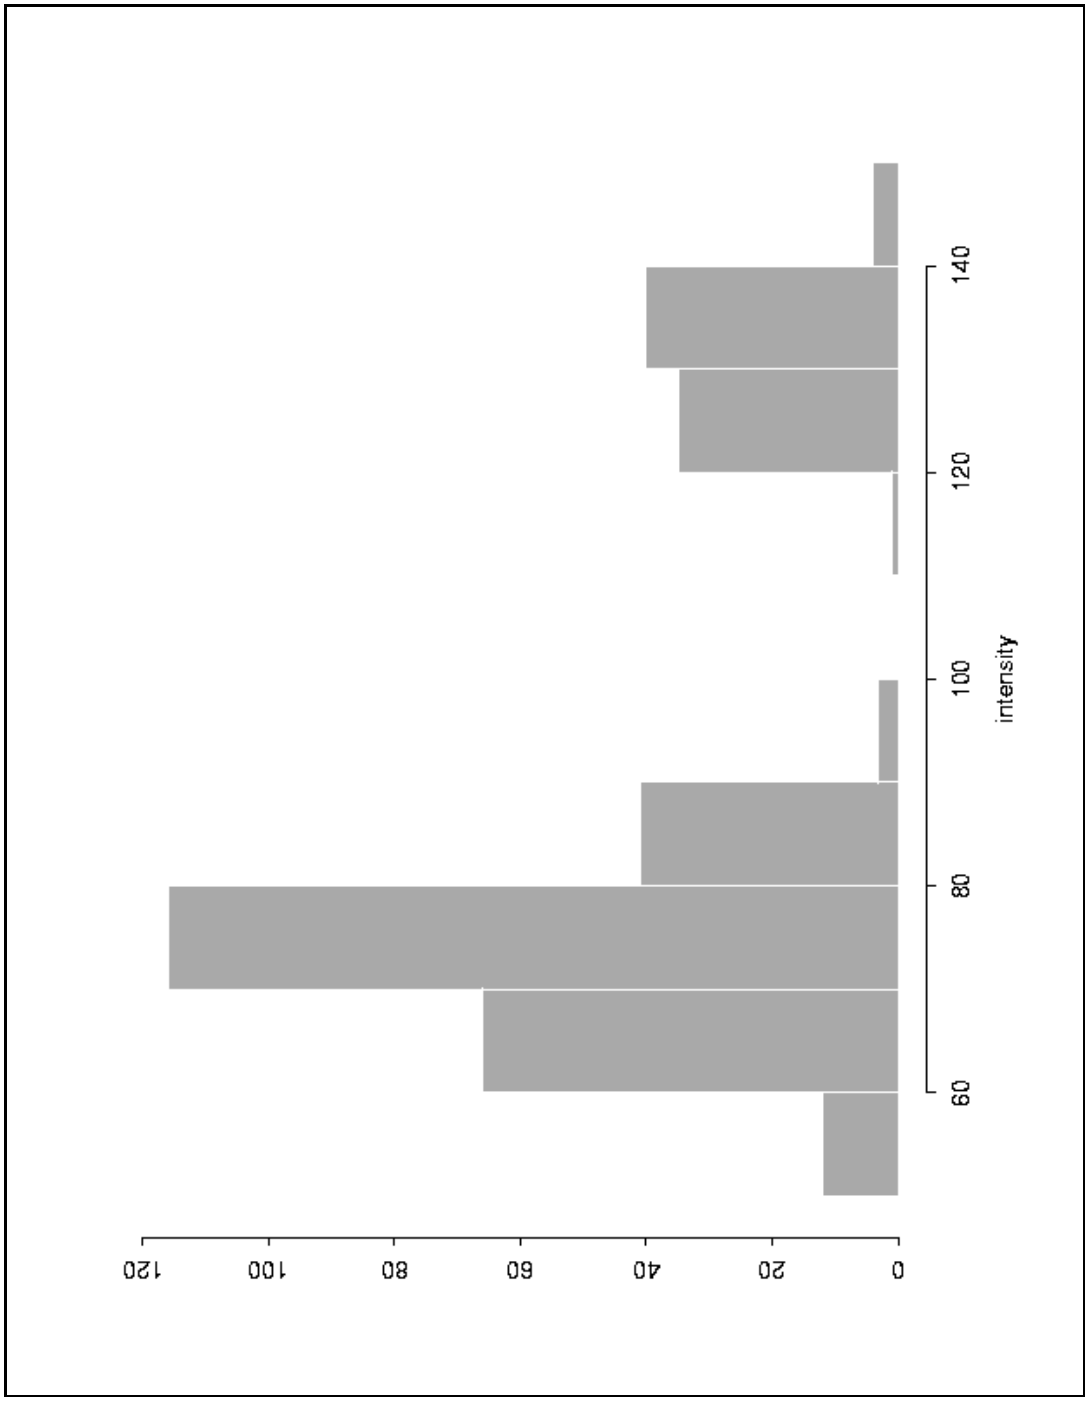

Figure 6: Intensity distribution of the magnified view of the lesion in Fig. 4. The left peak represents non–lesion pixel intensities while the right peak represents lesion pixel intensities. Only a sample of pixels has been taken into account.

system and the amount of dependence between neighboring pixels. Such an investigation is ongoing.

2.2 Active Contour Models – Snakes

Snakes can be thought of a special case of a more general technique of matching a deformable model to an image by means of energy minimization. Snakes are typically classified as active contour models because they dynamically alter their shape and position while trying to seek a minimal energy state [6].

Snakes are generally represented by a set of control points that define their contours.The actual definition of a snake comes as an energy functional, which is a weighted combination of internal and external energies. The snake's goal is to alter its shape and position to minimize its total energy.

The total energy of a snake can be expressed as:

$$\mathcal{E}_{snake} = \sum \mathcal{E}_{internal} + \sum \mathcal{E}_{external} \tag{1}$$

2.2.1 Internal Energy

Using a parametric definition of the snake contour, $v(s) = (x(s), y(s))$, the internal energy of a snake is expressed as two terms:

$$\mathcal{E}_{internal} = \alpha \left| \frac{dv}{ds} \right|^2 + \beta \left| \frac{d^2 v}{d^2 s} \right| \tag{2}$$

The two terms of the internal energy of a snake control it's smoothness. The effects of these terms are controlled by weights, α and β. Specifically, α controls the stretching or elasticity of the snake. Large values of α greatly increase the internal energy of the snake as the snake stretches. Small values of α, on the other hand, make the snake insensitive to the amount of stretching that the snake undergoes.

The second term of the snake internal energy represents the snake stiffness or flexibility smoothness constraint. Large values of β greatly increase the internal energy of the snake as it develops more curves. Small values of β, on the other hand, make the snake insensitive to curves in the snake contour. An extreme example exists where β is set to zero– the snake may then become second-order discontinuous and develop a corner.

2.2.2 External Energy

The external energy of a snake guides the snake toward image features such as lines and edges. The external energy can be defined as:

$$\mathcal{E}_{external} = \mathcal{E}_{edge} + \mathcal{E}_{ext1} + \ldots + \mathcal{E}_{extN} \tag{3}$$

The edge-based energy term causes the snake to be attracted to areas of the image with large gradients:

$$\mathcal{E}_{edge} = -\left|\nabla f(x,y)\right|^2 \tag{4}$$

Other external energy terms are also possible, and are usually defined with a problem domain in mind.

2.2.3 Suitability of Snakes to Our Problem–Domain

Having explained the details of snakes, it is useful now to consider why we have chosen this technique for our problem domain. The traditional weakness of snakes has been the need for a good local initialization of the snake. In addition, snakes are not normally suited to solving an entire segmentation problem autonomously. Typically, they depend on mechanisms such as interaction with a user.

Both of these traditional weaknesses turn out to be strengths within our problem domain. Firstly, the initialization of the snake fits in very well with the operation protocol we have agreed on with our medical partners at Baystate. Specifically, we have agreed that a physician/technician will draw at least one contour that surrounds a region of interest (lesion) in a single slice within a patient image data set. This initial contour will serve as the initialization for the first snake.

Related to user interaction with snakes is the inherent intuitive nature of snakes. Most energies that are defined on a snake can be equated to real energies within the physical world (i.e. spring energy). With this in mind, it becomes easier for a user to understand the behavior of snakes as the image context and parameters are varied.

Another feature of snakes is that the external energies, and forces derived from them, are highly customizable. Clearly, this allows us to take advantage of the properties of our imagery and other techniques that may contribute to improving snake segmentation. Furthermore, these external forces are always held in check by the snake's internal smoothness constraints, thus always maintaining a snake with a reasonable shape and size.

Finally, a property of snakes that is especially important to us deals with handling image degradation. Image degradation is particularly evident around lesion boundaries. Often these boundaries are not well defined and their contours are broken. Since the snake itself is a contour, it essentially "fills in" these breaks in lesion boundaries during energy minimization.

2.2.4 Summary of Overall Approach using Snakes

The overall approach to solving the 3–D lesion volume computation problem follows:

1. The user initializes the first snake in a single slice by manually drawing a contour surrounding the lesion.
2. The snake seeks an optimal configuration representing the lesion boundary.
3. The final position of the snake in the initial slice serves as an initialization for new snakes in adjacent slices.
4. These new snakes also seek optimal configuration and the process repeats from step 2 until the lesion is identified in all slices.
5. This process is repeated for the other one or two orientations of MRI scans. The final segmentation data in three dimensions is integrated to yield a lesion volume.

2.3 Features for a graphical user interface

Lesion identification is greatly simplified, as described above, if a human operator interacts with the process. Therefore, a user interface is required and must provide a means for the user to:

- Load all the images in a sequence.
- Choose the image which contains the largest area of the lesion.
- Zoom in on the region of interest (including the lesion).
- Draw an approximate closed contour around the lesion (including some area outside the lesion) to initialize the segmentation process.
- Monitor the results of lesion boundary detection on every slice of the sequence.
- Interact with the system to correct the boundaries in any of the image of the sequence. (Note that we do not expect this to be necessary in the production version of the system.)
- Examine a 3D visualization of the lesion.
- Calculate the computed volume of the lesion.

3 Progress

In the first six months of the project (from the end of January 1998 to the end of July 1998) we have implemented and tested promising segmentation techniques and developed a graphical user interface.

3.1 Lesion detection

The two methods – statistical segmentation and snakes – have been tested on three levels of difficulty presented in Fig. 7. The three images in Fig. 7 show how a lesion can be tiny with sharp edges or very large with fuzzy edges.

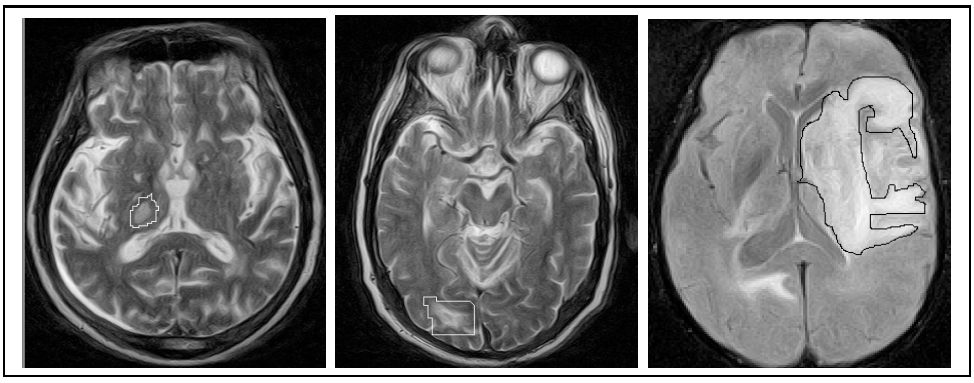

Figure 7: Axial views of three different brains showing lesion. The two left images have been taken in `T2` modality and the right one has been taken in `Flair` modality. In each image, the lesion has been marked and is demontrating how the shape of lesions varies. The lesion in the left image is small and has well–defined edges. The lesion in the middle image has a shape not well-defined and edges are not contrasted. The lesion in the right image occupies a large part of the brain, but edges are not well-defined and their detection is not simple.

The results of lesion detection of the three sample images in Fig. 7 are shown in Fig. 8, 9 and 10 respectively. In each of these three figures, the left image corresponds to a magnified view of the lesion, the middle image is the result of statistical segmentation in which white pixels represent the lesion, and the right image is the result of the snake algorithm. In the right image, the drawn contour represents the final boundary of the lesion. These images show that the detection either by statistical segmentation or by using snakes gives good results. In almost every case,

we can detect pixels belonging to the lesion with both algorithms, and often there is a reasonable segmentation of the full lesion.

However, these results must be improved. First, the statistical segmentation does not detect fuzzy edges in Fig. 10, middle. In addition, the statistical segmentation detects every region that has the intensity range of the pixels belonging to the lesion as white regions appear around the central lesion in Fig. 8 and 9. This implies that other regions outside the lesion are detected and would have to be removed by the expert which is not acceptable. The snake algorithm we have tested is based only on the computation of the local gradient of the image to constrain the snake and the snake is not stable. Other forces have to be applied to keep the snake stable. Also, the result depends on the initial contour drawn by the expert. If, as Fig. 8, right shows, the initial contour is too far away from the true contour of the lesion, intermediate and final snakes will not be accurate. Thus, there is significant research still to be carried out.

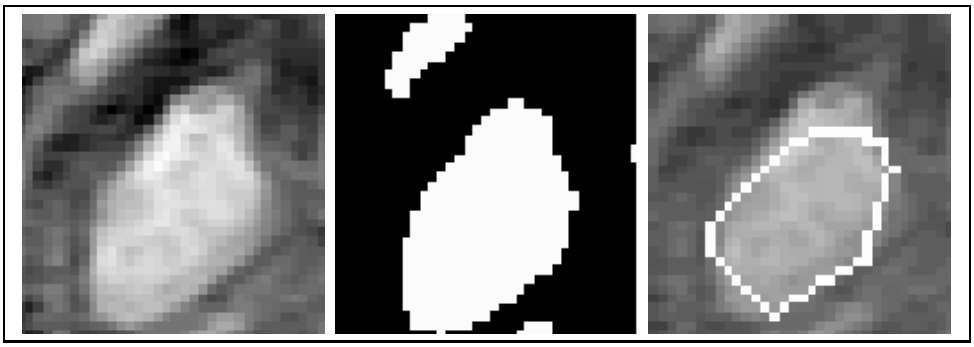

Figure 8: Detection of small lesion with sharp edges.
Left: Magnified image of the lesion.
Middle: Result of the statistical segmentation of the left image.
Right: Final snake outlining the lesion.
In this case the segmentation gives good results. One problem we notice is this technique detects not only the lesion pattern but also other fragments outside of the lesion that might or might not be lesion. In the case of the snake technique, the accuracy of the result depends on the initial contour drawn by the expert. The right image demonstrates that when the initial contour is too far away from the true contour as in the snake we have tested, problems occur.

Figure 9: Detection of lesion with not very well–defined edges.
Left: Magnified image of the lesion.
Middle: Result of the statistical segmentation of the left image.
Right: Final snake outlining the lesion.
The segmented image displays many other regions that are not part of the lesion, but still find the shape of the lesion. The snake was well–initialized and gives a nice contour of the lesion.

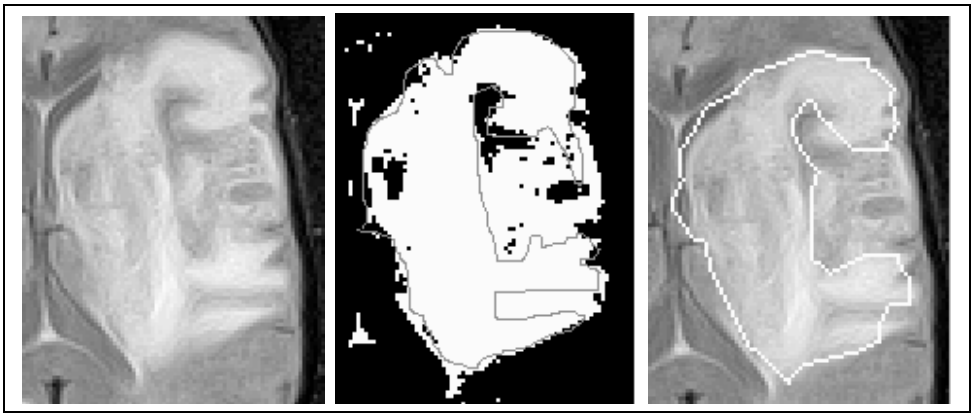

Figure 10: Detection of lesion with fuzzy edges.
Left: Magnified image of the lesion.
Middle: Result of the statistical segmentation of the left image.
Right: Final snake outlining the lesion.
In the segmented image, the detected shape does not match with the superimposed contour, drawn from doctor's indication. The snake fits the contour better than the segmented image.

3.2 Graphical User Interface (GUI)

The Java Image Prototype is being developed to provide a graphical interface as a tool for researchers so that the semi-automated process of computing the lesion's volume can be achieved. When this study is completed the GUI will be modified for more completely automated processing in a clinical production mode.

The Graphical Interface has been written in the Java 1.1.6 platform. This will allow us to use it on any computer, especially on PCs running Windows95, but also on any other computer running a graphical environment (such as XWindows). The features of the GUI that have already been implemented are:

- A browser allowing the user to load all the images on the sequence in one view (either axial or coronal or sagittal) from the number (or name) of the patient. Each image has a reduced size of 100x100 pixels and is labeled. The number of the patient and the number of the sequence appear on the top.
- From the browser or from the main menu the user can load an image of full resolution on the main screen.
- The user can select a rectangular region of interest and zoom in on it. The user is not allowed to zoom the whole image to avoid the limitation of small computer screens.
- On the main frame or on the zoomed image, the user can draw a free-hand contour, surrounding the lesion.
- From this contour, the user can select the detection of the lesion's boundaries on any or all images for a given patient.
- If the user is not satisfied with the results of the edge detection, he can select the images where he thinks the detection is correct.
- When the user is satisfied with the edge detection in every image of the sequence, he can compute the global volume of the lesion.

Other graphical features such as scaling the grey level intensity or applying a color model to an image to enhance the contrast have also been implemented.

4 Current Study and Future Efforts

We are continuing to refine our segmentation techniques by exploiting edge orientation information. Evidence for potential lesion edges is increased if the edges have

consistent orientation. That is, if adjacent candidate edges have similar orientation, the complete edge tends to be continuous and the edge components are consistent. Conversely, if adjacent edge segments have significantly different orientations, the resulting complete edge would tend to be fragmented, indicating that at least some edge segments are false. This information can be used to increase the sensitivity of lesion contour detection (often required for ambiguous edge regions) without significantly increasing false edge extraction.

Our work, to date, has focussed on lesion segmentation in single image slices. Our plan is to extend single slice segmentation to adjacent slices to include the full lesion extent. Our approach uses the segmentation results from the first slice to initialize segmentation in adjacent slices. Therefore, the first result serves the same role in adjacent slices as the human operator uses in the first slice. This process continues until the lesion size falls below a minimum threshold.

We will use a set of snakes, oriented orthogonally to the set of individually extracted lesion slices, to impose 3D smoothness on all snakes representing the lesion. This will provide a full 3D representation of the extracted lesion and will also provide global feedback forces to individual snakes that are not consistent with a smooth 3D shape.

As time permits, we will also investigate other promising techniques, such as scale-space decomposition using wavelets, which have recently been used for detection and 3D reconstruction/registration [5, 8]. Another area of investigation that could improve the quality of MRI imagery is super-resolution processing. The current imaging process results in non-isotropic imagery, where in-plane (the plane of the slice) resolution is much higher than out-of-plane resolution. By registering and integrating orthogonal views, it may be possible to generate composite imagery that has optimum resolution in all directions simultaneously.

5 UMass/Baystate Meeting on July 28, 1998 – Future Data

We had a major meeting on July 28, 1998 with Dr. Pleet and Dr. Hicks at the Baystate Medical Center. The goal of this this meeting was to show our progress to the doctors and discuss the recent upgrades to the MR equipment at Baystate.

The two doctors were very enthusiastic about the results we have obtained. We estimate that we can improve our current results with better resolution data, that is expected to be delivered in approximately one month. Dr. Hicks explained to us that it is possible to have 3D data, using `T1` protocol, and maybe `T2` protocol. This kind of acquisition gives a voxel equal to $1 \times 1 \times 1$ mm in the case of the

T1 protocol. He will provide this kind of data from a volunteer patient because he expects the patient to be in the MRI scanner for as much as two hours. He will also provide the usual 3 views scans but with varying resolutions; in particular he can reduce the thickness of the slices to 4 or 3 millimeters and reduce the gap to 1 or 0 millimeter. He will focus his attention on the part of the brain where the lesion is and mark the lesion on separate hard copy film of the images. Increasing the resolution will decrease the signal. Thus, from this set of data we will test how the detection is resolution and signal dependent and choose the best balance.

Conclusion

This six-months work has shown that the techniques we are testing and adapting in order to perform lesion detection could be useful for the physician to study the effectiveness of treatment in ischemic stroke. This work has also shown that it will be difficult to provide a completely–automated system to detect the lesion in the patient brain. The idea is to detect the lesion from cues provided interactively by a physician. From this starting point our current results indicate that we will be able to detect the lesion in the brain and compute it's volume.

Continuation funding for the next six months will allow us to develop and improve 3D detection of the lesion and to evaluate the accuracy and the robustness of the methodology.

Aknowledgments

This work was funded by the Collaborative Biomedical Research Program including the Baystate Medical Center, Springfield, Massachusetts and the Biotechnology Program, University of Massachusetts.

References

[1] *Proc. of the Workshop on Mathematical Methods in Biomedical Image Analysis*, San Fransisco, CA, June, 21–22 1996. IEEE Computer Society Press.

[2] N. Cressie. *Statistics for Spatial Data.* Wiley and Sons Inc., New York, 1993.

[3] D. Geman, J. Horowitz, and C. Li. Tissue Maps Labeling and Classification. Unpublished report, Univ. of Massachusetts, Amherst, 1997.

[4] J. Gong and J.P. Hornak. A Fast T1 Algorithm. *Magn. Reson. Imaging*, **10**:623–626, 1992.

[5] A. Grossmann and R. Kronland-Martinet. *In Signal Processing IV: Theories and Applications*, chapter Time and Scale Representation obtained throught Continuous Wavelet Transforms., pages 475–482. Elsevier Science Publishers, 1988. J.L. Lacoume., A. Chehikian., N. Martin and J. Malbos Eds.

[6] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active Contour Models. *Int. J. Comp. Vision*, 1:321–331, 1987.

[7] X. Li and J.P. Hornak. T2 Calculations in MRI: Linear vs Nonlinear Methods. *J. Imag. Sci. & Technol.*, **38**:154–157, 1994.

[8] B.S. Manjunath and R. Chellappa. A Unified Approach to Boundary Perception: Edges, Textures, and Illusory Contours. *IEEE Trans. on Neural Networks*, **4**:96–108, 1993.

[9] B.W. Silverman. *Density Estimation for Statistics and Data Analysis.* Chapman and Hall, London, 1986.

[10] D.D. Stark and W.G. Bradley. *Magnetic Resonance Imaging.* C.V. Mosby Co., St. Louis, MO, 1988.

[11] D. Terzopoulos and T. Mc Inerney. Deformable Models in Medical Image Analysis. In *Proc. of the Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 171–181, San Fransisco, CA, June, 21–22 1996. IEEE Computer Society Press.

Appendix A: Baystate/UMass CBR Program Conference

On Friday, May 8, 1998, the Baystate/UMass Collaborative Biomedical Research Program Oversight Committee hosted their annual half day conference. We presented our work there in a 20 minutes talk. Following are the slides shown during this presentation.

Development of Computer Vision Techniques to Support the Clinical Study of Ischemic Stroke Treatment

Baystate Medical Center, Springfield
University of Massachusetts, Amherst

People

BAYSTATE	UMASS	
Neurology: Dr. B. Pleet Radiology: Dr. R. Hicks	Math. Dept. : Prof. D. Geman Prof. J. Horowitz B. Stein	Comp. Sc. Dept : Prof. E. Riseman Dr. G. Whitten Dr. Y. Chitti K. Muhammad
How effective is the treatment of stroke lesions ?	• Graphical User Interface (GUI) • Allow the doctor to circle the lesion • Compute the change in volume of the lesion at two different times	

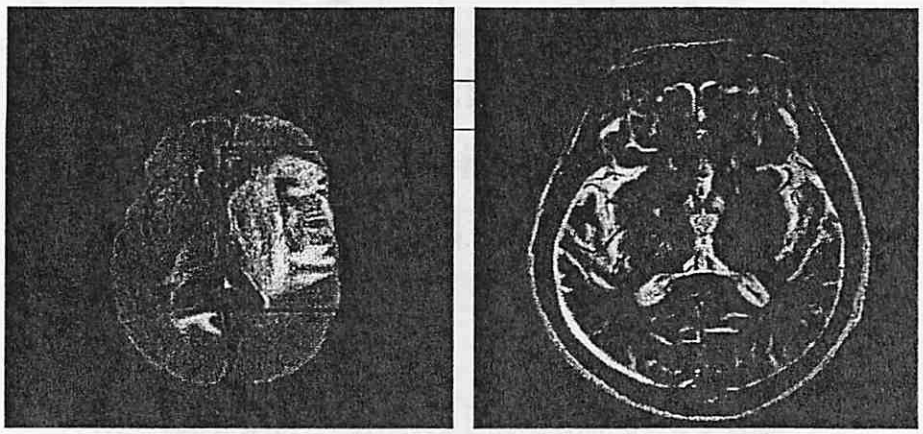

- 3 views of the brain
- around 20 slices a view
- lesion could be on only a few slices (3)
- lesions could be small
- contours of lesions are fuzzy

Our Approach

- Provide a GUI to the doctor : circle the lesion
- Lesion detection (from first contour)
  - provide a semi-automated detection of the lesion
  - Snake Segmentation
  - Statistical Segmentation
- Volume computation
  - Compute the volume from every slice of the brain
  - Extension of two-dimensional methods to 3D

Talk Overview

- Provide a GUI to the doctor : circle the lesion
- Lesion detection (from first contour)
  - provide a semi-automated detection of the lesion
  - Snake Segmentation (K. Muhammad)
  - Statistical Segmentation (B. Stein)
- Volume computation
  - Compute the volume from every slice of the brain
  - Extension of two-dimensional methods to 3D
  - Problem of gap between slices to reconstruct without errors

Snakes - Overview

- Active contour models
- Dynamically change shape and position
- Defined as energy functional

$$E_{snake} = \Sigma E_{internal} + \Sigma E_{external}$$

Snakes - Definition

- Internal energy
  - $v(s) = (x(s), y(s))$ where $x(s)$, $y(s)$ are coordinates on snake contour
  - Provides smoothness constraints
  - $E_{internal} = \alpha \mid dv/ds \mid^2 + \beta \mid d^2v/ds^2 \mid$
  - α, β are weights
- External Energy
  - Guides snake towards image features
  - Image gradient is often used

Snakes - Features

- Natural GUI implementation
  - User input defines initial snake
  - Interactive
  - Intuitive
- Deal well with image degradation
  - Broken contour, etc.

Snakes - Features

- External forces subject to smoothness constraints
  - External forces guide snake
  - Internal forces keep it smooth
- Customized external forces for problem domain
  - Statistical segmentation output

Snakes - Current Approach

- User defines initial snake contour
  - Contour includes lesion and is completely external to it
- Snake seeks optimal lesion contour
  - Statistical segmentation output
- Final snake state represents lesion in current slice
- Use final snake state in initial slice to initialize snakes in adjacent slices

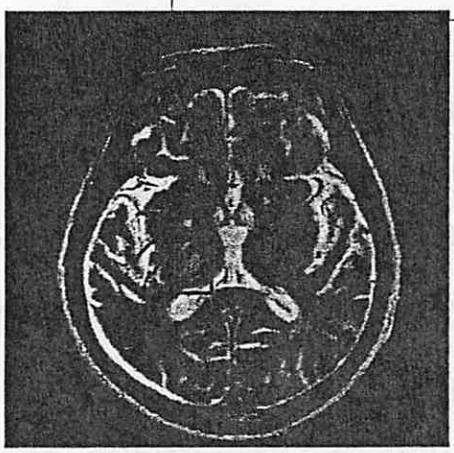

Snakes - Results

Initial Snake

Intermediate Snake

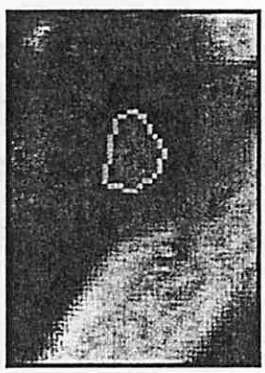

Final Snake

Statistical Segmentation

Given a user-defined region that contains the lesion, break it up into lesion and non-lesion regions

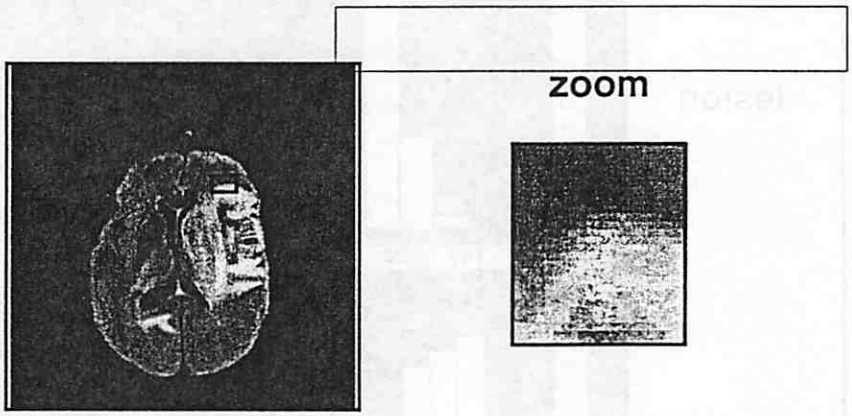

Strategy

- Sample pixel intensities from inside and outside lesion
- Use histograms, Q-Q plots, etc. to model the distributions

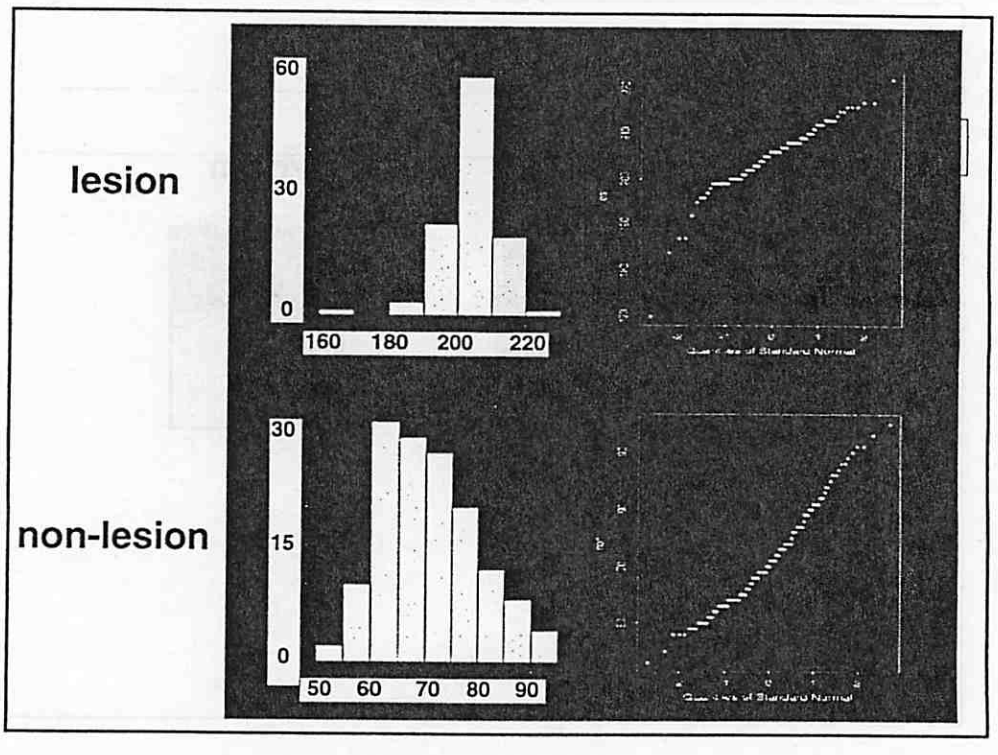

Analyses ➡ GAUSSIAN distributions

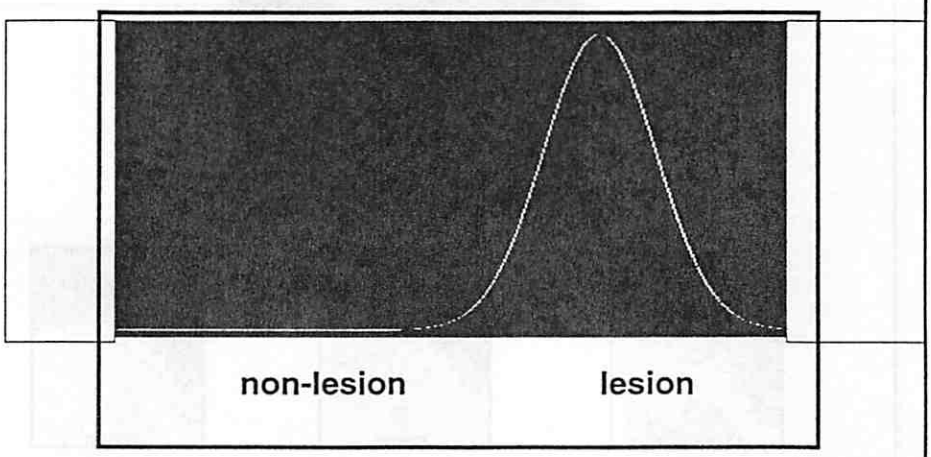

Classification

- Estimate means and standard deviations
- Determine distributions $f_1(x)$ and $f_2(x)$
- If $\frac{f_1(x)}{f_2(x)} < 1$, lesion

  otherwise, non-lesion

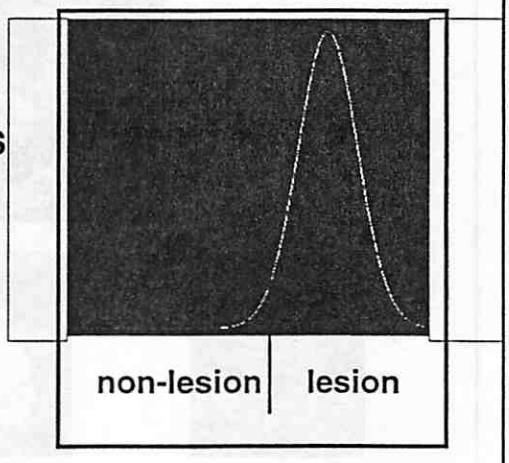

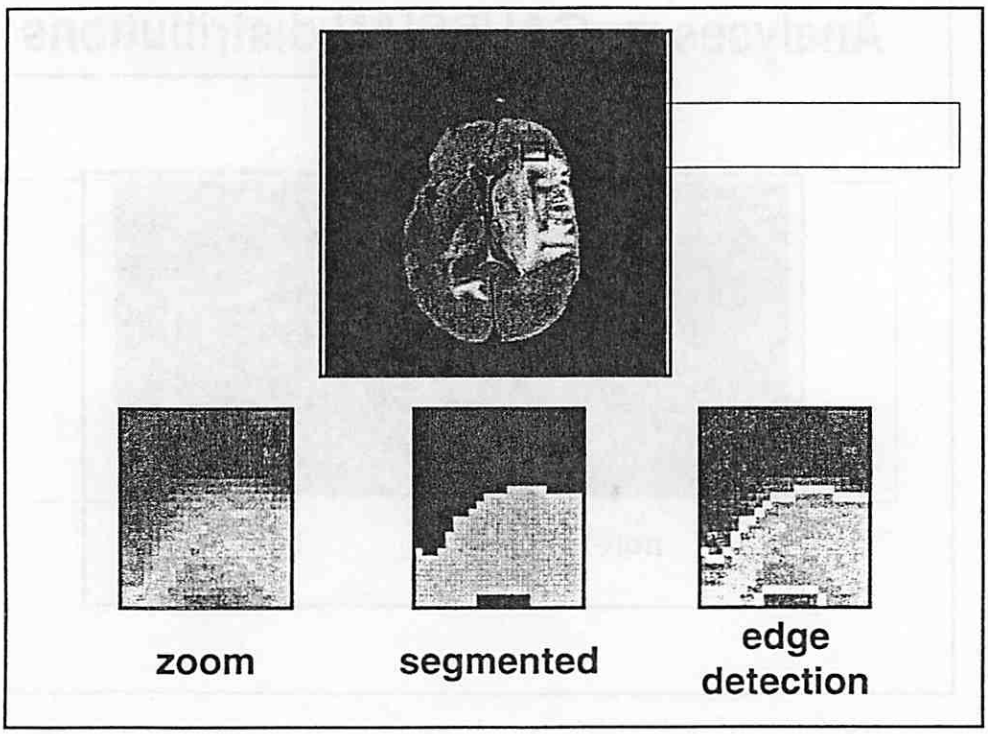
zoom
segmented
edge
detection

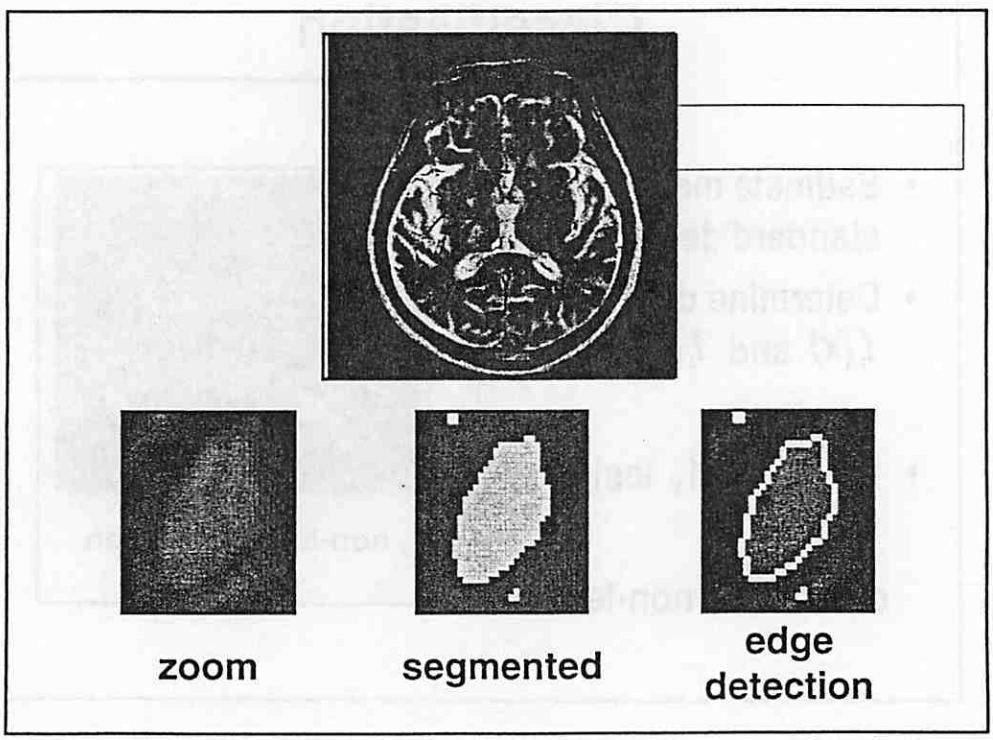
zoom
segmented
edge
detection

Conclusion

- Promising results
- Possible improvement: non-parametric approach
- Will integrate statistical with snake segmentation