

**The Baystate Stroke Project: Getting Started
on Segmentation of Stroke Lesions**

Justus H. Piater

CMPSCI Technical Report 96-48

August 1996

The Baystate Stroke Project: Getting Started on Segmentation of Stroke Lesions

Justus H. Piater

Computer Science Department
Lederle Graduate Research Center
University of Massachusetts
Amherst, MA 01003-4601

Abstract— This report describes the Stroke Lesion Segmentation Project of the UMass Computer Vision Laboratory in cooperation with Baystate Medical Center. Its purpose is to define the current state and next steps, and to familiarize the reader with existing data and tools. After a brief description of the project and its clinical goals, the UMass contributions and possible approaches are presented. Thorough literature reviews are related to our goals. Additional preliminary experiments are discussed. The report concludes with a description of our clinical data and existing tools.

Keywords— Stroke Lesion, Segmentation, Registration, Radio Frequency Correction, Volume Measurement, Markov Random Field

1 Introduction

This research is part of a project led by Dr. Pleet, Head of the Neurology Department, Baystate Medical Center (BMC), Springfield, MA. The main purpose of this project can be outlined as follows: People who have suffered an acute *ischemic stroke*¹ show an elevated blood pressure, the exact source and effects of which are not known. The goal of this clinical study is to determine whether patients should be treated with medication that lowers blood pressure.

¹Terms printed in *italics* are explained in the Glossary section of this project's Web page at URL <http://vis-www.cs.umass.edu/projects/stroke/stroke.html>.

In the course of the project, different blood pressure treatment protocols will systematically be performed on separate groups of patients. The development of brain *lesions* (as imaged by MR and CT methods), and other assessments of recovery, will be correlated with the treatment in order to define an optimal treatment protocol. More specifically, such lesions generally consist of a core of dead tissue (*infarct*), and a surrounding area (*penumbra*) of damaged tissue that either might recover or die. Whether the lesion imaged by MR and CT include none, part or all of the penumbra, usually is not clear. A more complete description of the project is presented on our Web page, and a brief introduction to Magnetic Resonance Imaging is provided by Kapouleas (1990).

The contribution of the UMass Computer Vision Laboratory will be the localization and volumetric measurement of lesions, the main problem being accurate identification of lesions and their boundaries. Difficulties arise from

- general fuzziness of the image data (caused by partial volume effects),
- low contrast between lesion and other tissue types,
- variations in intensity characteristics (intersubject and spatial),

some of which are readily illustrated in Fig. 1.

This work involves the following subtasks, not necessarily in this order:

Registration of volumetric data for

- obtaining multispectral volumes acquired with different parameters (eg., PD, T1, T2 MR and/or CT),
- data refinement by integrating slices acquired in different orientations (*axial, sagittal, coronal*),
- comparison of data acquired at different times, and/or
- relating a given scan to a brain atlas as a segmentation aid.

Segmentation of lesion from healthy brain tissue.

Volume measurement of the lesion.

A successful completion of this project may have a significant impact on the routine treatment of ischemic stroke patients.

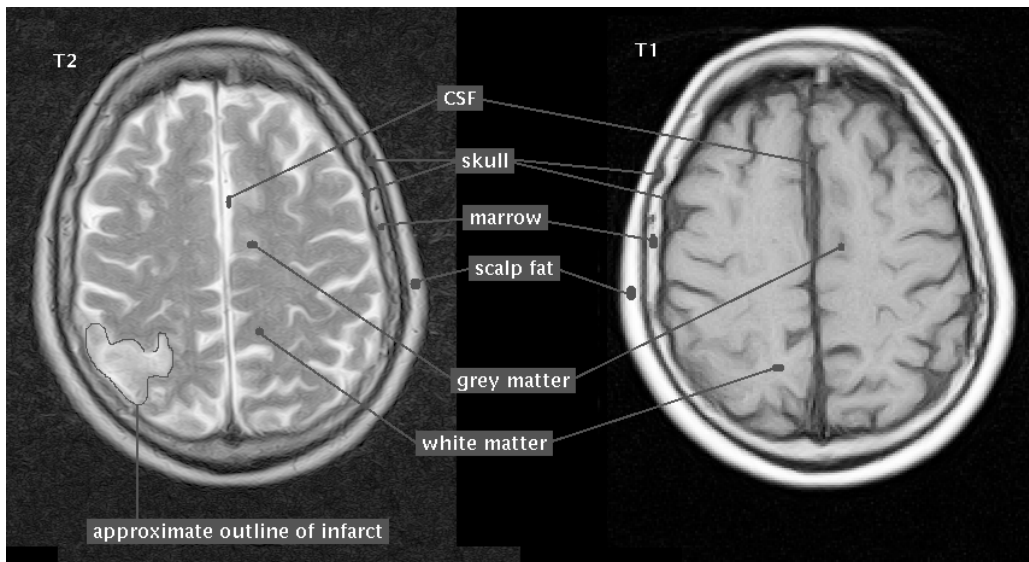


Figure 1: Two annotated sample MR images showing the same brain slice using different parameters. (This image is from the Project’s Web page.)

2 Related Work

We are not aware of any published work on automated stroke lesion segmentation. However, some significant work has been done on Multiple Sclerosis lesion segmentation. Here are a few characteristic examples:

Kapouleas (1990) was one of the first to present a comprehensive approach. He developed a highly specialized system, involving (among other features) a specially designed low-level 2-D segmentation algorithm based on a combination of edge detection, local thresholding, and region merging; and elastic registration of the brain volume to an atlas using the interhemispherical fissure as a landmark, moments of inertia for orientation, and spline surface fitting for elastic deformation. The result is used to aid false positive rejection. He operates on routine clinical data.

More recently, probabilistic segmentation paradigms predominate. Kamber et al. (1995), for instance, developed a brain tissue probability model from high-density MR data of healthy subjects. The model was built in a standardized reference coordinate system proposed by Talairach and Tournoux

(1988). They compare four different approaches to segmentation of MS lesions:

1. nonmodel-based (purely data driven),
2. data driven, but model used to confine the search to plausible white matter locations,
3. model used to provide spatial tissue type probability features in addition to the data features, and
4. model used for both purposes.

They employ four different classifiers, namely

- minimum distance clustering,
- Bayesian maximum-likelihood,
- decision tree, and
- pruned decision tree,

and run all of them on all four tasks mentioned above. In short, the best results were achieved when the model was used only to confine the search, and the Bayesian classifier achieved the most accurate results. The tests were performed on specially acquired high-density axial MR (PD and T2) data.

The most sophisticated approach that I am aware of was presented by Johnston et al. (1996). They use a non-model stochastic relaxation method for segmentation of MR slices obtained in clinical routine. First, all non-brain tissue is automatically masked out, and radio-frequency nonuniformities are corrected by homomorphic filtering (Mackiewicz 1995). Then, the intensity volume is segmented using a Markov Random Field (MRF) model, which initially requires minimal user interaction per data volume. The result is a volume of tissue-type probability vectors. (The MRF approach will be described in more detail below.) Those voxels that now have a lesion probability greater than a threshold are classified as lesion. Finally, a morphological opening operation is performed to reduce the number of isolated lesion voxels.

The problem of stroke lesion segmentation is closely related to this work. However, most of this previous work relies on the fact that MS lesions are predominantly found in the white matter of the brain. Also, it seems that stroke lesions tend to be less regular in shape and have fuzzier boundaries than typical MS lesions (personal observation). This makes both accurate segmentation and false positive rejection somewhat harder.

3 Our Approach

After removing any spatial irregularities from the data (in particular, radio-frequency correction), we propose to register all available data acquired at the same session. Multispectral segmentation will be performed in 3-D to extract any lesions. If successful, the volume of the lesions can then easily be computed. (It should be noted, however, that the accuracy is limited by the coarseness of the clinical data.) Key steps of this procedure will be outlined in more detail in the following sections. Furthermore, by registration of CT data to MRI, we may be able to define what parts of infarct and penumbra are actually imaged by CT and MR methods.

In clinical application, assessment of lesion size changes over time will then be possible by registering scans acquired at different sessions.

3.1 Data Acquisition

Our data is provided to UMass by Baystate Medical Center where it is acquired on a routine basis. A typical MR data set, acquired at a single session with one patient, at least includes one sagittal T1, and one axial T1/PD/T2 scan each. Each scan typically consists of roughly 20 slices of 5mm thickness, the distance between slices being 7mm (i.e. there is a gap of 2mm between slices). Each slice covers an area of 240×240 mm and is sampled into 256×256 (T1) or 512×512 (PD and T2) pixels, respectively. The gray values typically range from 0 to about 1600.

Every scan comes as a set of image files along with a machine generated ASCII parameter file. Perl scripts have been developed that read these files and perform a variety of operations on these files and the associated MR data, including anonymification of the data, sorting them by patients, display

of information, data conversion, reslicing of data volumes, and mapping of metric coordinates to slice and pixel indices.

3.2 Radio-Frequency Correction

The response of today's MRI machines is to some extent nonuniform over the scanned volume. Both inter-planar and intra-planar variations occur. Typically, the intensity of the MRI signal characterizing a particular tissue composition is lower near the borders of the scanned volume. In other words, MR images appear brighter near the center. (For an example and further references see Johnston et al. 1996). Therefore, every approach that relies on global features (e.g., MRF as described below) should start with the removal of intensity irregularities like radio-frequency inhomogeneities.

Since we have access to our MRI machine, we should try to obtain test scans of a so-called phantom and use the results to define a radio-frequency correction function. A phantom in this context is a homogeneous object of appropriate size which is scanned by an MRI machine for calibration purposes. It shows an intensity response similar to the human body. Any radio-frequency inhomogeneities of an MRI machine will show up as intensity variations within a phantom scan. These variations of the calibration scan (and therefore, any similar scan) can be reversed by an appropriately designed digital filter.²

If this turns out to be impossible, we will have to take another approach. Johnston et al. (1996) briefly mention methods proposed by other authors, including low-pass filtering of the original MR data to approximate a phantom and fitting a surface to interactively defined points of the same tissue type. Others suggest to suppress RF inhomogeneities directly by high-pass filtering of the MR data. Johnston et al. themselves employ a method equivalent to homomorphic filtering (Mackiewicz 1995) which we could follow.

3.3 Registration

In this context, there are three classes of registration:

²Apparently, at BMC a phantom is used for regular calibration of their MRI scanner, which may serve our purposes.

- intrasubject MR to MR,
- intrasubject MR to CT, and
- individual MR or CT scans to an electronic brain atlas.

All intrasubject registration tasks can be assumed to require only rigid transformations, while intersubject and scan-to-atlas registration will require elastic (non-rigid and, ideally, non-affine) transformations. Most published related work on MS lesion segmentation does not involve intrasubject registration because only MRI is used, and it is assumed that all data are already registered. In practice, however, this is only true if the different data sets are acquired simultaneously, which is generally not the case with our data.

3.3.1 Previous work on rigid registration

Most of the predominant approaches to rigid registration fall into one of the following categories, roughly in increasing order of complexity and accuracy:

1. Principal Axes Fitting
2. Surface-To-Surface / Point-To-Surface fitting
3. Feature Correlation

Dhawan et al. (1995) give a good example of the first approach. It works by first segmenting the volumes to be registered, creating binary representations of a reference volume (e.g., the brain), assuming it has a mass and computing its principal axes of inertia, which are then matched. It is apparent that the accuracy of the result heavily relies on the quality of the first segmentation. If the chosen reference volumes do not match or are incomplete in the data sets, then the outcome is likely to degrade.

Yan and Karp (1994) at the University of Pennsylvania use the principal axes transform only for the estimation of an initial estimate. Then, they use a Surface-To-Surface fitting approach: The brain surface is approximated by cubic spline surfaces, which can efficiently be matched. This procedure is more robust and quite accurate. However, it too requires segmentation of the brain surface.

Feature Correlation is the approach employed by Kumar et al. (1994). It starts with the extraction of 3-D edge features in both data volumes. These features are then iteratively matched at increasing resolutions to avoid local minima and speed up convergence. This procedure relies on the assumption that a sufficient number of common features occur in both data sets (which is always true in the cases discussed here). It is robust with respect to an additional large number of features that do not match (and thus may produce outlier matches).

The feature correlation approach seems to be the most promising one. However, it apparently has only been tried on data which is a lot denser than that obtained in our clinical routine. On the other hand, van den Elsen et al. (1995) in a previous paper experimented with 5 mm slices and no gaps between slices, and report some success. (Our Baystate data set is expected to usually consist of 5 mm slices plus 2 mm gaps.) This approach was an immediate predecessor of Kumar et al. (1994).

3.3.2 Plans for rigid transformation

It seems that the approach implemented by Kumar et al. (1994) represents the current state of the art, and is appropriate for our needs. A lot of manpower tends to be wasted by researchers at different institutions re-inventing or re-implementing existing solutions. To avoid this, and since there is a research agreement between the University of Massachusetts and the David Sarnoff Research Center, I propose that we make as much use of their existing software as possible.³ If it turns out that their solution does not produce satisfactory results (e.g., because of our thick slices), then I suggest we follow the approach of Yan and Karp (1994).

The performance of all iterative registration algorithms depends on an initial pose estimate. For our data, a good estimate is provided by the orientation coordinates given in the data (cf. Sec. 4.2). If a patient did not change his position during the entire scanning period, these coordinates represent the exact solution to this registration problem.

³On 6/4/96 Rich Weiss reported that Kumar is prepared to give us an executable version of his code.

3.4 Segmentation

Geman and Geman (1984) proposed to model a gray level image as a Markov Random Field. Local energy functions are defined on pixel similarities and edge elements. The potential coefficients characterizing the local characteristics of pixels and lines are given by hand.

I propose that we follow the approach taken by Johnston et al. (1996). They extend Geman and Geman's approach in two ways: (a) they operate on 3-D neighborhoods, and (b) instead of refined intensities, vectors of tissue type probabilities (rather than scalar class labels) are computed. Furthermore, only one type of local interdependency is used, which is derived from the probability of two (not necessarily distinct) tissue types neighboring each other. The vectors of tissue type probabilities are iteratively updated using Besag's (1986) Iterated Conditional Modes (ICM) algorithm. Simply speaking, one iteration of ICM constitutes one voxel-by-voxel sweep of the data volume in which, in this case, each voxel's tissue type vector is changed to maximize their likelihoods, given its observed data value and the current labels of its neighbors. This procedure is repeated until convergence, which is guaranteed if the voxels are swept according to a constant sweep pattern.

3.4.1 The approach of Johnston et al.

The approach is described in detail in Johnston et al. (1994a) and Johnston et al. (1994b).⁴ The following steps are performed for each (monospectral) MR volume:

1. First, the user is to interactively outline image regions which he believes are pure samples of one tissue type. At least one region is required per tissue type θ to be considered (white matter, gray matter, CSF, lesion).
2. Normalized gray level histograms H_θ are computed for each of these tissue types, providing initial probability estimates that a given gray level represents each tissue type.
3. A volume V is constructed consisting of vectors giving these probabilities for each voxel.

⁴It should be noted that in their terminology, the term "model-based" refers to the MRF model, whereas I used it referring to a brain atlas.

4. Based on the probabilities given by V , a maximum likelihood segmentation of the volume is performed. For the MRF model, the neighborhood interaction parameters β of the tissue types are taken as the inverse of the probability that two voxels of the given tissue types are adjacent.
5. One iteration of the stochastic relaxation process (ICM) is performed. Prior probabilities are given by the H_θ . This step yields an improved version of V .
6. If a given convergence criterion is met, terminate.
7. Recompute histograms H_θ from V . This bootstrapping step relies on the assumption that the assigned probabilities in V have become more accurate.
8. Go back to step 5, or, if refinement of the β is desired, to step 4.

To combine the volumes V obtained from the PD and T2 data sets, all corresponding probabilities are multiplied together (under the assumption of independence, which the authors justify by the independent acquisition of the PD and T2 data). Then, the vectors are renormalized to yield $\sum p = 1$, and a maximum likelihood segmentation of the data volume is performed.

It should be noted that it is unlikely that the construction of the β and the combination of probabilities can be formally justified. They must be regarded as reasonable heuristics. Obtaining the energy coefficients β is an important area of research.

3.4.2 Modifications and Comments

I suggest the following modifications:

- Rather than segment each data volume separately and then integrate the resulting probability vectors, I suggest we operate on one V only. This means that in steps 2 and 7 above, the histograms are computed for each of the volumes as before, but the resulting probability vectors are combined immediately in steps 3 and 5. Thus, all available information is used right from the start. I expect improvement in the reliability

of the segmentation result, because some tissue types cannot be differentiated in a monospectral image (e.g., lesion and CSF in T2), whereas this seems possible if all images are combined into a multispectral data set.

Johnston et al. noted that when they attempted this, they experienced serious performance problems due to the huge amount of data to be held in memory at the same time. I suggest that the local nature of the MRF model be exploited in such a way that at most three slices per volume are held in memory at any one time. This probably requires that sagittal and coronal data be axially resliced.

- While Johnston et al. use one PD and one T2 axial data sets only while preserving the resolution, the MRF model provides a natural way to integrate all available data at any orientation, modality, and resolution. Since we have various data sets at different orientations and resolutions, we should take advantage of that. If slices at different orientations are integrated, an appropriate definition of the energy coefficients effectively yields a refinement of the data to higher accuracy than is provided by scans at one orientation alone.
- If it turns out that a brain atlas is needed, then the prior tissue type probabilities given by the atlas can be included analogously in steps 3 and 5.

Johnston et al. exploit the highly parallel nature of the ICM algorithm by using a parallel architecture. Although they state that the procedure converges fast and can usually be terminated after 5–8 iterations, they do not indicate the order of magnitude of time required to process a data set.

The MRF/ICM algorithm seems a natural choice because it is relatively general, and most of the specific parameters can be derived from the input data and the desired segmentation objective. This is more desirable than a highly specialized approach, as proposed e.g. by Kapouleas (1990).

3.5 Volume Measurement and Visualization

Once the lesions have been identified in all slices, their volumes can trivially be estimated by counting voxels. This assumes that most of the voxels

involved are classified as highly pure lesion. It is not yet clear how significant the percentage of inhomogeneous voxels near the boundary of the lesions is. It might be necessary to employ mixture-models (as used in e.g. remote sensory classification of ground coverage type), especially for smaller lesions, to increase accuracy. Another possibility is the approximation of the lesion surface by a closed spline surface and to integrate the volume. The volumes and locations can then be compared between data acquired on different dates. Experiments will show whether our relatively coarse data will yield sufficiently accurate results. As Michael Scudder pointed out (personal communication), the goal of the project is to assess *changes* in volumes as opposed to absolute volumes. This implies that consistent, systematic errors in volume measurement can be tolerated.

For visualization, it might be helpful to interpolate between slices to increase the spatial resolution. This is nontrivial because a simple intensity interpolation of neighboring slices would just yield blurred images. Raya and Udupa (1990) proposed an interesting approach for shape-based interpolation of objects. It requires a slicewise binary representation of the object to be interpolated. First, a slicewise gray level representation is computed where the gray level of each pixel is given by the distance of that pixel to the closest boundary point of that object within the same slice. The resulting gray level is positive if the pixel belongs to the object, negative if not, and zero if it lies on the boundary. These gray level slices are then interpolated at the desired resolution. Thresholding the volume at zero yields a binary representation of the interpolated object voxels.

4 Preliminary Experiments

Some preliminary experiments were performed which may or may not play a role in future progress. They will briefly be discussed in the following sections.

4.1 Contrast Enhancement

While the intensity range extends over roughly 1600 gray levels, the interesting contrast is contained within very small gray scale ranges. This suggests the use of specially designed gray scale mappings both for visualization and

automatic segmentation. While this alone is obviously not sufficient for segmentation, in general it can be expected to aid a segmentation algorithm by focusing on the relevant gray level range. However, we do not expect this information to be useful in our approach since this information is implicitly incorporated in our proposed MRF model.

4.1.1 Global Thresholding

If a PD image is scaled uniformly over its entire intensity range, the resultant contrast is very low, and a lesion is hardly visible (Fig. 2). However, if an appropriate intensity window of a narrower range (say, 100 gray levels) is selected and mapped to the visible range, a lesion can be greatly enhanced for human inspection (Fig. 3).

In T2, lesion contrast from healthy brain tissue can be enhanced, making the boundary between lesion and healthy brain obvious. On the other hand, this is not the case for the boundary between lesion and CSF, since they are generally indistinguishable in T2 MRI. In T1, lesions can hardly be identified at all.

CT data contains very little detail if viewed at uniform scaling. A lot of detail is hidden in the lower intensity ranges (Fig. 3). On the other hand, now the bony structures are smeared out and cover parts of the brain near the skull.

As mentioned above, such intensity transformations are desirable for human viewing and can be helpful for many other segmentation algorithms, but they are probably useless in the context of our proposed MRF approach.

4.1.2 Image Arithmetic

Other researchers sometimes compute the difference between PD and T2 images to enhance certain features. This is related to multispectral processing of satellite data. Multiple bands can be combined in innumerable linear and nonlinear ways. An example is shown in Fig. 4. Experiments like this may be combined with gray scale remapping to enhance certain features. However, it will probably not be helpful in our case since it is unclear if and how the visibility of stroke lesions can be enhanced by such methods. A multispectral




Figure 2: PD, T2 and CT images at uniform gray level scaling.




Figure 3: PD, T2 and CT images at highly narrowed gray level scaling. The MR images have been inverted, causing the lesions to appear dark. In the top left PD image, one larger and one small lesion are clearly visible in the left hemisphere of the brain. These can also be identified in the T2 image on the right, but here they are indistinguishable from CSF. In the bottom CT image, some brain structures are visible, but the white areas representing bony structures are smeared, hiding part of the brain. Lesions are not clearly visible.

MRF approach will use all available data anyhow, rendering image arithmetic useless for segmentation.

4.2 Coordinate System Identification

For defining the position of a slice, a right-handed coordinate system is used with the origin roughly in the center of the head. The x axis extends to the right ear, the y axis down towards the neck, and the z axis forward towards the front of the head. This is the coordinate system used to identify the position and orientation of image slices in our data, and provides a good initial pose estimate for registration purposes.

This coordinate system is consistent with our parameter files. It was determined by inspection of angular values in our parameter files, interpreting them as Euler angles. It should be noted, however, that it may not be the definition that the designers of our parameter set had in mind.

If the patient does not move at all between scans, all data are perfectly registered with respect to this coordinate frame. Unfortunately, this is generally not the case. However, examples exist where the data seem well registered, as illustrated in Fig. 5. Here, the location at (10,20,30) millimeters with respect to this coordinate system was arbitrarily chosen as an example. From the data available in the parameter file, slices and pixel locations were computed that represent this point in the axial and sagittal scans. To show the registration of axial and sagittal scans, the appropriate sagittal slice was extracted from the axially scanned data volume. The location at (10,20,30) mm is marked and seems to correspond well in the two sagittal views.

4.3 MR Parameter Quantization

There does not seem to be a clear definition of the PD, T1, and T2 MR weightings. They are determined by the repetition time (TR), spin-echo time (TE) and inversion time (TI) parameters. Roughly, T1 is acquired using short TR and TE, and T2 from long TR and TE. If a proton density (PD) image is desired, TR is long and TE short. In all of these cases, TI is zero. A positive TI can be used to cancel out a certain tissue type. TR and TE are usually set as in T2, and such an acquired scan is called Flair.




Figure 4: Feature enhancement by image arithmetic. The first two images show PD and T2 weighted data, and the third one the difference between these two. Note how the boundary between brain and skull is slightly enhanced as compared to both raw images.




Figure 5: Localization of location (10,20,30) in 3-D millimeter coordinates in axial (top left), sagittally resliced axial (top right) and sagittal (bottom) data sets. Slice and pixel indices were computed from the coordinate information contained in the parameter files. Note that the top right image was produced from the same data volume as the top left one, but in a different orientation. By visual inspection, the sagittal data set seems to be well registered to the axial one.

Finally, pilot scans are performed using very short TR and TE. For automatic classification of our data, these heuristics were turned into crisp rules that are consistent with our data (Fig. 6).

<i>if</i>	TI	>	0			<i>then</i>	Flair
<i>else if</i>	TR	<	300			<i>then</i>	Pilot
<i>else if</i>	TR	<	800	<i>and</i>	TE	<	20
						<i>then</i>	T1
<i>else if</i>	TR	>	1500	<i>and</i>	TE	>	70
						<i>then</i>	T2
<i>else if</i>	TR	>	1500	<i>and</i>	TE	<	40
						<i>then</i>	PD

Figure 6: Mapping from MR parameters to weighting labels. Inversion (TI), Repetition (TR) and Echo (TE) times are given in milliseconds. The right-most column gives the label name for the respective data set.

5 Research Goals

This work has advanced to a stage where the steps described above can be formally defined, filling in the details, and implemented. Our goals include:

- Automatic detection and segmentation of stroke lesions.
- Modifications to Kumar et al.’s registration algorithm as necessary.
- Derivation of energy coefficients in an MRF segmentation model for segmentation of stroke lesions.
- 3-D multispectral relaxation segmentation.
- Integration of data from thick imaging slices at different orientations for resolution enhancement.
- Achievement of sufficiently accurate segmentation and volumetric measurements using routine (i.e., coarse) clinical data.
- Introduction of an automated or semi-automated system for stroke lesion detection and volumetric measurement into a clinical routine procedure.

A successful implementation and clinical evaluation may have a significant impact on the routine treatment of ischemic stroke patients.

Acknowledgements

Even though I did most of the work described in this paper, not all of the ideas expressed here are mine. In particular, I would like to thank Michael Scudder and Prof. Richard Weiss for their invaluable discussions, and Prof. Edward Riseman for suggestions on the presentation of this research. Thanks also to Dr. Bernard Pleet and Dr. Richard Hicks from Baystate Medical Center for their friendly and helpful cooperation.

References

- Besag, J. (1986). On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society* 48(3), 259–302.
- Dhawan, A. P., L. K. Arata, A. V. Levy, and J. Mantil (1995). Iterative principal axes registration method for analysis of MR-PET brain images. *IEEE Transactions on Biomedical Engineering* 42(11), 1079–1087.
- Elsen, P. A. van den, J. B. A. Maintz, E.-J. D. Pol, and M. A. Viergever (1995). Automatic registration of CT and MR brain images using correlation of geometrical features. *IEEE Transactions on Medical Imaging* 14(2), 384–396.
- Geman, S. and D. Geman (1984). Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6(6), 721–741.
- Johnston, B., M. S. Atkins, and K. S. Booth (1994a). Partial volume segmentation in 3-D of lesions and tissues in magnetic resonance images. In *Proc. SPIE – Medical Imaging*, Volume 2167, Bellingham, WA, pp. 28–39. SPIE – The International Society for Optical Engineering.
- Johnston, B., M. S. Atkins, and K. S. Booth (1994b). Three-dimensional partial volume brain tissue segmentation of multispectral magnetic resonance images using stochastic relaxation. In *Proc. SPIE – Non-Linear Image Processing V*, Volume 2180, Bellingham, WA, pp. 268–279. SPIE – The International Society for Optical Engineering.
- Johnston, B., M. S. Atkins, B. Mackiewich, and M. Anderson (1996). Segmentation of multiple sclerosis lesions in intensity corrected multispectral MRI. *IEEE Transactions on Medical Imaging* 15(2), 154–169.

- Kamber, M., R. Shinghal, D. L. Collins, G. S. Francis, and A. C. Evans (1995). Model-based 3-D segmentation of multiple sclerosis lesions in magnetic resonance brain images. *IEEE Transactions on Medical Imaging* 14(3), 442–453.
- Kapouleas, I. (1990). *Model Based Interpretation of Magnetic Resonance Human Brain Scans*. Ph. D. thesis, Rutgers Univ., Dept. of Computer Science, New Brunswick, NJ.
- Kumar, R., K. Dana, P. Anandan, N. Okamoto, J. Bergen, P. Hemler, T. Sumanaweera, P. A. van den Elsen, and J. Adler (1994). Frameless registration of MR and CT 3D volumetric data sets. In *Proceedings of the Second IEEE Workshop on Applications of Computer Vision*, Los Alamitos, Calif., pp. 240–248. IEEE: IEEE Computer Society Press.
- Mackiewich, B. (1995). Intracranial boundary detection and radio frequency correction in magnetic resonance images. Master’s thesis, Simon Fraser Univ., Computer Science Dept., Burnaby, British Columbia.
- Raya, S. P. and J. K. Udupa (1990). Shape-based interpolation of multidimensional objects. *IEEE Transactions on Medical Imaging* 9(1), 32–42.
- Talairach, J. and P. Tournoux (1988). *Co-Planar Stereotaxic Atlas of the Human Brain: 3-Dimensional Proportional System—An Approach to Cerebral Imaging*. New York: Thieme Med.
- Yan, M. X. H. and J. S. Karp (1994). Image registration of MR and PET based on surface matching and principal axes fitting. In *Nuclear Science Symposium and Medical Imaging Conference*, Volume 4, New York, NY, pp. 1677–1681. IEEE.