Change Detection
in
MRI Brain Scan Data

J. O. Nicholas Woodward

Master of Science
Department of Artificial Intelligence
University of Edinburgh
1990

Abstract

This dissertation describes research into change detection in MRI brain scan data using a point-to-point matching approach. The changes occur between two images of the same patient taken at different times. Problems to be overcome included noise, grey scale shifts, image registration and computational complexity.

Matching every pixel in one image to every pixel in the second image provides a vector mapping. From this mapping the image transformation can be calculated and removed to reveal the changes that have occurred in the brain. Such an approach assumes that images can be consistently imaged (or distortion removed), that the changes can be represented by the motion of pixels and that the changes are localised and small relative to the total brain volume.

The solution removes noise by conservative smoothing, matches pixels using local normalised correlation, smooths initial matches by maintaining neighbourhood trends and removes the image transformation by statistical minimisation. Results are shown throughout demonstrating the performance of the technique and there is an analysis of the computational complexity.

Acknowledgements

Throughout this project Bob Fisher and Martin Connell, my supervisors, have provided invaluable support and encouragement whenever needed. Without their willingness to discuss the project at all times I would not have achieved what I did. Special thanks to them.

Thanks also to Ann Ralls for loading the MRI images, Martin Waite for his assistance with X Windows and Douglas Howie, without whom there would be no photographs.

Lastly, my flatmates Sue and Emma deserve thanks for keeping me fed, watered and sane at crucial moments during the project.

Funding for this project was provided by the Science and Engineering Research Council of Great Britain (SERC), Award No. 890408345.

Table of Contents

List of Figures

Chapter 1

Introduction

In this chapter the problem, "Change Detection in Magnetic Resonance Imaging (MRI) Brain Scan Data", is introduced and the need for such a technique, along with potential problems, is discussed briefly. The approach taken in developing the technique and a summary of the results and subsequent conclusions are also presented.

In the remainder of the dissertation related work is outlined in Chapter 2 and the fundamentals of the technique and algorithms used in the implementation are covered in Chapters 3-6. Throughout those chapters an example is followed through to demonstrate the various stages of the change detection. In Chapter 7 the results of tests carried out on the change detection are described. Chapter 8 contains an analysis of the computational complexity of the problem. In Chapter 9 the performance of the change detection is discussed, the conclusions on the project are presented and possibilities for further work are suggested.

Appendix A contains information on the equipment used during the project and explains how the displays (photographs of which appear in the dissertation) were produced. In Appendix B the "C" source code used to implement the algorithms is documented.

1.1 Interpretation Of Medical Images

Medical imaging covers many different data capturing techniques; MRI, X ray, Positron Emission Tomography (PET) and photography are among these techniques. To obtain 3 dimensional data the patient is usually imaged in "slices" and these 2D slices can then be stacked up to form a 3D data set. For structural investigations of the brain MRI data is used since it shows more structural detail than other techniques. Fig 1–1 shows a typical MRI slice of a human brain. The colour is artificial. Higher pixel values show up in a brighter green. The background, has simply been thresholded and assigned a light grey colour for the purposes of clarity on the photographs. This accounts for the light grey patches between the skull and the brain. The change detection takes unthresholded images as input.

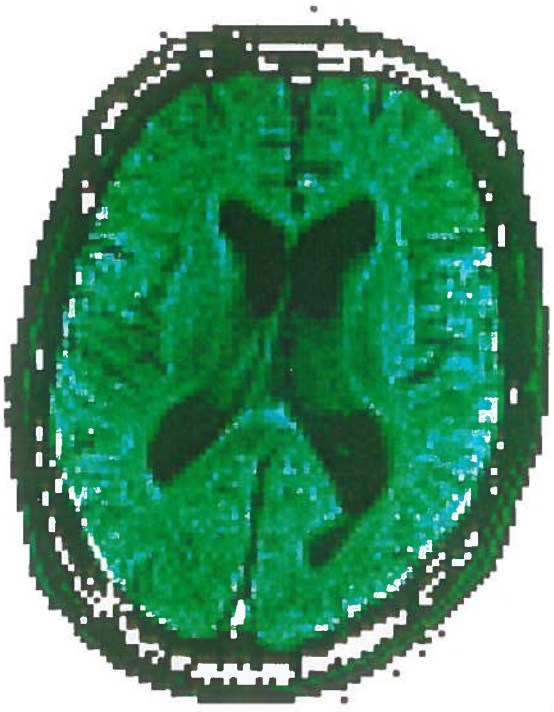

Figure 1–1: A typical MRI slice of the human brain.

The aim of this project was to investigate a technique for detecting structural changes in the brain by interpreting MRI images. At the present time, except in case of obvious gross structural degeneration, clinicians make no attempt to detect such structural changes; the task is too difficult.

These difficulties are a result of two problems. Firstly, in both the 2D and 3D cases, the images are unlikely to be registered. That is, there will effectively be a translation and rotation of the patient between the two images, making it harder for a clinician to diagnose any changes which might be indicative of either treatment success or disease progression. The automatic registration of images is the subject of much study at the present time; several methods are discussed in Chapter 2. There is an additional complication caused by the lack of registration when imaging in slices. Because the head will be in different positions within the scanner in different images taken at different times, the slices will not pass through the same planes in the head. Therefore 3D images cannot be interpreted slice by slice.

Secondly, in the case where 3D images are used, the clinician's ability to interpret 3D data from 2D slices is severely taxed and is a time-consuming process. As Apicella *et al* [Apicella 89] point out , "imaging matching has proven to be a difficult task for a human operator" and Wells *et al* [Wells 89] suggest that "with the current rapid increase in medical image data and the rising demands made on a clinician's time there is a strong case for developing automatic methods of recognising normal and abnormal regions in medical images." So any tools for automatically interpreting image data will be of use to the medical community.

Current methods of change detection, effectively measurements of change in cerebro-spinal fluid volume, are rather crude. A more detailed understanding of which areas of the brain are changing and by how much would be useful for diseases such as Alzheimer's, Creuzfeldt-Jakob and AIDS which cause structural degeneration of the brain.

For instance: AIDS compromises both the central nervous system (CNS) and the immune system. The drug AZT appears to alleviate the problems the disease causes to both these systems. However, it may be that, as understanding of the disease is improved, sufferers' immune systems can be restored, but they go on to develop dementia due to degeneration of the CNS. A technique to enable a detailed study of this degeneration would therefore be of benefit in understanding such diseases.

Based on discussions between members of the Departments of Artificial Intelligence (DAI), Medical Physics (DMP) and Medical Radiology (DMR) it was

felt that a technique relying on segmenting the brain into regions to develop a brain "map" and then comparing these regions with some standardised normal map was not the correct approach since at present there is not sufficient knowledge to indicate which areas of the brain change and by how much. Different diseases affect different parts of the brain in different ways so a system relying on particular knowledge might fail to detect some unexpected changes. For this reason the problem was seen as being similar to stereo matching; there are two or more images and a correspondence between feature points was required. From the correspondences, regions where significant changes had occurred could be isolated.

It was believed that the data from a hi-resolution MRI scanner would allow this type of analysis and the goal of the project was to investigate techniques for detecting the changes between two sets of images based on point-to-point matching.

1.2 Problems To Be Overcome

In addition to image registration, other problems which need to be accounted for when developing a change detection technique include; the imaging system, motion of the brain within the skull and the time required for computation.

The Imaging System

As with all imaging systems an MRI scanner suffers from noise, calibration errors and image distortion.

Noise can never be fully eradicated, but processing steps can be taken to limit its effects.

Image registration error is effectively a calibration error as it is a result of not aligning the body and the scanner's coordinate axes between the two images. There may also be shifts in the grey scale between images. For example, the same area of brain might have different pixel values, regardless of noise, in the two images. These grey scale shifts are non-linear and there are no standard techniques for compensating for them. Thus, the change detection technique will

have to match features on their structural content rather than simply on their grey scale values.

Image distortion is caused by ageing of the electrical and mechanical components within an MRI scanner. The electric and magnetic fields, within which the patient is imaged, alter over time and non-linear geometrical changes will occur in images over this time. For example, a circular object may not appear circular. This distortion will vary from machine to machine, but a study yet to be published by Saeed *et al* [Saeed 89] suggests that these distortions are much less significant than might be expected. For the purposes of this project it was assumed that images could be preprocessed to remove image distortion.

The ability to overcome the effects of noise and grey scale shifts was considered to be the most significant factor in determining the success or failure of the change detection algorithm.

The Brain and the Skull

The brain is not rigidly attached within the skull and so can move around to a small extent. Even if two images were perfectly registered, the skull was perfectly aligned between the two images, the brain would probably be in slightly different positions. This brain motion does not constitute a change of any importance and so must be eliminated when attempting to detect real structural changes within the brain.

Computational Complexity

Since a full data set from an MRI scanner was expected to be at least $128 \times 128 \times 128$ pixels and a correspondence between every point was required, a possible $10^{13000000}$ permutations exist. Hence the speed with which an output could be achieved was important. Stereo matching by microcanonical annealing shows some of the problems of computational complexity and one solution to them [Barnard 87]. That paper describes how two stereo problems, one with a problem space 10^{36000} times greater than the other, could be solved in approximately the same time by applying a hierarchical search technique, similar to abstraction, to the greater problem

space. With the existence of such techniques and the increasing use of parallel computing, the problem of computational complexity was given due consideration, but was not considered to be of major significance during investigations of the change detection technique.

1.3 Summary Of Design

Design Methodology

As has been mentioned, the data on which the technique was to be used was dense. It would be impractical, if not impossible, to analyse such data by hand so that results could be produced against which any change detection technique could be judged. Therefore it was important at each stage of the development to define the input, transformation, assumptions and expected output for that stage. The performance of the technique could then be assessed by comparing its output against the expected output with any discrepancies being discussed in terms of the input, transformation and assumptions. This approach to the design of a computational task is effectively that of Marr [Marr 82] [Hallam] as presented in his computational theory.

Design Outline

The change detection technique which was developed divided into five distinct stages; preprocessing, initial mapping, smoothing, transformation extraction and analysis of residual data.

The preprocessing constituted approximate image registration, noise reduction by non-linear smoothing and feature enhancement by measuring the localised pixel value standard deviation. The initial mapping indicated which pixel p_2 in an image $\mathbf{I}_2$ was most like a pixel p_1 in image $\mathbf{I}_1$. This was achieved by taking a neighbourhood, $\mathbf{R}_1$, around p_1 and finding the pixel, p_2 in $\mathbf{I}_2$ whose surrounding region $\mathbf{R}_2$ had the highest correlation with $\mathbf{R}_1$. The smoothing was to reduce the effect of bad matches in the initial mapping. The result of these three operations was a vector mapping indicating where pixels in $\mathbf{I}_1$ mapped to in $\mathbf{I}_2$. A statistical

analysis of this mapping allowed extraction of the image transformation parameters which were used to remove the effects of image transformation from the vector mapping. The residual vectors represented the changes which had been detected between the two images. No attempt was made to remove the effects of motion of the brain within the skull and residual data was only analysed qualitatively.

1.4 Summary Of Results And Conclusions

Although the project was intended to investigate 3D change detection, a lack of 3D data and the time taken to process 3D images restricted testing to 2D data. However, all the algorithms described in Chapters 4-6 extend naturally to 3D.

During the course of the project testing was carried out at three levels. Each individual stage was assessed to see that it produced the expected results, the change detection as a whole was tested on manufactured data and trials were carried out on real, although not entirely suitable data. The trials on the real data could only be assessed qualitatively, but the other tests demonstrated some weaknesses in the change detection technique. Fig 1–2 shows typical output derived from the change detection. The red lines, perceived as yellow on the green of the image, indicate where changes have been detected, the direction in which they occur and how far a pixel has been judged to move.

Using such output the performance of the change detection was investigated and assessed. The strengths and weaknesses of the change detection are outlined below.

The initial mapping, implemented by local normalised correlation, overcomes grey scale shifts between images successfully and the majority of matches are correct. However, in regions of low contrast the matches can be poor and when two images are rotationally misaligned, groups of pixels which lie on or near edges may suffer bad matches.

The smoothing process works as expected and does remove spurious matches from the initial mapping. Unfortunately its dynamic characteristics are not well understood so the control one has over the smoothing process is rather limited, basically to trial and error.

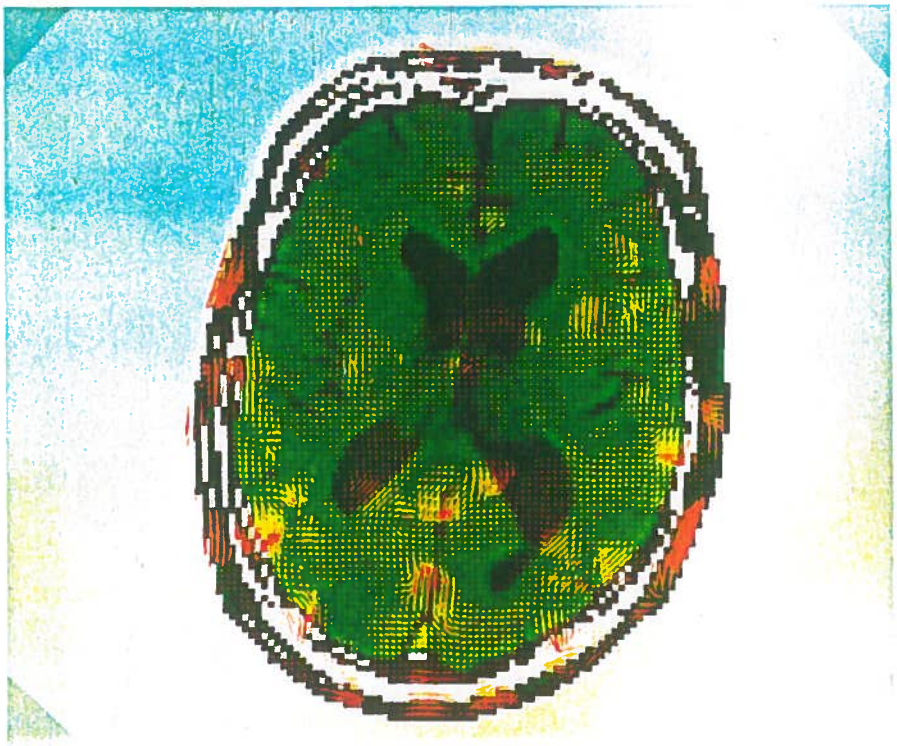

Figure 1–2: A typical display derived from the change detection.

Removal of the image transformation by statistical means is successful. As expected it relies on the initial mapping being accurate and as the initial mapping breaks down at large rotations, the accuracy of the estimation of the transformation parameters decreases.

Tests on real data, although only analysed qualitatively, were encouraging.

Overall the change detection displays the type of behaviour required. The difficulties appear to lie with the initial mapping. Groups of bad matches in the initial mapping can propagate through the smoothing and affect the accuracy of the calculation of the image transformation parameters. Even if the image transformation is calculated correctly, the fact that poor matches have not been removed means they will show up as changes, which is not desirable. The structure of the change detection appears to be satisfactory and for practical purposes could readily be used on a parallel machine to increase the speed of operation.

To conclude, a method of detecting changes between two MRI images based on point-to-point matching has been investigated and the results are encouraging. The overall structure of the change detection algorithm produces the required behaviour, but more work on the initial mapping and computational complexity is needed if such a technique is to be pursued further.

Chapter 2

Review of Associated Work

Attempting to detect changes on a point-to-point basis is a new approach and so there is no literature available of direct relevance. However, there is much literature on the registration of medical images with many different techniques proposed. In this chapter four papers which take different approaches to image registration are reviewed. A fifth paper suggesting that correlation is "inappropriate for the registration of dissimilar images"[Herbin 89] is also outlined. Two further papers, one concerned with matching a medical image to model data and one proposing a signal-symbol approach to change detection are discussed as possible alternative approaches to change detection, other than the point-to-point one investigated in this dissertation.

2.1 Global Image Registration

Image registration attempts to estimate the rotation, translation and magnification between two images so that they can be aligned. There are many different techniques for registering medical images, but generally they take one of three approaches; using properties of the whole image, using properties of features and using statistical methods. Papers advocating each these approaches are described below.

Properties of the Image

Apicella *et al* [Apicella 89] and Junck *et al* [Junck 90] both described image registration techniques using correlation methods. By correlating one set of data against another the highest value of correlation is expected when the data are aligned.

Junck *et al* detailed how one image, $\mathbf{I}_1$, was rotated and translated over certain ranges and correlated with another image, $\mathbf{I}_2$. The translation and rotation which yielded the largest correlation value were the parameters required to register $\mathbf{I}_1$ with $\mathbf{I}_2$.

The approach of Apicella *et al* was to take the fast fourier transform of both images first, before correlating the power spectrums of the images. This had the effect of decoupling the translation, rotation and magnification. The parameters could then be found individually and therefore in a faster time.

Correia [Correia 90] suggested that such methods "use all the available data and hence show good stability. A limitation of this approach, however, is the implicit assumption that both images have the same structure." That is, the structures show no gross changes. This assumption also applies to the local normalised correlation method used in finding an initial mapping between images, as described in Chapter 4.

The work of Junck *et al* was described in 2 dimensions, as was that of Apicella *et al*, but the latter suggested that their work should be extended to 3 dimensions.

Properties of Features

Rather than taking a whole image and trying to match it against another one, certain features within an image can be extracted and matches made on those features.

Saeed *et al* [Saeed 89] described how certain features could be extracted from MRI images and how the transformation parameters could then be calculated. Features such as the brain area and inner scalp boundary were identified by image processing algorithms developed specifically for such purposes. Two methods of determining the rotation between two images, based on the features extracted, were outlined. The first technique used the contour of the extracted feature and the second dealt with the mass of the feature. By coding the contours of features as a complex waveform the waveforms could be cross correlated and where a peak in correlation occurred, the rotation between the two images was indicated. The mass of the features could be used to find the line of minimum inertia, one of the principle axes of moments. The angle between two such lines of minimum inertia was the angle of rotation between the two images. These methods relied on good feature extraction, but it was stated that the extraction algorithms had been "shown to be robust over a wide range of data." A more serious criticism [Apicella 89] suggested that the "method of moments is satisfactory for normal anatomy and physiology, but its results deteriorate when applied to abnormal cases that most require reliable diagnosis."

The above work was developed and intended for use in 2 dimensions only.

Statistical Methods

Pelizzari *et al* [Pelizzari 89] described a statistical method of registering images. The procedure was to find a series of contours in one image which defined some distinct surface, usually the external skin of the head. From the second image a series of points were taken (it was not made clear how these were chosen). Finding a match between the points and the surface was described as being analogous to fitting a hat (the points) to a head (the surface). This was achieved by a least squares minimisation between the points and the surface described by the contours.

An assumption of this method was that there exists a unique transformation between the two images, which highly symmetrical objects would not give. "Fortunately such perfect symmetry , however, seems sufficiently lacking in human heads" [Pelizzari 89]. Such calculations "tend to have some sensitivity to errors in the selection of the edge points." [Correia 90]

The above work was developed for 3 dimensions and a similar least squares fit method was used to extract image transformations as explained in Chapter 6 of this dissertation.

Criticism of Correlation

Herbin *et al* [Herbin 89] argued that correlation is unsuitable for the registration of dissimilar images. They proposed as alternatives various mathematical modelling similarity measures. The unsuitability of correlation is to be expected if the images are drastically dissimilar since that would negate the underlying assumption of correlation. Given that the point-to-point matching, described in Chapter 4, attempts to correlate regions, it may be that in areas of the brain where changes have occurred the regions being matched would not be sufficiently similar for correlation to be reliable. If this were to be the case similarity measures such as those described by Herbin *et al* could be investigated. For the work described in this dissertation it can be assumed that similarity holds at large scale between images and changes are localised and small relative to the unchanged portions of the brain.

2.2 Change Detection; Alternative Approaches

Elastic Matching

Although not specifically developed for change analysis the work of Bajcsy and others [Bajcsy 83][Bajcsy 89][Dann 88] used an elastic matching technique to match patient data to an idealised brain atlas, a method which might lend itself to change analysis. It is this technique that allows a system of computer programs

"to automatically locate, measure and describe anatomical features of interest with accuracy and consistency."[Bajcsy 83]

A series of steps allowed the creation of a computerised brain atlas from the brain of a normal, but deceased, male. The brain was sliced into sections and stained to enhance subcortical structures and grey/white interfaces. The slices were then photographed and digitised. By manually tracing contours around various structures on these digitised images, using a mouse, the brain atlas was created.

This brain atlas could then be elastically matched to fit another patient being imaged, and thus create an individualised anatomy atlas. This process was described as being "analogous to using a standard atlas of contours which had been drawn onto a sheet of elastic rubber. Such a rubber sheet may be deformed or stretched so that the atlas and the image match each other point for point"[Bajcsy 83]. The process of multi-resolution elastic matching is detailed in [Bajcsy 89].

It would not seem unreasonable that a technique such as this would allow an analysis of the change between two images of a patient taken at different times by comparing the brain atlases produced.

In concluding, Dann and Bajcsy *et al* [Dann 88] stated that "the current findings were promising" comparing favourably to experts analysis of image slices. However, they suggested that while it was not yet time "to turn the task of data analysis over entirely to the computer, .. results do suggest that this kind of system may be very useful to enhance analysis, display, interpretation and communication of brain images."

Although taking a different approach from point-to-point matching the work of Bajcsy *et al* and this project both made matches based on local information only and the motivation for using a multi-resolution approach was that it reduced the computational complexity of the problem, as the hierarchical approach suggested in Chapter 8 would.

A Signal-Symbol Approach

Lee *et al* [Lee 86] described how a signal-symbol approach could be used for detecting significant changes in images obtained by aerial photography. They suggested

that previous efforts to automate change detection had focussed on implementations in either the signal or the symbolic domain, but not both. By taking a dual approach a signal based change detection algorithm could pass on information to a symbol based change interpreter. The signal was the raw image and the symbols were tokens containing slots indicating properties of the regions cued by the change interpreter. Properties such as size, shape and area were used.

The change detection algorithm was based on a linear prediction model which used small patches from a reference aerial photograph to locally model the corresponding areas in a second image and vice versa. The change interpreter contained a set of "physical cause frames" which attempted to determine whether the changes were physically significant or not. Non-significant changes included cloud cover, shadows, parallax effects and partial occlusion. Changes for which no physical cause could be found were passed on to an image analyst for manual inspection.

The work outlined above makes no reference to image registration, nor does it attempt to show how changes have occurred (e.g. a certain region is moving in a particular direction), but the change detection algorithm which makes matches on small image patches is not unlike the approach taken in this project. The main difference is the use of a change interpreter to cue significant changes. The algorithm described in this dissertation implicitly assumes that all changes which are signalled, after the image transformations have been removed, are significant.

In this chapter several techniques of image registration have been reviewed and attention has been paid to some criticism of correlation. The alternative approaches to change detection have similarities to the work of this project in that all methods rely on matching small patches of images. The work described in this dissertation differs from that outlined above in that changes are detected by finding a mapping for all pixels in one image into a second image. On the basis of this point-to-point mapping and the assumption of global and local consistency of shape, image transformations are removed to reveal real structural changes that have occurred between the two images.

Chapter 3

Outline of the Change Detection

This chapter outlines the change detection technique and explains why it was structured as it is. Chapters 4-6 discuss in detail the algorithms, associated mathematical considerations and limitations on performance.

3.1 Outline

Input/Output

The change detection technique developed takes as its input two images, $\mathbf{I}_1$ $\mathbf{I}_2$, approximately registered, taken of the same subject time t apart. The output is a list of vectors indicating where pixels in $\mathbf{I}_1$, representing part of the brain, have moved to in $\mathbf{I}_2$. No vector indicates that no change has occurred. In order to be successful, this relies on being able to accurately remove the effects of image transformations between $\mathbf{I}_1$ and $\mathbf{I}_2$ and on finding reasonable matches for all the pixels in $\mathbf{I}_1$.

Assumptions

A change detection technique based on point-to-point matching which also has to overcome image registration problems has three underlying assumptions.

An image of the brain is only a digitised representation, with pixel values representing some property of a small region of the brain being imaged. (Usually

a function of water molecules in the brain in MRI data.) Therefore if structural changes are to be detected in the brain, the imaging process must accurately and consistently reflect the structure of the brain. Assumption one:

A1) The imaging process consistently and accurately reflects the structure of the brain.

This explains why it was stated in Section 1.2 that images would require image distortion to be removed by preprocessing. Pixels can now be considered to represent a particular part of the brain and changes within the structure of the image necessarily imply changes in the structure of the brain.

Consider two images $\mathbf{I_1}$ and $\mathbf{I_2}$ of the same brain, taken at times t_1 and t_2 respectively. If $\mathbf{I_2}$ is rotated, translated and magnified with respect to $\mathbf{I_1}$ then every pixel, p_{1i} at point (x_{1i}, y_{1i}) in $\mathbf{I_1}$ will have moved to some new position, (x_{2i}, y_{2i}), in $\mathbf{I_2}$. If, in addition, there has been a change of the shape of the brain between t_1 and t_2 then the pixel p_{2i} which matches p_{1i} will be located at $(x_{2i} + d_x, y_{2i} + d_y)$, where (d_x, d_y) represents the change that has occurred. This can be expressed in vector notation as

$$\underline{\mathbf{P}}_{\mathbf{2i}} = \underline{\mathbf{P}}_{\mathbf{1i}} + \underline{\mathbf{T}}_{\mathbf{i}} + \underline{\mathbf{D}}_{\mathbf{i}} \tag{3.1}$$

Where $\underline{\mathbf{P}}_{\mathbf{xi}}$ represent image points and $\underline{\mathbf{T}}_{\mathbf{i}}$ and $\underline{\mathbf{D}}_{\mathbf{i}}$ the effects of transformation and change respectively. Eq. 3.1 can be rearranged to give the change:

$$\underline{\mathbf{D}}_{\mathbf{i}} = \underline{\mathbf{P}}_{\mathbf{2i}} - \underline{\mathbf{P}}_{\mathbf{1i}} - \underline{\mathbf{T}}_{\mathbf{i}} \tag{3.2}$$

This is the second assumption:

A2) Changes in the structure of the brain can be represented as motion of the parts of the brain and hence motion of pixels within an image.

After smoothing of the initial mapping $\underline{\mathbf{P}}_{\mathbf{1i}}$ and $\underline{\mathbf{P}}_{\mathbf{2i}}$ are known, so to calculate $\underline{\mathbf{D}}_{\mathbf{i}}$, $\underline{\mathbf{T}}_{\mathbf{i}}$ must be found. If it is assumed that:

A3) Changes are small compared with the total brain volume.

Then a measure of the overall motion of the pixels, perturbed in places by changes, will allow an estimation of the image transformation parameters; rotation, magnification and translation.

In addition, the ability of point-to-point matching to represent changes should be questioned. If a region expands then one pixel in the first image must map to several in the second image. Conversely if a region contracts then many pixels in the first image should map to one pixel in the second image. To represent expansions and contractions the matching should account for both these one to many and many to one matches. Because the brain is an enclosed object, if one part expands another must contract so both types of change should be representable. However, given the computational complexity of the problem (See Chapter 8), searching for more than one match per pixel was felt to be impractical, so only contraction can be properly represented. A possible way to overcome this is suggested in Chapter 9.

The three assumptions stated above apply whatever algorithms are used in the change detection.

Stages of Processing

Five stages of processing were identified which would enable a vector map of changes to be produced. These were: preprocessing, initial mapping, smoothing of the initial mapping, transformation extraction and analysis of residual data. The first four stages were fully implemented and the residual data was extracted for visual qualitative analysis, but a full analysis was not developed since there was not sufficient data available. In the following section the requirements of these stages are stated. It should be noted that at this level of discussion the details of the change detection are unimportant. The five stages described should be sufficient to specify a change detection system. The performance of the algorithms chosen to execute the five stages of the change detection and the resultant implementations should be discussed at the algorithmic and implementational levels.

Having split the change detection into stages the requirements of these stages can now be discussed.

3.2 Processing Requirements

The individual requirements of the five processing stages are outlined below.

Preprocessing

The preprocessing is expected to register the images approximately, reduce noise and enhance features which are likely to match well.

There are many suitable methods of image registration, both interactive and automatic. Approximate registration was achieved using the HIPS and ISIS interactive image processing tools, but solving this problem was not directly part of the project.

There are many noise reduction techniques available, but some of these techniques remove data which are essential for the later stages of the processing. For instance, if the initial mapping relies on sharp features and the image noise is reduced with a low pass filter, which removes sharp image features, this makes the noise reduction and initial mapping incompatible. Whatever noise reduction technique is chosen must be compatible with the initial mapping and so must preserve local feature structure. In this project conservative smoothing was used.

Feature enhancement is similar to noise reduction in that it brings data out above the noise. How this feature enhancement is used depends on the later stages of processing. A measure of the local standard deviation around a pixel was used here to increase the reliability of the smoothing stage of the processing.

Initial Mapping

If a pixel, p_1 in $\mathbf{I}_1$, is to be found a match in $\mathbf{I}_2$ how should this be done? Searching for a pixel with the same grey scale value is insufficient for three reasons; noise perturbs pixel values, imaging causes shifts in the grey scales and there may be several pixels with the same value. To overcome the grey scale shifts and possibility of similar pixels the structure of a neighbourhood around a pixel needs to be taken into account. So requirement one is:

R1) A matching function based on local structure can overcome grey scale shifts which may occur between images.

A normalised correlation function, as explained in Chapter 4, offers the required behaviour.

Smoothing of the Initial Mapping

As a result of the effects of noise, image transformations, grey scale shifts and changes within the brain not all pixels in $\mathbf{I}_1$ can be expected to be matched to the correct pixel in $\mathbf{I}_2$. To overcome this a smoothing function is needed to identify spurious matches in the initial mapping and correct them appropriately. It is important that the initial matching has a large proportion of good matches otherwise identifying spurious matches will be a difficult task. Another function of the smoothing process is to ensure that the changes and transformation represented by the initial mapping are smooth. The brain is a single natural object so there will be no abrupt changes that occur naturally with time. Two more requirements have now been identified:

R2) The majority of initial matches are accurate so that spurious vectors may be readily identified.

R3) A smoothing function can remove the spurious vectors on the basis of neighbouring matches.

In Chapter 5 the proposed smoothing function is described. This technique aims to remove local discrepancies from the initial mapping while maintaining the trends of neighbouring matches.

Transformation Extraction

The smoothing process produces a vector mapping indicating which pixels in an image, $\mathbf{I}_2$, match pixels in image, $\mathbf{I}_1$. Each of these mappings is the result of the image transformation and any changes which might have occurred. A method of calculating and extracting the image transformation from the vector mappings

would leave vectors indicating the changes that had taken place. A requirement of this stage is that:

R4) There are enough accurate matches to reliably estimate the transformation parameters.

Chapter 6 describes how the transformation parameters can be estimated by a least squares minimisation technique.

Analysis of Residual Data

The analysis of the residual data requires the removal of changes which have been registered as a result of motion of the brain within the skull. The vectors remaining will then identify real changes which have occurred within the structure of the brain. These changes should then be grouped and analysed for their extent, direction, volume and characterised as expansion or contraction.

3.3 Summary

Having outlined a solution to the problem it can be seen that this technique is reliant on three assumptions:

A1) The imaging process consistently and accurately reflects the structure of the brain.

A2) Changes in the structure of the brain can be represented as motion of the parts of the brain and hence motion of pixels within an image.

A3) Changes are small compared with the total brain volume.

Assumptions **1** and **3** stipulate the kinds of data on which the technique should be successful and assumption **2** must underlie a point-to-point matching technique.

For the structure of the proposed solution four requirements have been identified:

R1) A matching function based on local structure can overcome grey scale shifts which may occur between images.

R2) The majority of initial matches are accurate so that spurious vectors may be readily identified.

R3) A smoothing function can remove the spurious vectors on the basis of neighbouring matches.

R4) There are enough accurate matches to reliably estimate the transformation parameters.

Of these, **R2** and **R4** require the initial mapping to be accurate, so it can be expected that the performance of the change detection will rely on the performance of the initial mapping. Additionally, the smoothing function must not affect the accuracy of the initial matches otherwise **R4** will be violated.

In the following chapters it will be shown that the algorithms chosen for the processing stages and the subsequent implementations display the appropriate behaviour. The test results show that there are some difficulties still to be overcome.

Chapter 4

Preprocessing and Initial Mapping

In this chapter it is explained how image noise reduction was achieved by conservative smoothing and how the measure of local standard deviation could be used as a "feature enhancer". In the final section the requirements of the initial mapping are discussed before local correlation matching is presented along with an example of an initial mapping.

4.1 Conservative Smoothing

Conservative smoothing [Fisher] is a method of non-linearly smoothing images. It removes spot noise by looking at the neighbourhood of a pixel and if the pixel value is larger (or smaller) than the neighbouring values the pixel value becomes the maximum (or minimum) of the neighbouring values.

```
forall pixels i in an image
        if p(i) > max(N(p(i)))    p(i)=max(N(p(i)))
        else
        if p(i) < min(N(p(i))) p(i)=min(N(p(i)))
```

Where `p(i)` is the value of pixel `i`, `N(p(i))` defines a set of neighbouring values and `max` and `min` find the maximum and minimum of a set.

The smaller the neighbourhood defined the more pixel values will be smoothed. Fig. 4–1 shows the 3×3 neighbourhood which was used in the conservative smoothing and Fig. 4–2 shows the pixels which were changed in white. The image after

smoothing is not shown because the effects are not easy to identify. (Smoothing of the background is not shown.) Both images input to the change detection were conservatively smoothed before the initial mapping took place.

n1	n2	n3
n8	p	n4
n7	n6	n5

Figure 4–1: A 3x3 neighbourhood.

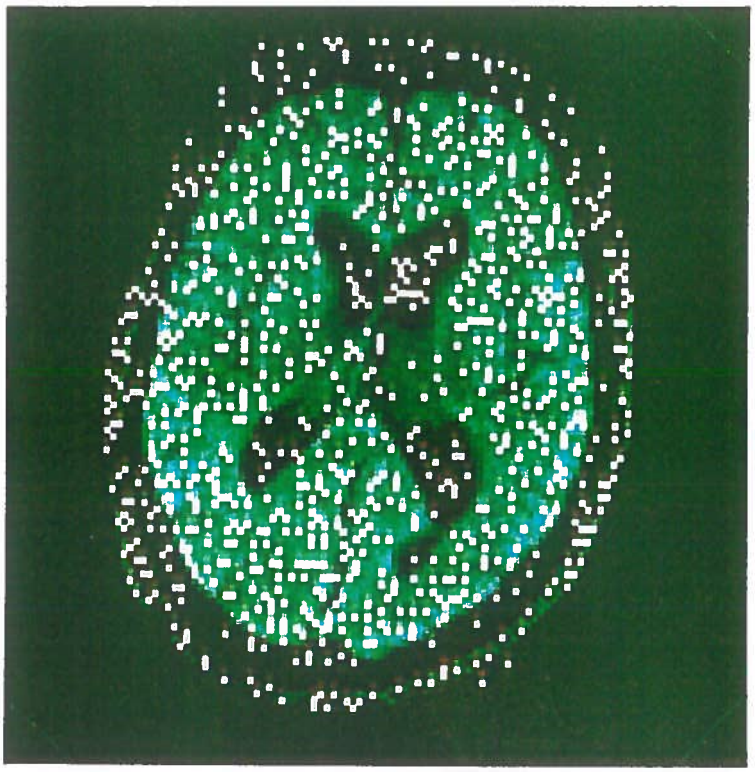

Figure 4–2: Pixels changed by conservative smoothing.

4.2 Feature Enhancement

In the context of this project, features were not considered to be unique objects or parts of objects, but regions of an image which showed structure. For a region to contain structure it must show variation in the pixel values. A measure of this variation is the standard deviation of the pixels over the region.

$$\sigma(\underline{i}) = \left[\frac{\sum_{\underline{k} \in \mathbf{R}} [p(\underline{i} + \underline{k}) - \overline{p}(\mathbf{R}(\underline{i}))]^2}{n(\mathbf{R})} \right]^{\frac{1}{2}} \tag{4.1}$$

$\mathbf{R}$ is a set of vectors defining points in a region, $\mathbf{R}(\underline{i})$ defines a region about a pixel at $\underline{i}$, $p(\underline{i} + \underline{k})$ is the value of a pixel at $(\underline{i} + \underline{k})$, $\overline{p}(\mathbf{R}(\underline{i}))$ is the average pixel value of the region centred on a pixel at $\underline{i}$ and $n(\mathbf{R})$ is the number of pixels in the region.

Fig. 4–3 shows how the local standard deviation varies over the image of Fig. 1–1. It can be seen that the edges of the skull, brain, ventricles (dark lobes in the centre of the brain) and some other features have been crudely enhanced. Pixels with raw data values less than 20 were not used and are shown in white. The local standard deviation was used as an input to the smoothing stage as part of a measure of reliability of a match between two pixels. The region over which the standard deviation was measured was a 5×5 square with the corners removed (See Fig. 4–5). This region was the same shape as the region used to obtain correlation values for matching pixels and the reason for the shape is explained in Section 4.3.2.

Fig. 4–4 shows the distribution of local standard deviations for Fig. 4–3. The large peak is the background local standard deviation, values between $\approx 3 - 10$ are those of the inner brain and above ≈ 10 are the skull, brain and ventricle edges.

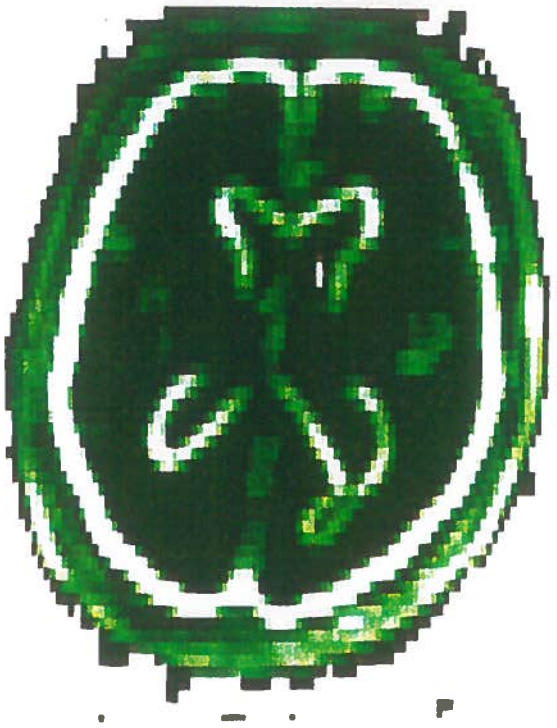

Figure 4–3: Feature Enhancement: Brighter = Higher Standard Deviation.

4.3 Producing An Initial Mapping

The first, and most important stage of the change detection, is to find an initial mapping between the two images being investigated. This initial mapping indicates where every pixel in the first image, $\mathbf{I}_1$, has a "most similar" pixel in the second image, $\mathbf{I}_2$. What constitutes a "most similar" pixel depends on the method used to test for similarity. In this section the mathematical considerations of trying to match two images are explained and the normalised correlation function is presented as the similarity test. The performance of the initial mapping is discussed with reference to an example.

4.3.1 Mathematical Considerations

Consider trying to find a match, region R_2 in image $\mathbf{I}_2$, for region R_1 in image $\mathbf{I}_1$. If the best match is considered to be the region "most similar" according to

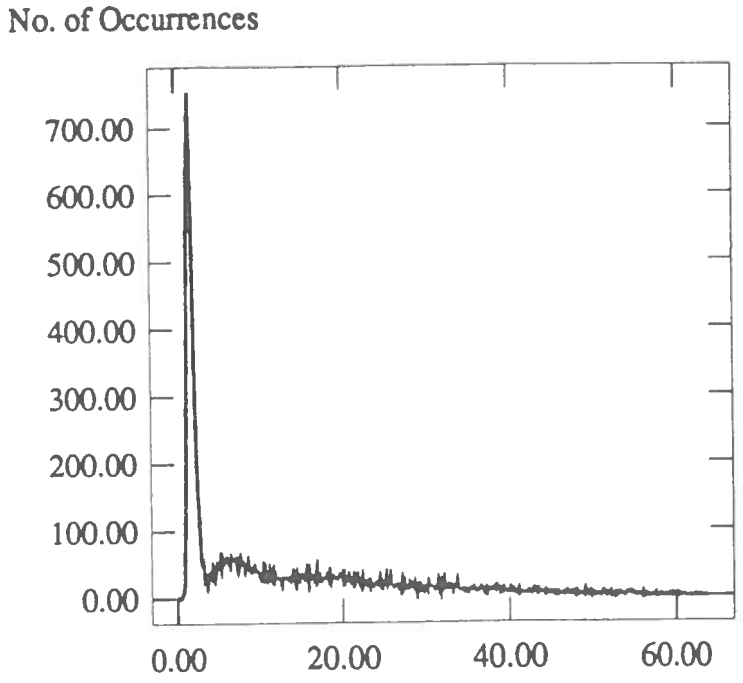

Figure 4–4: Histogram of Local Standard Deviation.

some function, where in $\mathbf{I}_2$ should R_2 be sought and what transformation must R_1 undergo if the correct match is to be found?

If it is assumed that the only operations that can be applied to I_1 to produce I_2 are magnification, m, rotation, $\mathbf{R}$, and translation, (x_t, y_t), then there are six different orderings of these operations. RTM,RMT,TRM,TMR,MRT.MTR. Additionally, the centres of magnification, (x_m, y_m), and rotation, (x_c, y_c), must be taken into account. Grey-level scaling is not taken into account here since the normalised correlation compensates for grey scale shifts between images.

Consider a rotation, translation and magnification applied to $\mathbf{I}_1$ to produce $\mathbf{I}_2$ by studying a point (x_1, y_1) and its match (x_2, y_2).

$$(x_2, y_2) = m[\mathbf{R}(x_1 - x_c, y_1 - y_c) + (x_t, y_t) - (x_m, y_m)] \tag{4.2}$$

Since rotation and magnification are linear operators [1] Eq. 4.2 can be rewritten

[1] $f((0,0)) = (0,0)$ and $f(\lambda(a+b)) \equiv \lambda f(a) + \lambda f(b)$

as:

$$(x_2, y_2) = m\mathbf{R}(x_1, y_1) + (x_{new}, y_{new}) \tag{4.3}$$

$$(x_{new}, y_{new}) = m[-\mathbf{R}(x_c, y_c) + (x_t, y_t) - (x_m, y_m)] \tag{4.4}$$

Similarly it can be shown for the other five orderings that the resultant transformation is equivalent to a magnification and rotation centred on (0,0) with the translation being a function of the ordering, translation, magnification, rotation and the centres of magnification and rotation. How does this affect an attempt to match R_1 to R_2?

R_1, being a portion of an image can be treated like an image. So if $\mathbf{I}_1$ undergoes a transformation $m\mathbf{R} + (x_{new}, y_{new})$ to become $\mathbf{I}_2$ then R_1 must undergo the same transformation to become R_2. Furthermore, it can be shown that if (x_r, y_r) is the centre of region R_1 this can be incorporated as part of the translation. Thus all regions in $\mathbf{I}_1$ will undergo the same magnification and rotation, but the amount of translation will depend on the position of the centre of the region from (0,0) and the values of the other transformation parameters. Therefore to find a match for R_1, R_1 must be rotated, translated and magnified (over certain ranges) and the values of these parameters at the best match, R_2 in $\mathbf{I}_2$, will indicate the transformation between the two images.

4.3.2 Local Normalised Correlation

Of the many methods of matching images or image regions intensity based normalised correlation [Ballard 82], which has been used successfully for both image registration and stereopsis, seems particularly suitable for this application. It has a history of successful use on images and the measure of similarity is based on the structure of an image or image region.

If a region R_1 centred at $\underline{i} = (x_1, y_1)$ is correlated with a region R_2 centred at $\underline{j} = (x_2, y_2)$ the value of the correlation is given by Eq. 4.5:

$$c(\underline{i}, \underline{j}) = \frac{\sum_{\underline{k} \in \mathbf{R}} [p(\underline{i} + \underline{k}) - \overline{p}(\mathbf{R}(\underline{i}))] \times [p(\underline{j} + \underline{k}) - \overline{p}(\mathbf{R}(\underline{j}))]}{\left[\sum_{\underline{k} \in \mathbf{R}} [p(\underline{i} + \underline{k}) - \overline{p}(\mathbf{R}(\underline{i}))]^2 \times \sum_{\underline{k} \in \mathbf{R}} [p(\underline{j} + \underline{k}) - \overline{p}(\mathbf{R}(\underline{j}))]^2 \right]^{\frac{1}{2}}} \tag{4.5}$$

Where the symbols of Eq. 4.5 have the same meanings as in Section 4.2.

This correlation is sensitive to the signal-to-noise content of the two image regions, $\mathbf{R}(\underline{i})$ and $\mathbf{R}(\underline{j})$, and high uncorrelated noise decreases the value of the correlation. For normalised correlation to be meaningful the standard deviation of the regions to be matched should be significantly greater than that of the noise. Fig. 4–4 shows that the local standard deviation of the background (effectively only noise) lies between $\approx$ 0-3 whereas the local standard deviation of the brain data lies between $\approx$ 3-65, so the normalised correlation is valid.

Before an initial mapping can be produced a correlation region, $\mathbf{R}$, needs to be defined and functions to magnify and rotate the region are required. If the noise and signal characteristics are known, then the region size may be optimised using that information and certain simple statistical arguments. However, such optimisation neglects the effects of non-statistical, systematic error such as rotation, magnification and image distortion. The statistical errors grow with the region size and can swamp the statistical advantages of large regions. Earlier it was stated that to find a match for, R_1, R_1 must to be rotated, magnified and translated. Having to rotate and magnify a region over certain ranges increases the size of the search space compared with translation alone, because more matches are tried. Since increases in the search space are undesirable, with the appropriate choice of region size is it be possible to make a valid match by only translating R_1?

If images $\mathbf{I}_1$ and $\mathbf{I}_2$ are approximately registered, then the magnification and rotation between them will be small. It follows that a region will only have to be rotated and magnified slightly to find a match. By making the correlation region small the effects of magnification and rotation are minimal compared with the size of a pixel and so can be ignored. The search for a match for R_1 then reduces to translation only. The information on the magnification and rotation are not lost, for those parameters partly determine the translation each region has to undergo to find a match. It should be remembered from Section 4.3.1 that the translation a region undergoes is a function of all the translation parameters, not just the overall image translation. By reducing the search to translation it is assumed that the rotation and magnification will become small systematic errors in the normalised correlation, which therefore requires that a correlation region is small. There is an additional restriction on the size of the correlation region if it is to be

used to detect changes. Any changes that occur will only be small (Assumption 3, Section 3.1) and so the region size to detect these changes will have to be small. Larger regions will only detect larger scale changes in an image. For the above reasons the correlation region, **R**, was defined as shown in Fig. 4–5. It was felt that this region size would be large enough to provide reasonable correlation values without suffering from the effects of magnification and rotation and would be small enough to detect the expected changes.The corners were removed to make the region more rotationally symmetric than a square. The same region was used to calculate the local standard deviation since this was used in conjunction with the correlation value as an input to the regularisation.

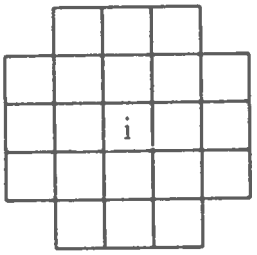

Figure 4–5: The correlation region.

Fig. 4–6 shows the values of correlation[2] produced by trying to match a pixel from the centre of the image in Fig. 1–1 with pixels in a region ± 12 pixels in the x and y directions in the image of Fig. 7–9. It can be seen that although there is a central peak, which indicates the best match, there are other peaks. It had been hoped that the correlation would produce a more distinct match than Fig. 4–6 suggests. It may be postulated that when extending this problem to 3D a more distinct match will be found since the matches will be made on a larger 3D correlation region. (81 points in 3D as opposed to the 21 of Fig. 4–5 in 2D)

---

[2]Throughout Chapters 4-6 the statistical data is taken from a match between two images taken of the same patient at different times, real data. This real data match is discussed in Chapter 7. The examples pictured in Chapters 4-6 are of tests on images transformed by rotating, magnifying and translating a real image, to produce manufactured data, since these show more clearly the operation of the algorithms.

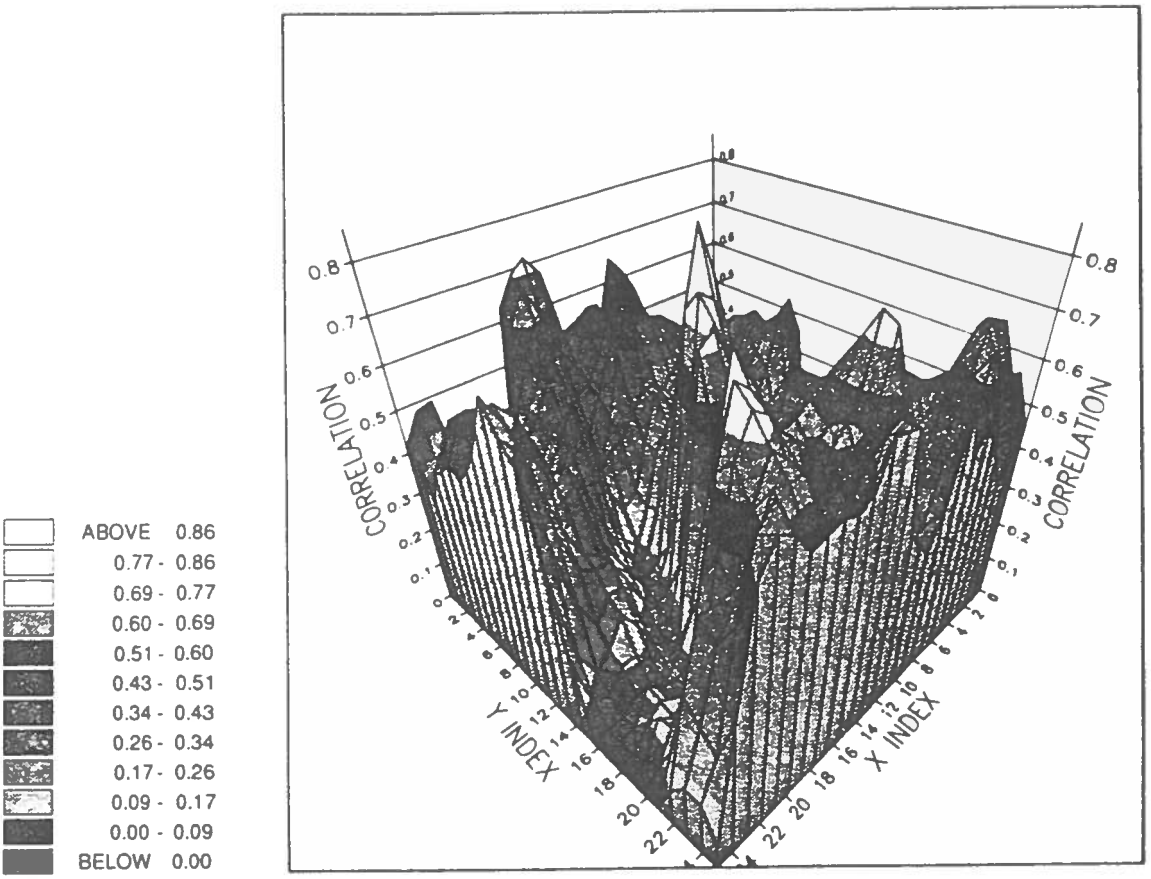

Figure 4–6: Correlation values over a 25x25 region.

When trying to find a match R_2 for R_1 it has now been established that R_1 need only be translated until the best match is found. This is equivalent to a search across image $\mathbf{I}_2$ for the best match. To limit the search to an area less than the whole image a search region, $\mathbf{S}$, was defined. If a match is to be found for a pixel $\underline{i}$ in $\mathbf{I}_1$ then the correlation region around $\underline{i}$, $\mathbf{R}(\underline{i})$, is correlated with all the regions $\mathbf{R}(\underline{i}+\underline{k})$, $\underline{k} \in \mathbf{S}$ in $\mathbf{I}_2$. The region $\mathbf{R}(\underline{i}+\underline{k}_m)$ which produces the highest correlation value is taken to indicate that the pixel at $\underline{i}$, matches the pixel at $\underline{i}+\underline{k}_m$ in $\mathbf{I}_2$.

Writing this in a more structured form gives an algorithm for producing an initial mapping.

```
Search region = S
Correlation region = R

forall pixels, i, in image 1
{
        take the region R(i)
        forall pixels, j, in the region S(i) in image 2
        {
                take the region R(j)
                correlate R(i),R(j) to give c(i,j)
                when c(i,j) is a maximum i matches j
        }
}
```

The theoretical aspects of local normalised correlation matching have been discussed, all that remains is to see how the algorithm performs in practice.

4.3.3 Example

Fig. 4–7 shows the initial mapping after matching Fig. 1–1 with a rotated (-0.05 radians) version of itself. The red lines (seen as yellow on the green background) are vectors mapping a pixel from the initial to the second image. (The "blocking" effect, is an artifact of the rotation algorithm. Small rotations of each pixel in a block lead to the same quantised vector in the mapping.) Since the two images matched were identical (apart from the transformation) the initial mapping would be expected to look something like a "whirlpool" about the centre. That is, all mappings would lie on the circumference of concentric circles.

There are three important characteristics of Fig. 4–7:

1) Most of the mappings are correct.

2) There are some spurious matches (single wayward lines).

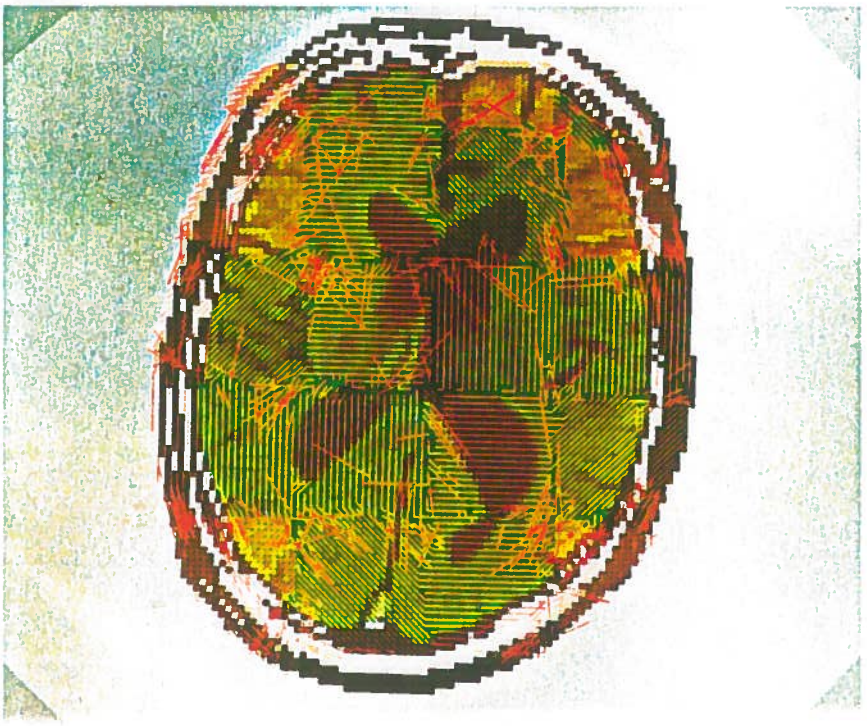

Figure 4–7: An Initial Mapping.

3) There are occasional groups of bad matches.

The fact that the regularity appears is encouraging and shows that the overall conditions of the change detection are being met, namely; requirement **2** of Section 3.2, that the majority of the initial matches are accurate.

The regularity also supports the choice of local normalised correlation to make the initial mapping:

> Most of the initial matches are correct so the local normalised correlation works as a measure of similarity.
>
> The decision to search by translation alone is vindicated since good matches have been found despite there being a small rotation between the matched images.
>
> Similar matches are detected in local neighbourhoods.

The spurious matches were to be expected and it is the role of the smoothing of the initial mapping to remove them.

The reasons for the groups of bad matches are not immediately apparent. At the top left hand side of the bottom right hand ventricle there is a group of 5 or 6 lines which are pointing in the wrong direction (down and to the right). There are other such groups in the image, but they do not show up clearly on the photograph. It was observed that these groups of bad matches tended to lie along image edges. A simple explanation for this could be that the rotation algorithm produced peculiar image artifacts. Another reason could be that after rotation certain portions of curved edges have ambiguous matches. For instance, how would a portion of the circumference of the circle be matched against another portion? There would be many ambiguous matches. This might demonstrate a weakness in searching for matches without rotating regions.

Whatever the reasons, as a consequence of the bad matches occurring on edges the matches are not removed by the smoothing process. Exactly why this occurs is explained in Section 5.2.

What it is important to note here is that the initial mapping produces predominantly the expected correct mapping and some spurious matches, but also some groups of bad matches when images are rotationally misaligned. The consequences of this are discussed and possible solutions are suggested in Chapter 9.

Fig. 4–8 shows the values of correlation when the image of Fig. 1–1 was matched with the image of Fig. 7–9. This suggests that the values of correlation are uniformly high across the brain and this is confirmed by Fig. 4–9, which shows the distribution of the correlation values in Fig. 4–8. The broad peak at ≈ 0.7 contains the background correlation values and the sharp peak at ≈ 1.0 contains the matches in the brain. The sharpness of this peak indicates that the value of the correlation over the area of the brain has a very limited range $\approx 0.9 - 1.0$ and may not be an ideal measure of how reliable matches are. This is of relevance to the smoothing of the initial mapping described in Chapter 5.

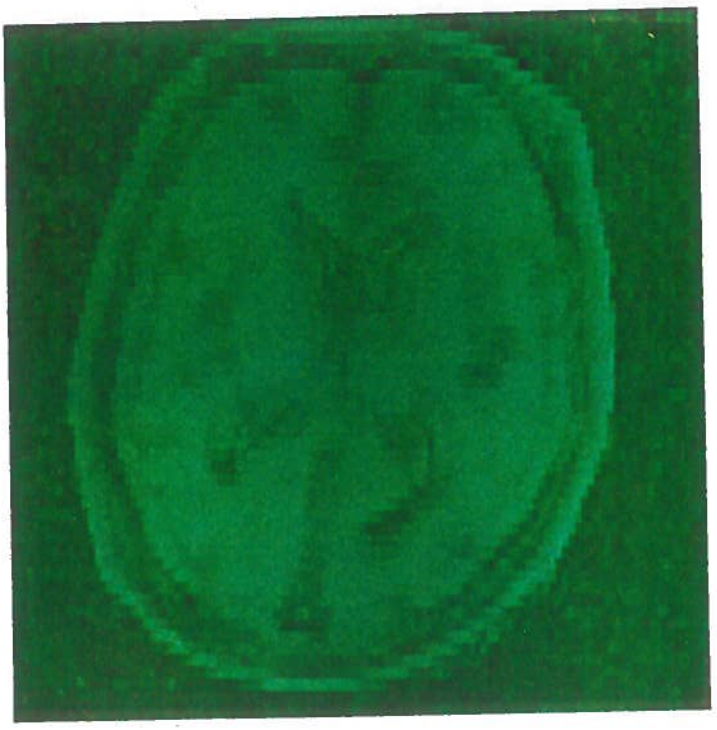

Figure 4–8: Correlation Values: Brighter = Higher.

4.4 Summary

In this chapter conservative smoothing and feature enhancement have been explained. The process of initial mapping by using local normalised correlation matching has been presented and its performance discussed and shown to meet the requirements of the change detection. The following chapter explains how the initial mapping can be smoothed to remove the effects of spurious matches.

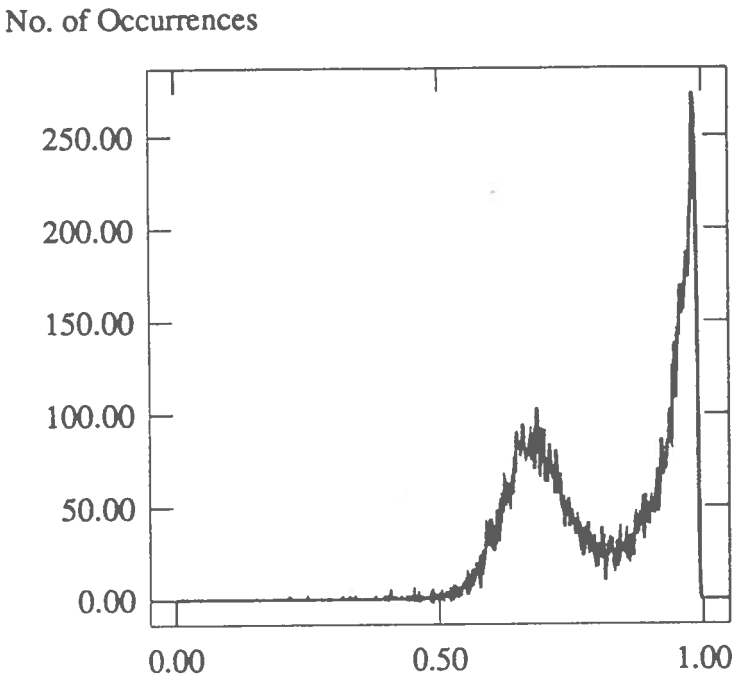

Figure 4–9: Histogram of Correlation Values.

Chapter 5

Smoothing of the Initial Mapping

This chapter details the requirements of a smoothing function and explains the algorithm which was developed to meet these requirements. An example and series of test results are presented to demonstrate the behaviour of the smoothing process.

5.1 Requirements

The output from the initial mapping is a list of vectors indicating where pixels in image $\mathbf{I}_1$ have a best match in image $\mathbf{I}_2$. Associated with each vector is the value of correlation from the best match. It was expected that there would be spurious matches in the initial mapping as Fig. 4–7 shows. The object of the smoothing function is to identify these spurious matches and alter them, or alter a measure of their reliability, in some way which takes into account the neighbouring vectors and their reliability. The output from the smoothing process can then be used to extract information regarding the image transformation between $\mathbf{I}_1$ and $\mathbf{I}_2$.

Whatever smoothing function is used it must have certain characteristics. It is not sufficient to only look at vectors whose correlation is below a certain value. For instance, a vector with a strong correlation might point in the wrong direction and a vector with a weak correlation may nevertheless be pointing in the correct direction. There is a trade off required between the direction of the vectors in the neighbourhood and the reliability associated with the vector. Also, if the input to the smoothing process is a list of identical vectors the output must be that

same list of identical vectors. The following section explains how the smoothing algorithm achieves the above requirements.

5.2 The Smoothing Algorithm

The smoothing of the initial mapping was achieved by an iterative, weighted average of a vector and its neighbours. Eq. 5.1 is the mathematical expression of this smoothing:

$$\underline{V}(i,t+1) = \frac{\underline{V}(i,0) \times w(\underline{i}) + \alpha \sum_{\underline{k} \in \mathbf{N}} \underline{V}(\underline{i}+\underline{k},t) \times w(\underline{i}+\underline{k})}{w(\underline{i}) + \alpha \sum_{\underline{k} \in \mathbf{N}} w(\underline{i}+\underline{k})} \tag{5.1}$$

where $\underline{V}(\underline{i},t)$ is the vector mapping at a pixel in position $\underline{i}$ after t iterations ($t = 0$ is the original data). $w(\underline{i})$ is a weighting associated with the vector $\underline{V}(\underline{i},0)$. $\mathbf{N}$ is a set of vectors defining a neighbourhood and α is a gain factor given to the neighbourhood of vectors. This gain factor determines how much influence the neighbouring mappings have. Eq. 5.1 incorporates the neighbourhood vectors while also maintaining a measure of the original matches and a measure of the reliability. The form of this equation is inspired by typical algorithms resulting from regularisation [Tikhonov 77] assumptions. These assumptions are that the observed data is a corrupted version of a known class of function. Many regularisation models are implemented by iterative functions such as Eq. 5.1. In this case the class of function which the corrupted data, the initial mapping, has to match is unknown so there remain four problems:

1) How should a weight be chosen for each vector?

2) How large should a neighbourhood be?

3) What should the value of α be?

4) How many iterations should be executed?

Choosing Weightings

The weighting given to each vector is intended to reflect the "reliability of match" for that vector. That is, vectors with a higher weighting are considered to be more reliable and have more influence in the smoothing process. It was hoped that the correlation values would provide a reasonable measure of the reliability of a match. However, as Figs. 4–8 and 4–9 show, the distribution of values is limited and therefore not a good measure of reliability. To improve upon this, the correlation value for a pixel was combined with the local standard deviation about that pixel. The normalised correlation function relies on a high image standard deviation to work well so if a region has a high standard deviation it can be expected to have a good match. By multiplying the correlation value and the local standard deviation the weightings would take into account how well a pixel was expected to match and how well it actually matched. The resultant weightings were normalised to lie within the range 0-1.

$$w(\underline{i}) = \frac{c(\underline{i}) \times Sd(\underline{i}) - min(\mathbf{cSd})}{max(\mathbf{cSd}) - min(\mathbf{cSd})} \tag{5.2}$$

$min(\mathbf{cSd})$ and $max(\mathbf{cSd})$ find the minimum and maximum of the set of all $c(\underline{i}) \times Sd(\underline{i})$.

It was because the standard deviation and correlation value were multiplied together that the standard deviation was measured over the same region as the correlation.

Fig. 5–1 shows a histogram of the weights when the image shown in Fig. 1–1 was matched with the image shown in Fig 7–9. The large peak is the background data (i.e. air around the head), points with a weighting of $\approx 0.05 - 0.2$ lie in the central regions of the brain and points with a weighting bigger than ≈ 0.2 tend top lie on edges of the skull and brain. This suggests that weightings are given to appropriate points since the edges, which are expected to provide the best matches, have the higher weightings. The intention behind assigning the weightings was that the strong matches would propagate through the smoothing and pull round vectors which had been poorly mapped.

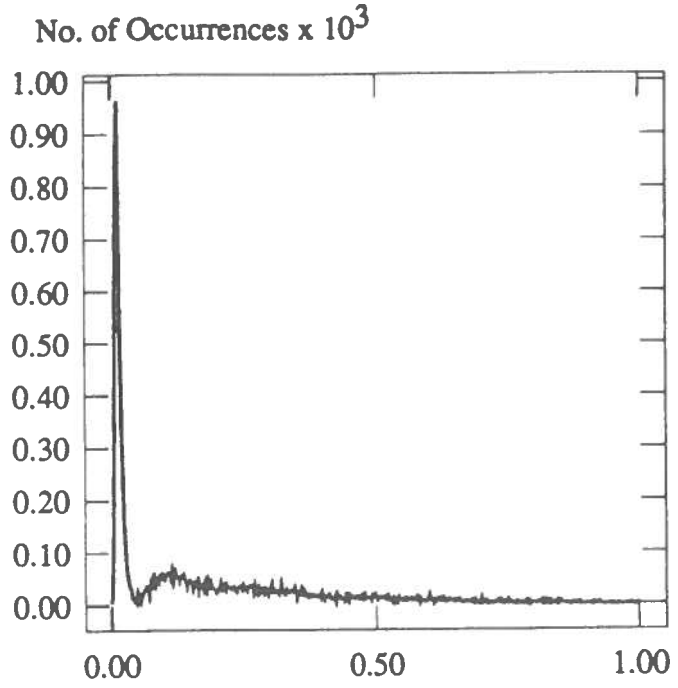

Figure 5–1: Histogram of weightings.

Neighbourhood Size

At each iteration a vector will be affected by the neighbourhood defined around it. Neighbouring vectors can be expected to be different for three reasons; spurious matches, gradual changes across an image due to the effects of rotation and magnification, and changes to the brain which may have occurred. If the neighbourhood used in the smoothing is too large the gradual differences between vectors due to rotation, magnification and changes will be engulfed by the neighbourhood. In order to minimise this effect only a 3×3 neighbourhood was used. Thus spurious vectors would be removed, but the smooth differences of rotation and change would remain. For a 3×3 neighbourhood $\mathbf{N}$ of Eq. 5.1 is defined as:

$$\mathbf{N} = ((1,1),(1,0),(1,-1),(0,1),(0,-1),(-1,1),(-1,0),(-1,-1)) \qquad (5.3)$$

The problems of what value of α to choose and how many iterations to execute remain unresolved. A formal mathematical analysis of the smoothing would be required to determine rigorously what values to choose for weighting and iterations. Failing this a series of tests were carried out and these are described in Section 5.3.

5.2.1 Example

Fig. 5–2 shows the vector mapping of Fig. 4–7 after smoothing, with $\alpha = 1$ and 3 iterations. It is evident that spurious vectors have been removed and so requirement 3 of Section 3.2 has been met. Also, the "smoothness" of the mapping suggests that the neighbourhood trends have been taken into account appropriately. There is, however, still the problem of the bad matches discussed in Section 4.3.3. Because these matches lay on image edges they were given high weightings. This, combined with the fact that they were in groups, caused them to propagate through the smoothing. Such behaviour is to be expected, but highlights the need to overcome the bad matching.

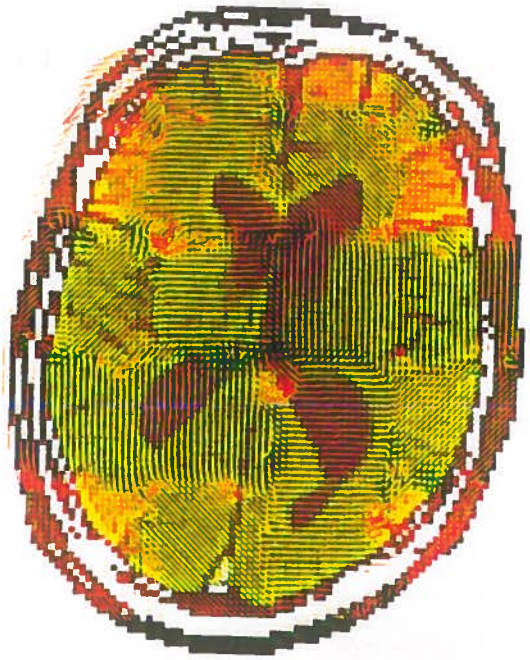

Figure 5–2: Vector mapping after smoothing

Overall the smoothing works well and provides the required behaviour. More work is required on understanding the dynamic behaviour of the smoothing if it is to be used reliably. This requires a formal method of determining α and the number of iterations so that particular types of smoothing behaviour can be obtained from the smoothing process. During the project values were chosen on a trial and error basis.

The final section presents a series of tests carried out on the smoothing to obtain a feel for the dynamic behaviour.

5.3 Testing Of The Smoothing

How the smoothings affected a vector mapping as the number of iterations, weightings and gain factors were changed (the dynamic behaviour) was not well understood. In an attempt to understand this dynamic behaviour a known vector mapping was smoothed several times, with different values of the parameters each time.

The mapping consisted of 900 vectors. Of these 900, the 800 external vectors pointed up and to the left and the 100 central vectors pointed down and to the left. Fig. 5–3 shows this mapping with the vectors indicated in white. All the vectors were the same length, but where vectors overlap the appearance is that of one long vector. The various tests are described below.

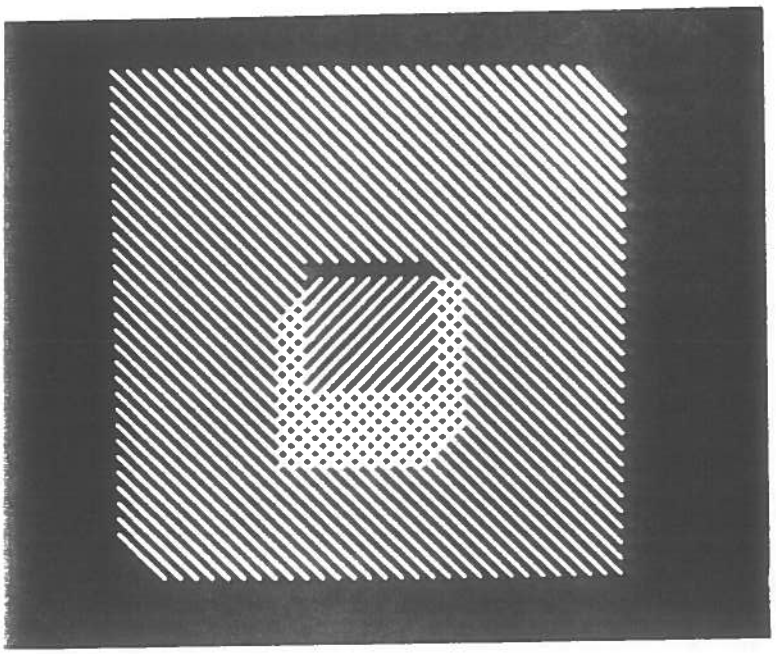

Figure 5–3: The test vector mapping.

Test 1 - Equal weightings, Equal Gain, Varied Iterations

The number of iterations performed effects the amount of smoothing and to what extent the "influence" of vectors extend. Keeping the weightings of the central and outer pixels equal and setting the gain at 1 the vector mapping of Fig. 5–3 was smoothed at 4, 16, 64 and 256 iterations. Figs. 5–4 and 5–5 show the effects of 4 and 256 iterations. It can be seen that after 4 iterations vectors on

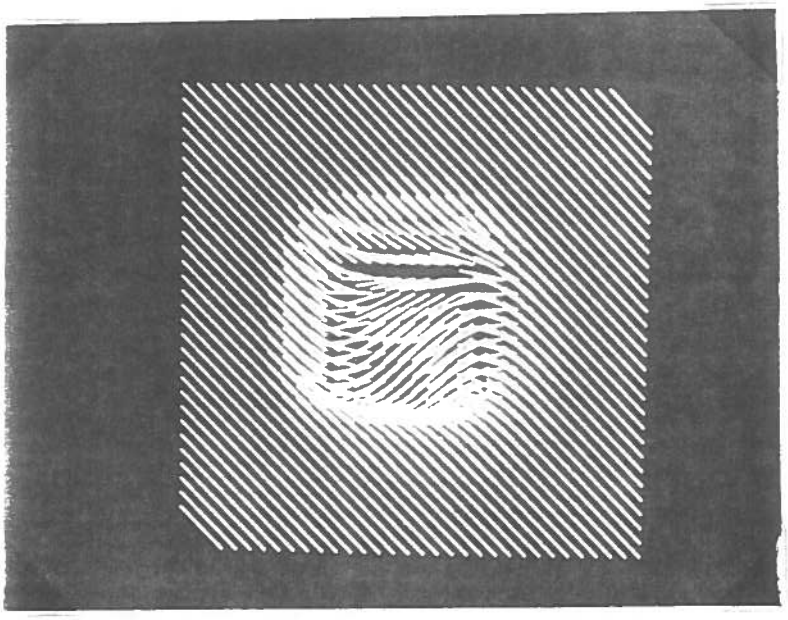

Figure 5–4: Test 1 - 4 iterations.

the border between the central and outer regions have been affected. After 256 iterations all the vectors in the central region have been pulled round to line up in approximately the same direction as the outer region. Vectors in the outer region closest to the central region have also been affected slightly. Despite the fact that all the weightings were the same the central vectors have all been affected.

From the above behaviour it may be concluded that:

- More iterations increase the smoothing effect and spread the influence of vectors further.

- Small regions of vectors will be engulfed by the surrounding vectors after large numbers of iterations.

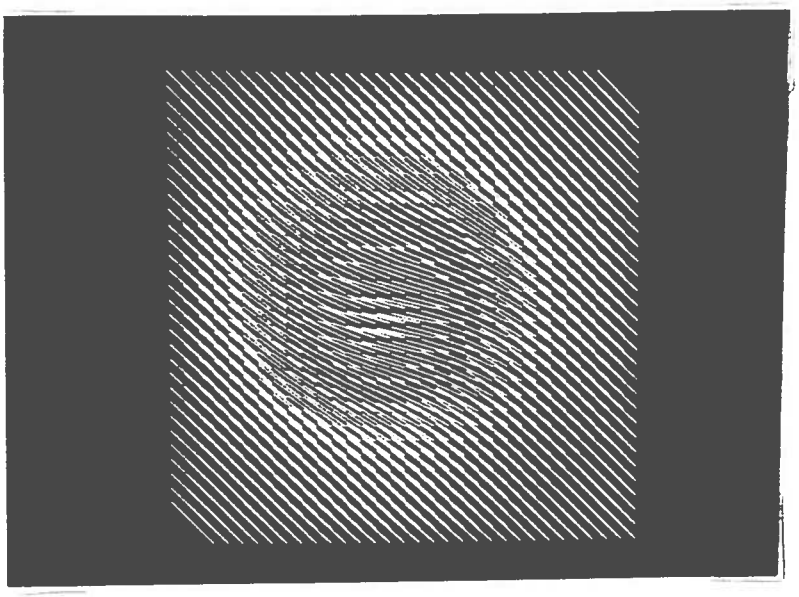

Figure 5–5: Test 1 - 256 iterations.

Test 2 - Unequal weightings, Equal Gain

Although the weightings to be given to vectors were automatically calculated it is still useful to understand how different weights affect the smoothing process. In Test 2a the outer vectors had a weighting twice that of the central vectors, Fig. 5–6, and in Test 2b the central vectors had a weighting twice that of the outer vectors, Fig. 5–7. In both cases 256 iterations were performed. In Fig. 5–6 the central

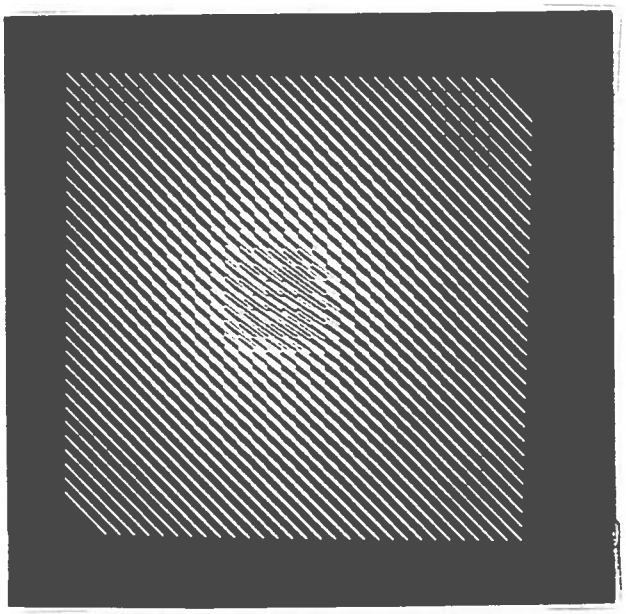

Figure 5–6: Test 2a - 256 iterations.

vectors have been pulled round almost completely. This would be expected given the results for Test 1. In Fig. 5–7 the central vectors have been pulled round,

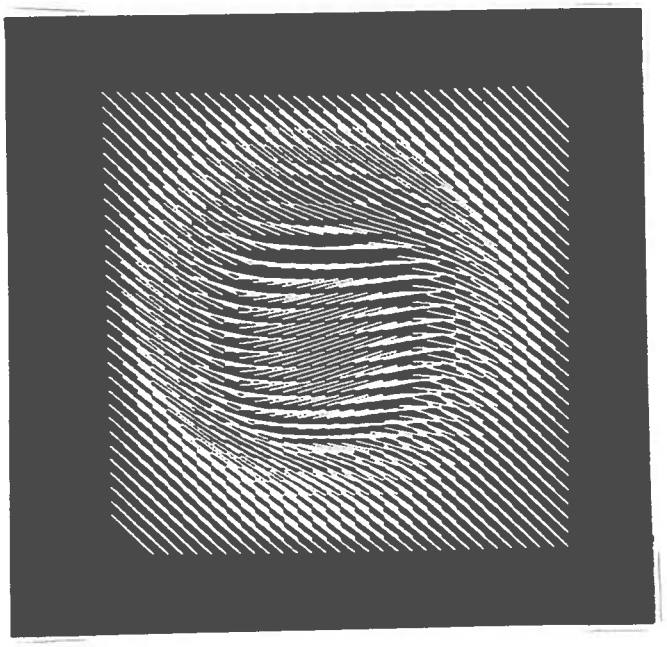

Figure 5–7: Test 2b - 256 iterations.

but only half as much as in Test 2a. Some of the outer vectors have also been pulled downwards. What is evident here is that despite the central vectors having a higher weighting, the mass of outer vectors has still influenced them.

From Test 2 it can be concluded that:

- Weightings do influence how neighbouring vectors interact, but the effects of the neighbourhood can dominate after many iterations.

Test 3 - Equal weightings, Unequal Gain

The final test was to investigate the affect of the gain factor, α, on the smoothing. In Test 3a $\alpha = 4$ and in Test 3b $\alpha = 0.25$. In both tests the weightings of all the vectors were equal and 256 iterations were performed.

Fig. 5–8 shows the results of Test 3a. All the central vectors have been pulled round, not as far as in Test 2a, Fig. 5–6, and about the same as in Test 1a, Fig. 5–5. This suggests that the weightings and number of iterations are having a greater influence than this particular value of α. This is expected given the results in Tests 1 and 2 where neighbourhood structures were already affecting the result and increasing α would only enforce the neighbourhood effect.

Fig. 5–9 shows the results of Test 3b. This is particularly notable because the central region of vectors has remained relatively undisturbed. At the border

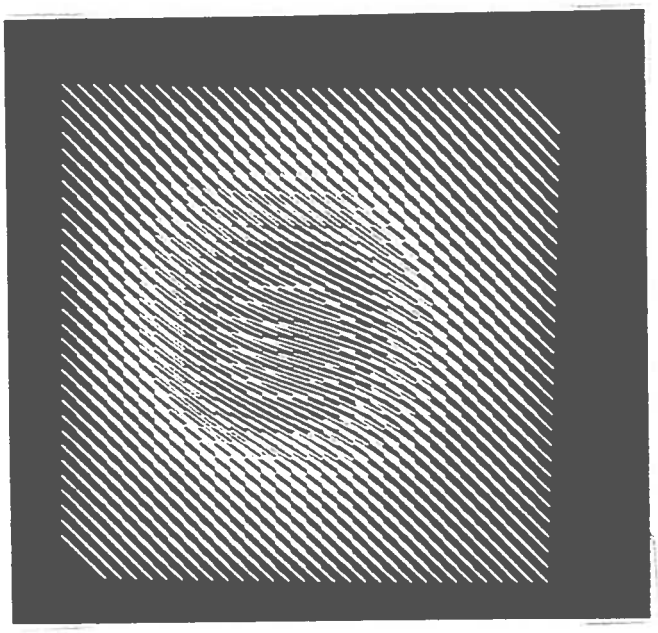

Figure 5–8: Test 3a - 256 iterations.

between the central and outer vectors there have been changes, but other than that most of the vectors have remained close to their original values. This behaviour suggests that the gain factor might act as some kind of filter frequency cut off selector. If smoothing is considered as being a spatial filter, then it must have a spatial frequency response. The smaller α becomes the higher spatial frequencies (sharper changes) will be permitted to pass through the filter. Fig. 5–9 looks very similar to Fig. 5–4 so it is likely that different combinations of gain and iterations can give similar behaviour.

From Test 3, the most informative in the light of the other tests, it may be concluded that:

- The gain factor acts like a cut off for sharp changes. The smaller the gain the sharper the changes
- The number of iterations can still dominate the process for certain gain factors
- It is likely that the number of iterations and the gain are dependent variables. High gain will move more quickly to a result similar to that achieved by lower gain over more iterations.

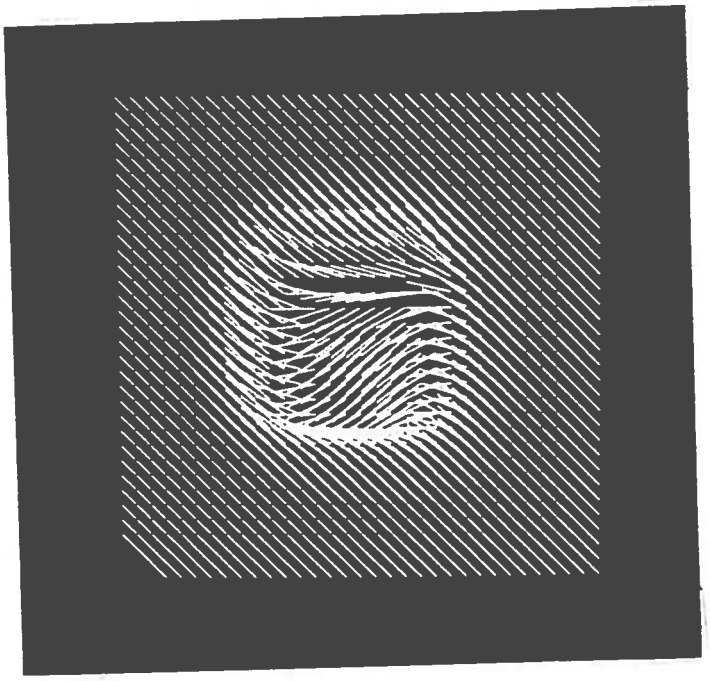

Figure 5–9: Test 3b - 256 iterations.

From the tests carried out a qualitative understanding of the properties of the smoothing has been gained. Ideally a rigorous mathematical analysis indicating precisely how the gain and number of iterations interact and how to chose values for these parameters to obtain certain behaviours would be needed. However, failing this certain general statements can be made about the smoothing.

- The gain and number of iterations are dependent variables.
- Large regions can engulf smaller regions.
- The weightings do affect the output of the smoothing appropriately, however a large region may have more influence due to the neighbourhood effect after many iterations.
- The more iterations are performed the further the influence of vectors spread, but there does seem to be a limit.

With reference to the problem of smoothing an initial mapping, the requirement is that spurious vectors are removed and groups of spurious vectors will be maintained, since these represent changes. It is expected that these groups will be small, since the changes will be small (Assumption **3**, Section 3.1), so a small gain to prevent large neighbourhood effects and a few iterations would seem reasonable. In practice $\alpha = 1$ and iterations between 3 and 5 produced acceptable results.

5.4 Summary

The smoothing algorithm has been explained in this chapter and tests carried out to investigate the dynamic behaviour have been presented. The example given demonstrates that the smoothing successfully removes and corrects spurious matches. The removal of the image transformation from the smoothed initial mapping to reveal the residual data is explained in the next chapter.

Chapter 6

Image Transformations and Analysis of Residual Data

The expected output from the smoothing process is a vector mapping indicating where points in $\mathbf{I_1}$ map to in $\mathbf{I_2}$. The mapping contains information on both the image transformation and the changes which have occurred. This chapter describes how the effects of image transformation can be removed by statistical methods, leaving only the information on the changes. An example of this process is given. The analysis of residual data, although not implemented, is discussed briefly.

6.1 Calculation Of The Transformation

[Meertens 90]
If an image $\mathbf{I_1}$ undergoes a translation (t_x, t_y), rotation θ, defined to be +ve anticlockwise, and magnification λ to become $\mathbf{I_2}$ then points P_i, (x_{1i}, y_{1i}) in $\mathbf{I_1}$ will be expected to map to the points $\hat{E}_i$, (x_{ei}, y_{ei}) in $\mathbf{I_2}$ according to the following equation:

$$\begin{pmatrix} x_{ei} \\ y_{ei} \end{pmatrix} = \lambda \begin{pmatrix} cos\theta & -sin\theta \\ sin\theta & cos\theta \end{pmatrix} \begin{pmatrix} x_{1i} \\ y_{1i} \end{pmatrix} + \begin{pmatrix} t_x \\ t_y \end{pmatrix} \tag{6.1}$$

As was shown in Section 4.3.1 the centres of rotation and magnification can be incorporated into the translation, so that the rotation and magnification are treated as being centred on $(0, 0)$.

The transformation parameters between $\mathbf{I_1}$ and $\mathbf{I_2}$ are unknown, but the output from the regularisation is a list of points M_i, (x_{2i}, y_{2i}) in I_2, which match the points P_i, (x_{1i}, y_{1i}) in $\mathbf{I_1}$. If more than two matching pairs are known the system is overdetermined for the unknown parameters t_x, t_y, θ and λ. These parameters will not usually have a solution because the matching points are not exact and will be perturbed by noise and changes. The values of the parameters are then determined by minimising the sum of the squared distances between the actual match points M_i and the expected match points $\hat{E}_i$.

Defining:

$$m_x = \lambda \quad cos\theta \tag{6.2}$$

$$m_y = \lambda \quad sin\theta \tag{6.3}$$

Eq. 6.1 can be written as:

$$x_{ei} = m_x x_{1i} - m_y y_{1i} + t_x \tag{6.4}$$

$$y_{ei} = m_x y_{1i} + m_y x_{1i} + t_y \tag{6.5}$$

and in vector notation:

$$\hat{\mathbf{E}} = \mathbf{Pt} \tag{6.6}$$

where:

$$\hat{\mathbf{E}} = \begin{pmatrix} x_{e1} \\ y_{e1} \\ \vdots \\ x_{en} \\ y_{en} \end{pmatrix}, \mathbf{t} = \begin{pmatrix} m_x \\ m_y \\ t_x \\ t_y \end{pmatrix}, \mathbf{P} = \begin{pmatrix} x_{11} & -y_{11} & 1 & 0 \\ y_{11} & x_{11} & 0 & 1 \\ \vdots & \vdots & \vdots & \vdots \\ x_{1n} & -y_{1n} & 1 & 0 \\ y_{1n} & x_{1n} & 0 & 1 \end{pmatrix}$$

x_{ei} and y_{ei} are linear functions of m_x, m_y, t_x and t_y. The solution to this linear least squares problem is:

$$\mathbf{t} = (\mathbf{P}^T\mathbf{P})^{-1}\mathbf{P}^T\mathbf{M} \tag{6.7}$$

where:

$$\mathbf{M} = \begin{pmatrix} x_{21} \\ y_{21} \\ \vdots \\ x_{2n} \\ y_{2n} \end{pmatrix}$$

The actual, rather than expected match points

The inverse of the matrix $\mathbf{P}^T\mathbf{P}$ is found by an LU decomposition routine [Press 89] and so, knowing $\mathbf{P}$ and $\mathbf{M}$, $\mathbf{t}$ can be calculated to find the values of m_x, m_y, t_x and t_y[1]. From Eqs. 6.2 and 6.3 it follows that:

$$\lambda = (m_x^2 + m_y^2)^{\frac{1}{2}} \tag{6.8}$$

and:

$$\theta = \tan^{-1}\left(\frac{m_y}{m_x}\right) \tag{6.9}$$

6.2 Extraction Of The Transformation

Having calculated m_x, m_y, t_x and t_y the points $\hat{E}_i$ can be found from Eqs. 6.4 and 6.5. The effects of global transformation can then be removed with any residual representing the change.

$$x_{ri} = x_{1i} + x_{2i} - x_{ei} \tag{6.10}$$

$$y_{ri} = y_{1i} + y_{2i} - y_{ei} \tag{6.11}$$

The vector from (x_{1i}, y_{1i}) to (x_{ri}, y_{ri}) indicates the change that has occurred at the point (x_{1i}, y_{1i}) between $\mathbf{I_1}$ and $\mathbf{I_2}$. It should be realised that being a statistical process the transformation calculation and subsequent extraction can only be as good as the data it is presented with.

---

[1]The values of t_x and t_y calculated are not the real image translation, but the result of treating the rotation and magnification as if they were centred on $(0,0)$. See Section 4.3.1.

6.3 Example

Fig. 6–1 shows the vectors indicating the changes detected after the effects of image transformation have been removed from the smoothed mapping of Fig. 5–2. The actual rotation was -0.05 radians and the calculated rotation was -0.049 radians. It can be seen that the majority of the vectors have been reduced to

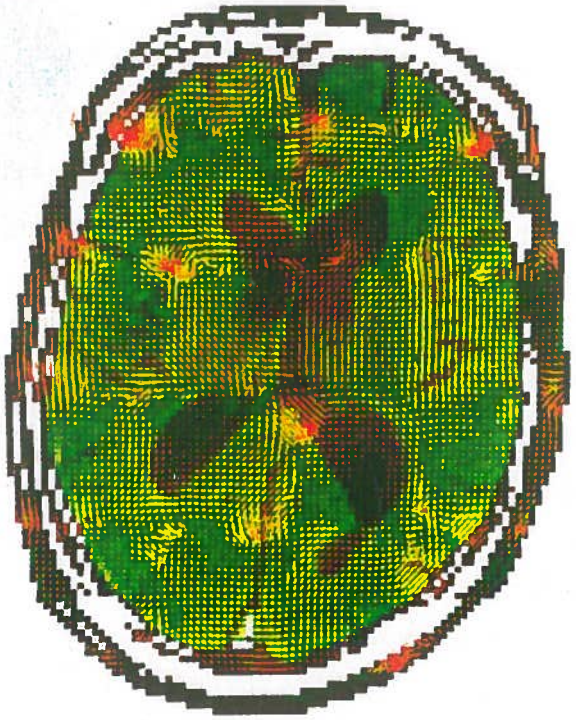

Figure 6–1: Vector mapping after transformation extraction.

≈ 0. The clumps of larger vectors are the result of the groups of bad matches highlighted in Section 4.3.3.

The method of image transformation extraction detailed above is one of many statistical methods. It would be expected to work well. What is noteworthy, however, is the fact that despite there being some bad matches the correct rotation has been calculated and that most of the vectors are showing little or no change.

The entire process of change detection has now been described and the examples shown in Chapters 4-6 are for a test where one image was matched with a rotated version of itself. All that remains is to analyse the residual data. Possible types of analysis are outlined below.

6.4 Analysis Of Residual Data

After removal of the effects of image transformations from the vector mappings what remains should be vectors indicating the changes which have occurred plus vectors indicating where the brain has moved within the skull. The first stage of the residual analysis should be to remove the effects of the brain motion, leaving only vectors indicating real structural changes. The removal of the brain motion could be complicated by the fact that the surface of the brain may be changing.

Before an analysis of the residual vectors could be implemented a clinician would have to be consulted to discuss the form the analysis should take. Possibilities might be; bunching vectors together to indicate regions which were changing, calculating volume changes from the vectors, classifying regions as contracting or expanding and calculating rates of change from a sequence of images.

A symbolic change analysis based on the region parameters described above might be able to remove the bad matches which propagate from the initial mapping. Since most of the bad matches occur along edges and changes would normally be expected to occur in directions orthogonal to edges, bad match groups could be eliminated on this basis.

6.5 Summary

In this chapter a statistical method for calculating image transformations from known match points has been described and future work for the analysis of the residual vectors has been suggested. The change detection has now been fully described. In the next chapter the performance of the change detection algorithms is examined further by evaluating the results of testing.

Chapter 7

Test Results

Throughout Chapters 4-6 an example has been shown of the change detection working on a manufactured data pair and statistical measures of the change detection working on a real data pair have been presented. The results of further tests on manufactured and real data are shown below.

7.1 Tests On Manufactured Data

A manufactured data pair was formed by taking an MRI image (Fig. 1–1) and transforming it in some way. The transformation was one for which the change was known so the results could be predicted. The transformations were; 1) magnification, 2) rotation, 3) translation and 4) translation of a small region within the image. The test results are described below and shown where informative.

Test 1: When the image was matched with itself there was no residual data left. This demonstrated that the algorithms were working correctly. The residuals are not shown, since there were none.

Test 2: When the image was matched with a translated version of itself no residual data was left. The output from the initial mapping was a list of identical vectors, which the smoothing left unchanged, as expected. This test confirms the operation of the smoothing algorithm on a uniform mapping and reiterates the correct operation of the algorithms. The residuals are not shown.

Test 3a: When the image was matched with a rotated version of itself (-0.05 rads, $-2.5°$) some residuals were left . These are shown in Fig. 6–1. The results of this test were discussed in Chapters 4-6. The calculated rotation was -0.049 rads., which compares well with the actual -0.05 rads.

Test 3b: When the image was matched with a rotated version of itself (-0.1 rads, $-5°$) there were more residuals than in Test 3a. The residual mapping shown in Fig. 7–1, which should be compared with Fig. 6–1, displays more residuals and larger patches of bad matches. The calculated rotation was -0.091 rads. compared with -0.1 rads., not as accurate as Test 3a. This suggests that the correlation produces larger bad patches as the rotation increases, which should be expected given the assumption that the rotations would only be small and thus would only cause a small error in the correlation value.

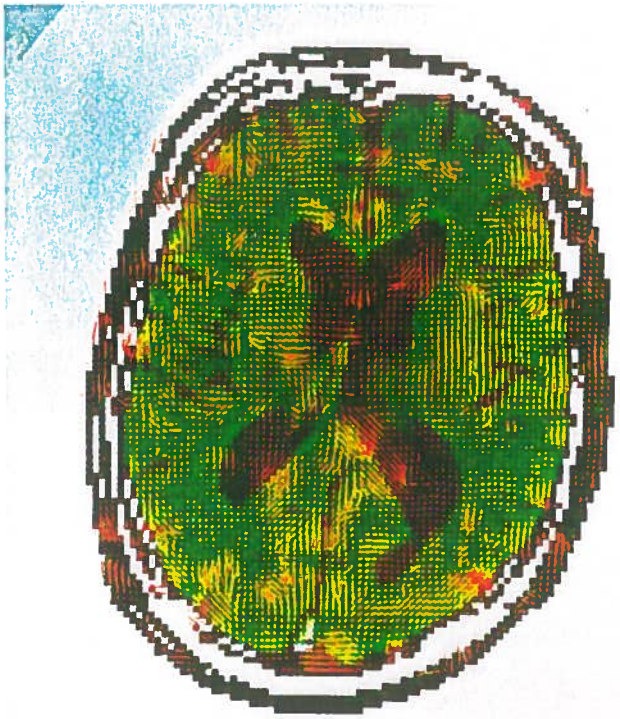

Figure 7–1: Residual mapping: Test 3b, Rot. -0.1 rads.

Test 4: The result of matching the image with a magnified (1.05) version of itself is shown in Fig. 7–2. Overall there are almost no residual vectors, but there are occasional bad matches near edges. The calculated magnification was the same as the actual magnification of 1.05.

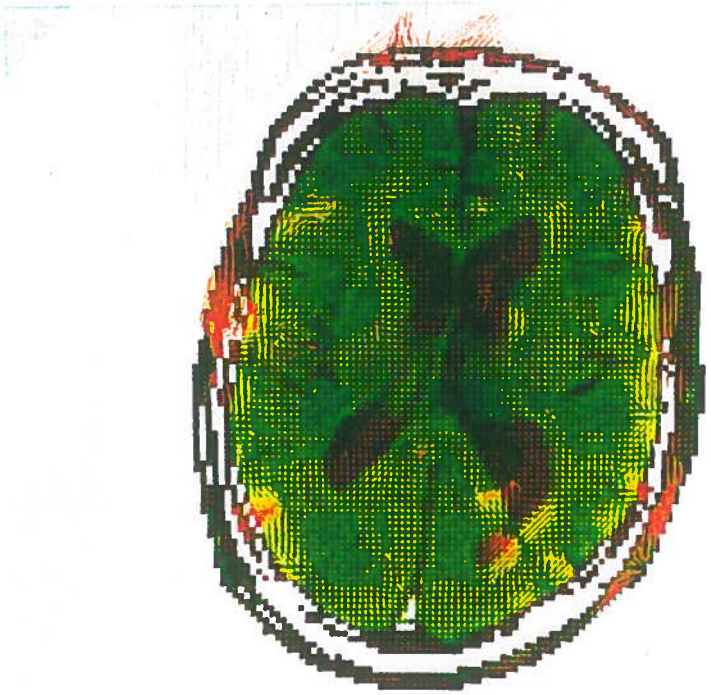

Figure 7–2: Residual mapping: Test 4, Mag. 1.05.

Tests 1-4 were on images where only one of the possible transformations was performed. These results show that the change detection has no difficulty in deducing and removing translation and magnification has very few residual vectors. The rotations cause the most bad matches resulting in residuals, with more residuals as the rotation is increased. The next set of tests, 5-8, combine the possible transformations.

Test 5: The result of matching an image with a rotated (-0.05 rads.) and translated[1] version of itself is shown in Fig. 7–3. The result is almost identical to that of Fig. 6–1 indicating that the translation has been successfully removed, but the image rotation caused the same occasional bad mappings. The calculated rotation was -0.049 rads.

Test 6: Matching the image with a rotated (0.05 rads.) and magnified (1.01) version of itself resulted in the residual mapping shown in Fig. 7–4. As for previous

[1]The value of the image translation is not relevant. See Section 4.3.1

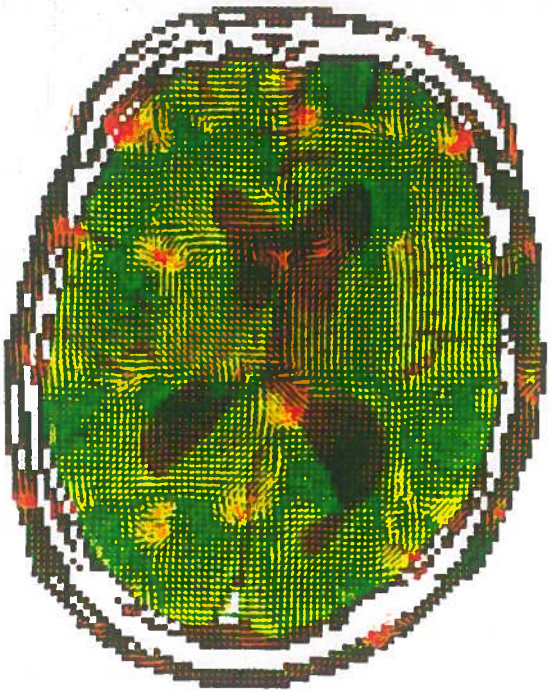

Figure 7–3: Residual mapping: Test 5, Rot. -0.05, Trans.

tests, the residuals are mostly zero with occasional bad matches. The rotation was not accurately calculated, -0.043, but the magnification was, 1.01.

Test 7: The result of matching the image with a magnified (0.95) and translated version of itself is shown in Fig. 7–5. Almost all the residual vectors are close to zero. This confirms the fact that the initial mapping copes better with magnification than rotation. Once more the effect of translation in this test has been completely removed. The calculated magnification was 0.95.

Test 8: In the final transformation test, the image was matched with a rotated (0.03 rads.), magnified (1.025) and translated version of itself. The residual mapping is shown in Fig. 7–6. As usual there were bad matchings resulting in residuals, but fewer than Tests 3a, 3b, 5 and 6. The rotation in this test was smaller than previous tests and this will account for the fewer bad matchings. The calculated magnification was 1.025, and the rotation 0.024 rads.

Summary: Tests 1-8 Throughout tests 1-8 an image was matched with a transformed version of itself, so the expected output would be no residual vectors (no changes). The factor which caused the most residuals to appear was the rotation, more residuals appearing with larger rotations. There was no serious

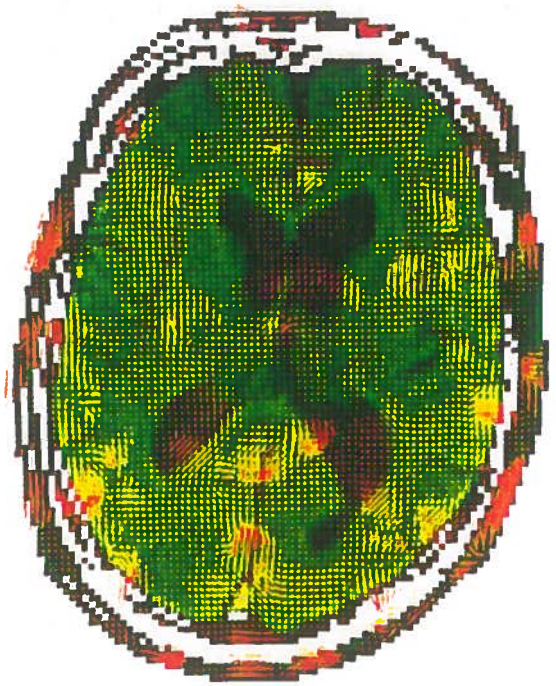

Figure 7–4: Residual mapping: Test 6, Rot. 0.05, Mag. 1.01.

degeneration of the residual mapping when rotation, magnification and translation were performed compared with rotation alone. The magnification, which was consistently calculated and removed accurately, caused far fewer groups of bad matches than the rotation. The translation was always removed and left no residual vectors.

The above tests demonstrate that the change detection technique can successfully calculate and remove image transformation parameters and the accuracy decreases as the rotation increases, as is to be expected. The magnification causes less bad matches than the rotation and is accurately calculated. The magnification and rotation can be compared in this way because 0.05 rads. rotation causes approximately the same size changes as a magnification of 1.05. The translation causes no difficulties whatsoever, as expected. The final test on manufactured data involves an artificial change inserted into an image.

Test 9: The manufactured change was effected by taking a small region (12×12 pixels) at the top of the right, upper ventricle, translating it up and to the left and laying it back over the image. The residual mappings for this test are shown in Figs. 7–7 and 7–8. Fig. 7–7 shows the effect when smoothing with $\alpha = 1$ and

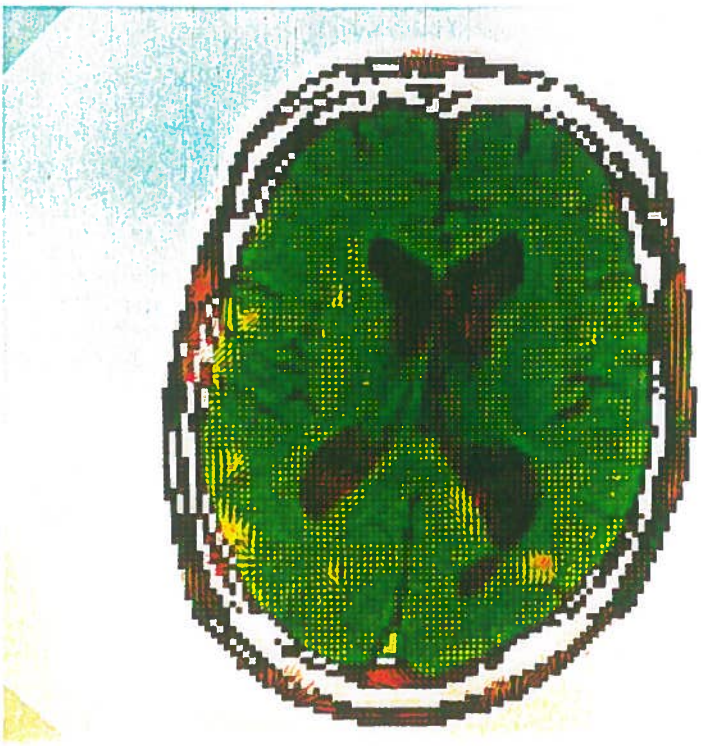

Figure 7–5: Residual mapping: Test 7, Mag. 0.95, Trans.

3 iterations was carried out on the initial mapping and Fig. 7–8 when $\alpha = 1$ and 10 iterations were performed. Both results show that the displaced region has been correctly matched. The effects of different smoothing iterations show that after more iterations the effects of neighbouring mappings have caused a smoother residual mapping. Although it is not immediately evident from Fig. 7–8 the residual region is slightly larger than in Fig. 7–7, indicating that the extent of influence of the vectors has spread. This test also shows that the local distortions have been detected irrespective of the smoothing, i.e. they have not been smoothed out.

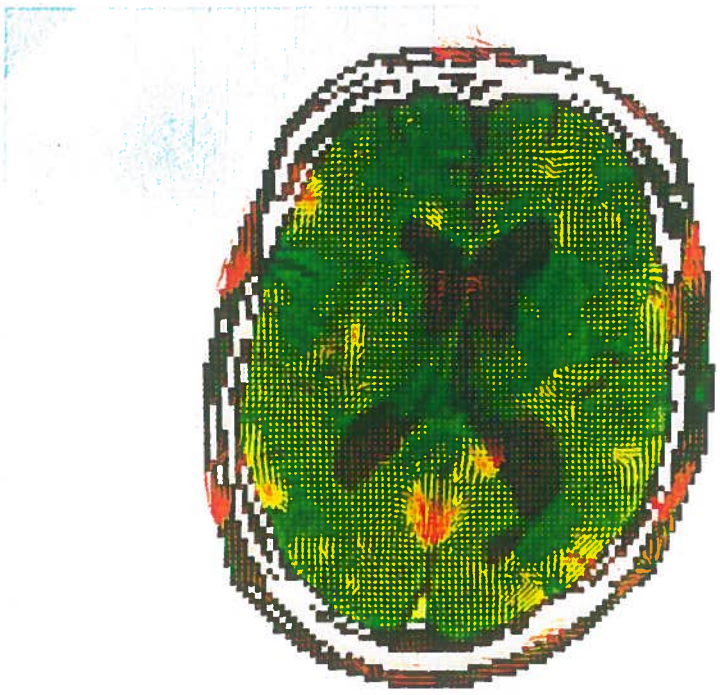

Figure 7–6: Residual mapping: Test 8, Mag. 1.025, Rot. 0.03, Trans.

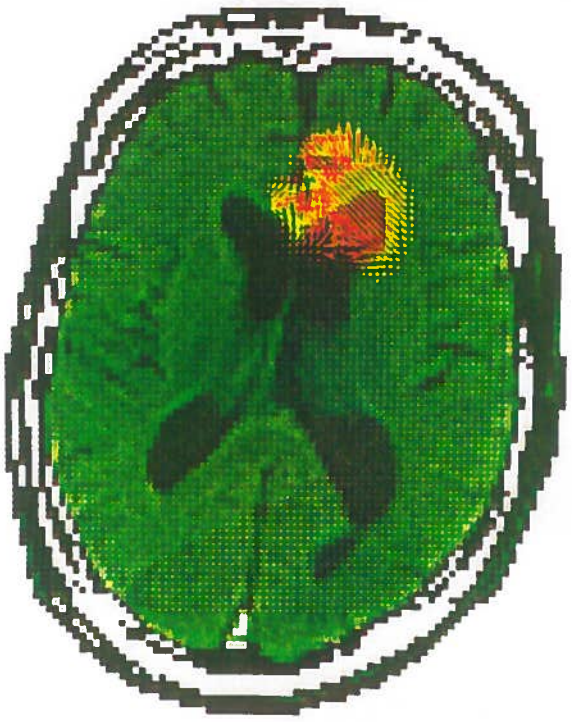

Figure 7–7: Residual mapping: Test 9 - 3 Smoothing iterations

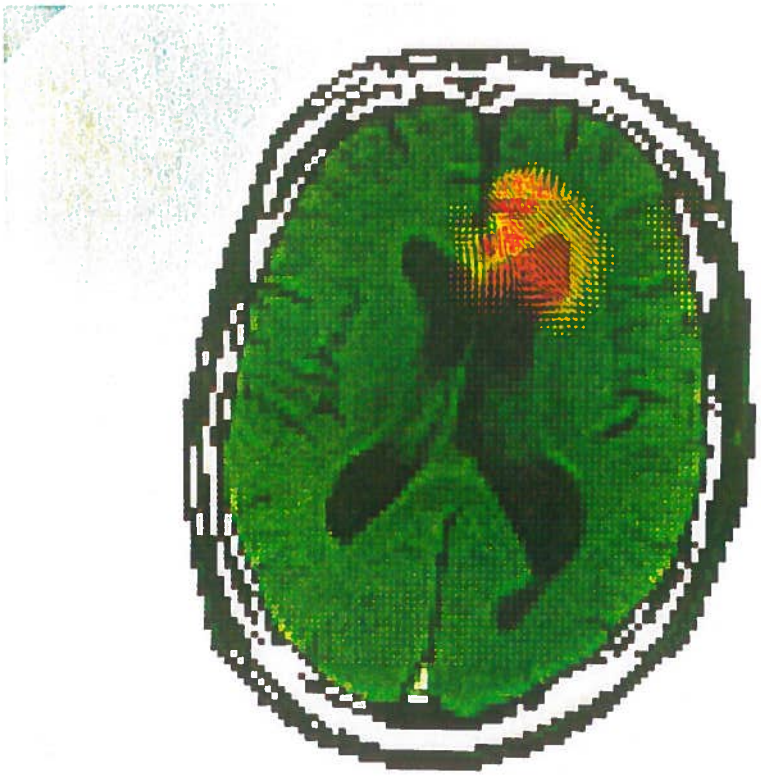

Figure 7–8: Residual mapping: Test 9 - 10 Smoothing iterations

7.2 Real Data Pair

The only real data pair available consisted of two images, taken of the same patient an unknown period of time apart. The image shown in Fig. 1–1 was taken after that shown in Fig. 7–9. These images are not ideal since there is distortion between them (evident by the altered skull shape), but it can be seen that there are changes that have occurred to the ventricles (expansion between Fig. 7–9 and Fig. 1–1) and new features have appeared at the middle of the sides of the brain. The ventricle changes are larger than the change detection would normally be expected to work on since they are obvious to an observer. The change detection is intended for use where the observer has difficulty detecting the changes.

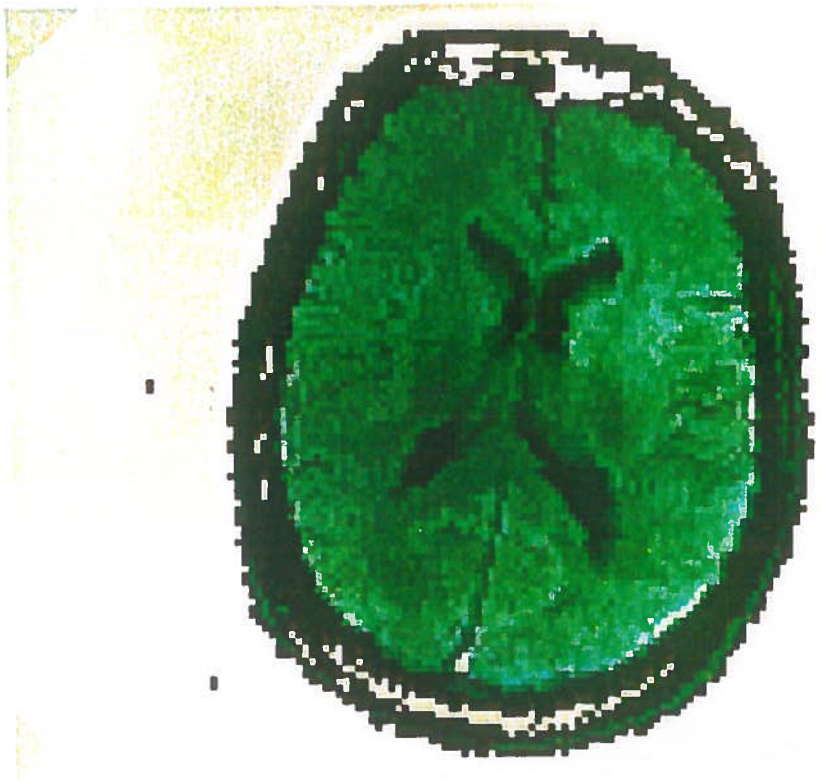

Figure 7–9: Second MRI slice. Separated in time from Fig. 1-1.

Fig. 7–10 shows the residual mapping when the smoothing process is iterated 3 times before transformation extraction and Fig. 7–11 shows the residual mapping when the smoothing process is iterated 10 times before transformation extraction. What can be interpreted from these pictures? It can be seen that in Fig. 7–10 the vector mapping is not as smooth as in Fig. 7–11 and for both there is a lot of vector

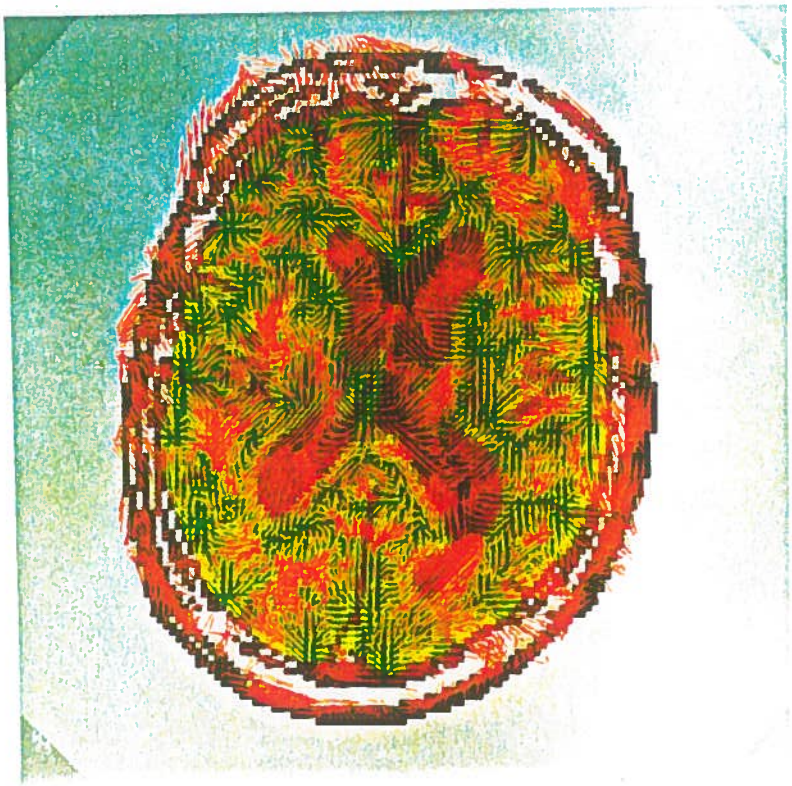

Figure 7–10: Residual mapping, Real data, 3 Smoothing iterations.

"activity" around the ventricles. Comparing the original data, dramatic changes can be seen in the regions of the ventricles, so this vector motion is expected. Another region of activity is the middle of the sides of the brain where new features have appeared in the later image of Fig. 1–1. While the correctness and relevance of these results cannot be accurately assessed, regions of large residual content do seem to correspond to regions of strong change between the two test images. In order to make a qualitative assessment of these mappings new images were created by mapping pixels in Fig. 1–1 to where the residual vectors indicated that they had moved[2]. Fig. 7–12 shows the image constructed by using the residual mapping of Fig. 7–10 and Fig. 7–13 shows the image constructed from Fig. 7–11. There are many empty pixels because there were a considerable number of many to one mappings in the residual mapping. However, the ventricles reconstructed in Fig. 7–12 compare favourably with the real ventricles of Fig. 7–9. The ventricles

[2]In this test the later picture was matched to the earlier picture because of the representational problems of expansions. (See Section 3.1) Mapping the later image to the earlier one is equivalent to looking for contractions.

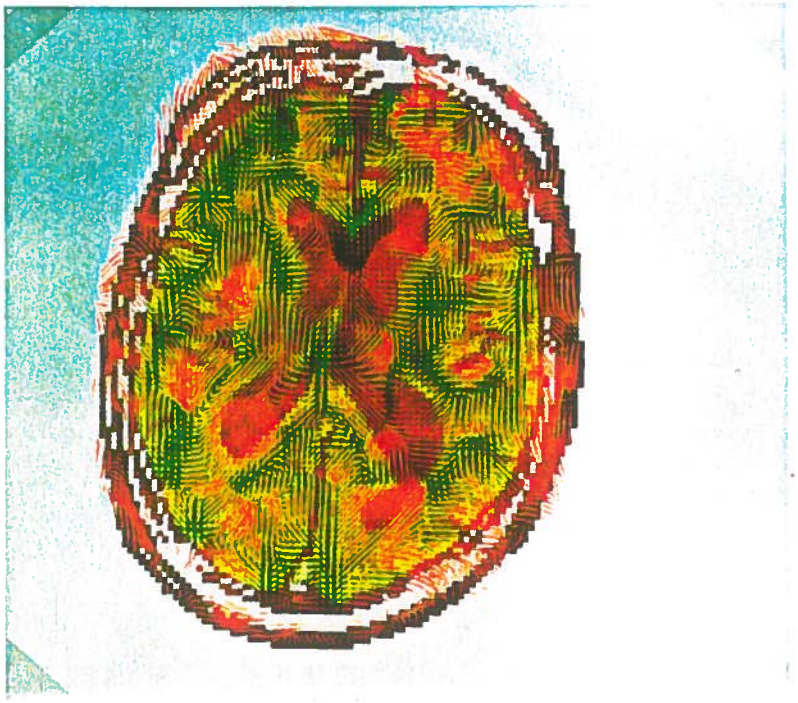

Figure 7–11: Residual mapping, Real data, 10 Smoothing iterations.

in Fig. 7–13 are slightly large compared with those in Fig. 7–9 and there are also fewer many to one (blank) mappings in Fig. 7–13.

This test has highlighted two difficulties. Firstly, the method of assessing the change detection is rather crude and should be improved before it can be used practically. Secondly, the differences that the smoothing can make to the residual mapping emphasise the importance of being able to correctly choose parameters for the smoothing process. This also shows how the smoothing process extends neighbourhood structure since more smoothing results in a more uniform mapping.

7.3 Summary

The tests carried out have shown that the change detection technique can successfully calculate and remove image transformation parameters. The accuracy decreases for larger rotations. Magnification, while not removed perfectly, causes fewer residuals than rotations. Translation causes no difficulties whatsoever. The validity of the smoothing function is shown by the fact that the image transfor-

Figure 7–12: Image constructed from Fig. 1-1 (3 Smoothing iterations).

mations are calculated accurately, so the accuracy of the initial mapping is not disturbed by the smoothing.

The test on the real data shows the difficulty in assessing the performance of the change detection and therefore the difficulty in assessing its practical use. It is also apparent that more knowledge is required on the dynamic operation of the smoothing if it is to be used confidently without oversmoothing results.

In the following chapter the practical problem of computational complexity is addressed and in Chapter 9 the overall performance of the change detection and the test results are discussed. Chapter 9 also contains conclusions on the achievements of the project and suggestions for further work.

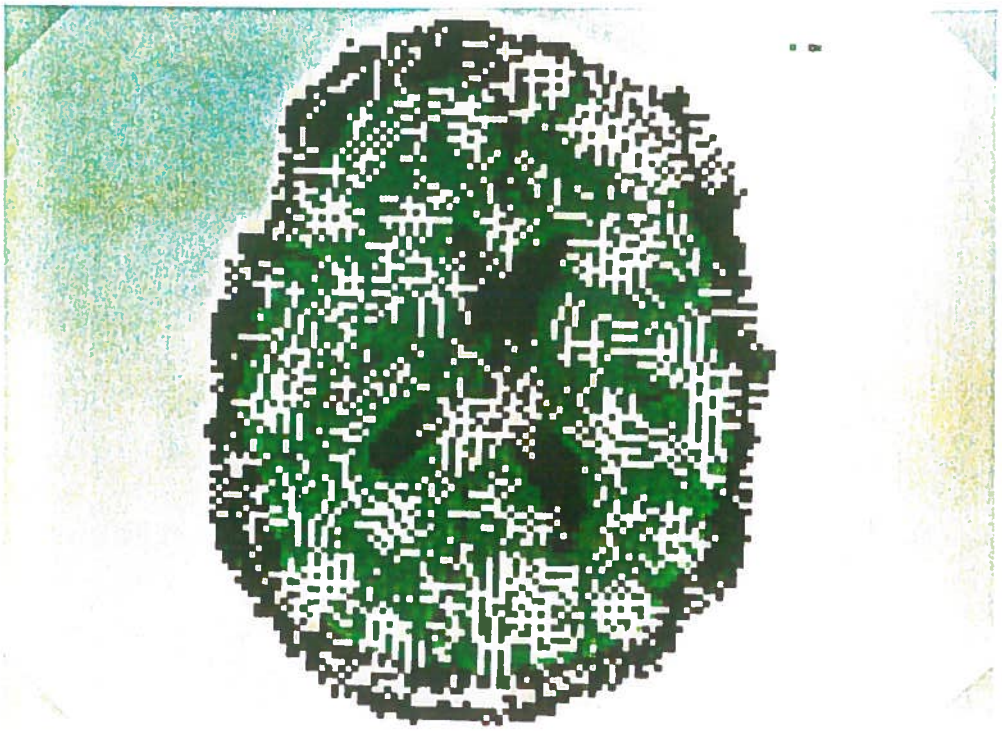

Figure 7–13: Image constructed from Fig. 1-1 (10 Smoothing iterations).

Chapter 8

Overcoming the Computational Complexity

To be of any practical use the change detection technique must produce results within an "acceptable" time. This chapter discusses the factors which affect the operating time and suggests ways of improving it. Examples of operating time are given along with the potential reductions.

8.1 Size Of Search Space

The initial mapping is effectively a search through an image $\mathbf{I}_2$ to find matches for the pixels in image $\mathbf{I}_1$. It is this stage of the processing which takes most time due to the calculation of many correlation values. By considering potential reductions in the search space an estimate can be made of reductions in the time taken to produce the initial mapping.

If an image has sides of length n and is d dimensional it contains n^d pixels (voxels if 3D). Searching all of $\mathbf{I}_2$ for a match for every pixel in $\mathbf{I}_1$ would require $n^d \times n^d$ matches to be tried. This is the slowest, brute force search strategy. However, the whole of $\mathbf{I}_2$ need not be searched, since under any transformation pixels in $\mathbf{I}_1$ will map to a certain region in $\mathbf{I}_2$. Rather than searching all of $\mathbf{I}_2$ only a region of r^d need be searched. This assumes that the image transformation is known approximately, which is reasonable. The number of matches to be tried is now reduced to $n^d \times r^d$. Such a search policy (Basic Region Search) was adopted

during the development of this project. A further reduction can be achieved by only searching for matches for pixels which contain relevant data, i.e. their value lies above a certain threshold. If, of the n^d pixels, p% contain relevant data then the size of the search space becomes:

$$S = n^d \times r^d \times \frac{p}{100} \tag{8.1}$$

$$\text{If } n = 128, d = 2, r = 25, p = 50 \text{ then } S = 5,120,000 \tag{8.2}$$

A more significant reduction in the size of the search space can be effected by searching for matches from coarse to fine resolution in a hierarchical manner (Hierarchical Search). This involves reducing an image size by scaling it down, searching for matches and using the results from coarse resolution to guide the search at the next finer level of resolution. If every time an image is reduced it is reduced by a factor of 2, then after s reductions there will only be:

$$\frac{n^d}{2^{d \times s}} \tag{8.3}$$

pixels to search. If the image is reduced then the search region can be reduced by the same amount. Although in the reduced image less pixels will be searched the same real data area is covered, but at a coarser resolution. The size of the search region becomes:

$$\frac{r^d}{2^{d \times s}} \tag{8.4}$$

Assuming that a match at one resolution will be close to a match at the next level of resolution then the search can remain nearby to the solution found at the previous level. Therefore the search region can remain the same size. Fig. 8–1 illustrates how such a hierarchical search works. The total search for s reductions is:

$$S = \frac{n^d}{2^{d \times 0}} \times \frac{r^d}{2^{d \times s}} + \frac{n^d}{2^{d \times 1}} \times \frac{r^d}{2^{d \times s}} + \ldots + \frac{n^d}{2^{d \times s}} \times \frac{r^d}{2^{d \times s}} \tag{8.5}$$

which is equivalent to:

$$S = \sum_{k=0}^{s} \frac{n^d}{2^{d \times k}} \times \frac{r^d}{2^{d \times s}} \tag{8.6}$$

and can be written as:

$$S = \frac{n^d \times r^d}{2^{d \times s}} \sum_{k=0}^{s} \frac{1}{2^{d \times k}} \tag{8.7}$$

The summation can be replaced by a closed form and allowing the inclusion of relevant pixels only Eq. 8.7 becomes:

$$S = \frac{n^d \times r^d}{2^{d \times s}} \times \frac{1 - 2^{-d \times (s+1)}}{1 - 2^{-d}} \times \frac{p}{100} \tag{8.8}$$

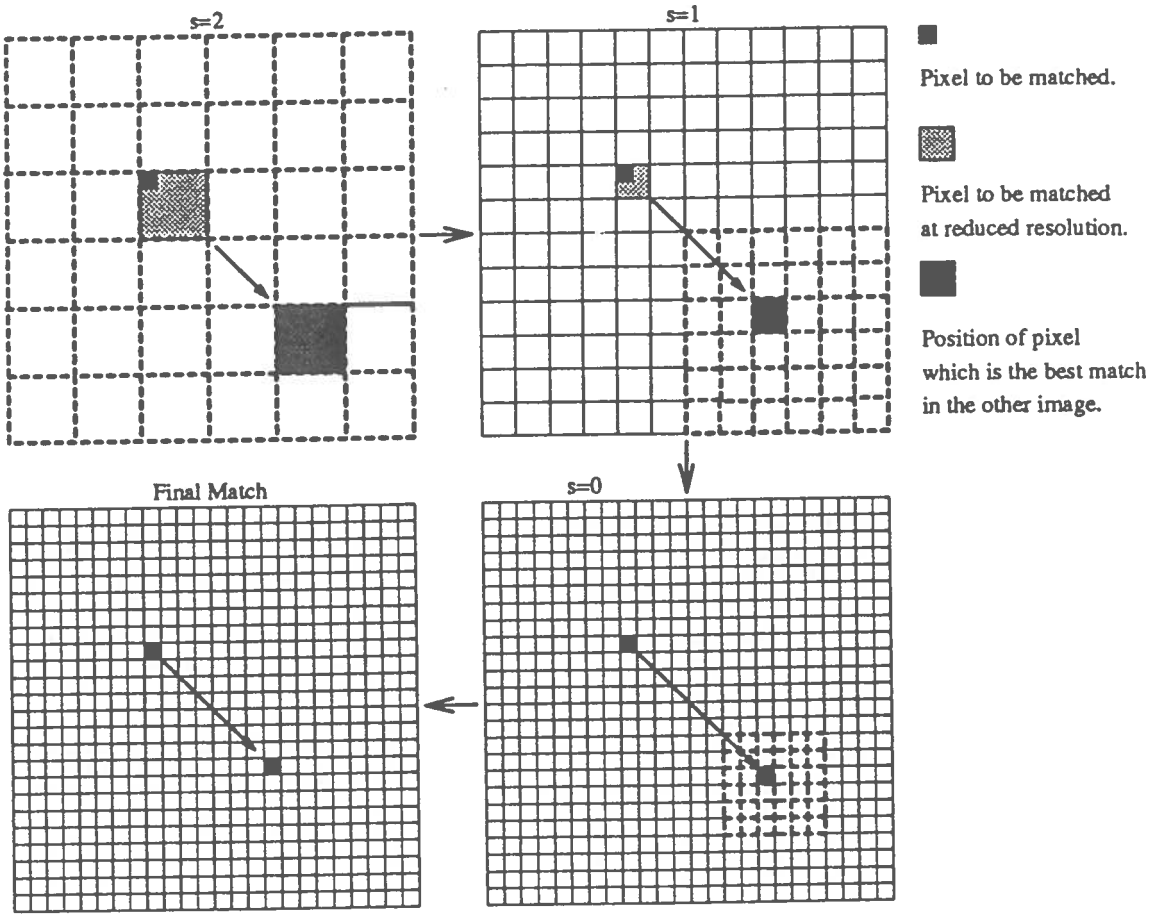

Figure 8–1: Hierarchical Search

The maximum reduction that can be achieved by this method is to reduce the final search region to 6^d. At each level of resolution there can be an expected error in a match of ± 1 pixel in each dimension (eg. x, y and z). At the next level of resolution this is equivalent to ± 2 in each dimension. The best match after $s + 1$ reductions will expand to 2 pixels in each dimension at the finer level of resolution, s reductions, so the minimum search region after reducing an image s times, having made matches on the image after $s + 1$ reductions, is 2 ± 2 pixels in each dimension, which is a region of size 6^d pixels.

If the problem of Eq. 8.2 is turned into a hierarchical search the maximum number of reductions will be $s = 2$ ($r = 25 \rightarrow 12.5 \rightarrow 6.25$) and this gives $S = 420,000$, a reduction of $\approx \frac{1}{12}$.

Edge effects have not been taken into account in this analysis, but as n becomes large the minimum search space will tend to that in Equation 8.8.

A hierarchical search could be expected to be faster than a basic region search by a factor of:

$$\text{t-factor} = \frac{1 - 2^{-d}}{2^{d \times s} \times (1 - 2^{-d \times (s+1)})} \times \frac{100}{p} \tag{8.9}$$

Eq. 8.9 is derived by dividing size of the Basic Region Search space by the Hierarchical Search space. No account has been taken of the possible computational overheads accrued in executing a hierarchical search.

The time taken to make 117^2 matches in an image 128^2, searching over a region of size 25^2, running on an HP9000/300 series with a maths coprocessor, averaged 8hrs. Of the 117^2 pixels only 49% contained relevant data so an immediate reduction in time to 4hrs (2 times faster) could be expected. If a hierarchical search were to be executed the expected running time would reduce to approximately 20 minutes (12 times faster).

Extending the above calculations to 3D, the timing estimates are:

Basic Region Search 2 years 11 months.

Relevant Data Only 1 yr 5 months. (2 times faster)

Hierarchical Search 228 hrs / 9 days 12 hrs. (112 times faster)

As can be seen, even with the quickest hierarchical search this technique would still take an unacceptable amount of time to process 3D data. The hierarchical search described is the fastest search that can be undertaken if matches for all relevant data points are to be found individually. To improve the operating time further the implementation of the algorithms must be investigated.