DOKUZ EYLÜL UNIVERSITY

GRADUATE SCHOOL OF NATURAL AND APPLIED SCIENCES

AUTOMATIC REGISTRATION OF MEDICAL IMAGE VOLUMES OF DIFFERENT MODALITIES

by

Umut DENİZ

September, 2011

İZMİR

AUTOMATIC REGISTRATION OF MEDICAL IMAGE VOLUMES OF DIFFERENT MODALITIES

A Thesis Submitted to the
Graduate School of Natural and Applied Sciences of Dokuz Eylül University
In Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Electrical&Electronics Engineering, Electrical&Electronics Program

Umut DENİZ

September, 2011
İZMİR

Ph.D. THESIS EXAMINATION RESULT FORM

We have read the thesis entitled **"AUTOMATIC REGISTRATION OF MEDICAL IMAGE VOLUMES OF DIFFERENT MODALITIES"** completed by **UMUT DENİZ** under supervision of **ASST. PROF. DR. HALDUN SARNEL** we certify that in our opinion it is fully adequate, in scope and in quality, as a thesis for the degree of Doctor of Philosophy.

Asst. Prof. Dr. Haldun SARNEL

Supervisor

Asst. Prof. Dr. Metehan MAKİNACI

Thesis Committee Member

Asst. Prof. Dr. Adil ALPKOÇAK

Thesis Committee Member

Prof. Dr. Tuncay ERTAŞ

Examining Committee Member

Asst. Prof. Dr. Güleser K. Demir

Examining Committee Member

Prof.Dr. Mustafa SABUNCU

Director

Graduate School of Natural and Applied Sciences

ACKNOWLEDGEMENTS

I would like to thank to my supervisor, Asst. Prof. Dr. Haldun Sarnel, for his undying patience, confidence, encouragment and valuable suggestions.

I would like to thank my committee members, Asst. Prof. Dr. Metehan Makinacı and Asst. Prof. Dr. Adil Alpkoçak, for their helpful comments and talents to guide me trough the doctoral program.

I also would like to thank my colleagues, Güven İpekoğlu, Şebnem Seçkin Uğurlu, Neslehan Avcu and Buket Turan for their supports and encouragments. I appriciate for my dear friends', Merve Özdemir, Yeliz and Durukan Çolak, valuable and helpful emotional supports.

Finally, I would like to thank my family for their patience and sensibilities.

Umut DENİZ

AUTOMATIC REGISTRATION OF MEDICAL IMAGE VOLUMES OF DIFFERENT MODALITIES

ABSTRACT

As an important part of clinical knowledge, biomedical images have critical importance for researches and healthcare. With the advances in medical imaging technologies, multimodal images are playing a more and more important role in improving the quality and efficiency of healthcare. Multimodal biomedical volume image registration is becoming increasingly important in diagnosis, treatment planning, functional studies, computer-guided therapies, and in biomedical research.

Registration by maximization of normalized mutual information is an intensity based approach to register images, which is derived from principles of information theory. Because of its robustness and accuracy, mutual information has become common in intensity based multimodal volume registrations. However, the approach involves time consuming intensitive computations to perform geometrical transformations of a volume image and estimate mutual information between two images.

Based on the theory of mutual information, a fully automated, robust and fast registration method was developed in this study, which has high level accuracy in registering various sets of multimodal volume images of brain.

In this thesis, a novel method is introduced, which is named as 3-Slice Method, to reduce the computational burden of the geometrical transformations and normalized mutual information (NMI) cost function.

Proposed volume image registration method optimizes normalized mutual information measure using a very fast algorithm known as simultaneous perturbation stochastic approximation (SPSA) algorithm. SPSA is applied first time by this study in registering 3-D biomedical volume images. Proposed method uses only three

orthogonal central slices of the volumes to build joint histogram instead of entire volumes. Performance results show our method provides remarkable speed improvements over the methods using conventional mutual information computations and other optimization algorithms. In addition to these, an approach is suggested to calculate the SPSA gain parameters automatically, instead of experimental pre-registration trials.

Keywords: Medical image volume registration, multimodal image registration, normalized mutual information, SPSA optimization.

FARKLI MODALİTELERDEKİ HACİMSEL MEDİKAL GÖRÜNTÜLERİN OTOMATİK ÇAKIŞTIRILMASI

ÖZ

Klinik bilginin bir parçası olarak, biomedikal imgeler araştırmalar ve sağlık hizmetleri bakımından kritik önem taşımaktadır. Medikal görüntüleme teknolojilerinin gelişmesiyle beraber, çok-modaliteli görüntüler, sağlık hizmetlerinin kalite ve verimliliğinin geliştirilmesinde giderek artan önemde rol oynamaktadır. Çok-modaliteli hacimsel görüntü çakıştırması, teşhis, tedavi planlama, fonksiyonel çalışma, bilgisayar destekli terapilerde ve biyomedikal araştırmalarda giderek daha önemli hale gelmiştir.

Normalize ortak bilginin doruklaştırılması ile çakıştırma, enformasyon teorisi prensiplerinden türetilen yoğunluk tabanlı bir görüntü çakıştırma yöntemidir. Ortak bilgi, güçlü ve isabetli oluşundan dolayı, yoğunluk tabanlı çok modaliteli hacimsel çakıştırmalarda yaygın hale gelmiştir. Ancak, bu yaklaşım, hacimsel bir görüntünün geometrik dönüşümünü gerçeklemek ve iki görüntü arasındaki ortak bilgiyi hesaplamak için, zaman alan yoğun hesaplamalar içermektedir.

Ortak bilgi teorisi temel alınarak, çeşitli çok-modaliteli hacimsel beyin görüntü setlerinin çakıştırılmasında yüksek düzeyde isabete sahip, tamamen otomatik güçlü ve hızlı bir yöntem geliştirilmiştir.

Bu tezde, geometrik dönüşüm ve normalize ortak bilgi fonksiyonun hesaplanmasındaki külfetin azaltılması için 3-Slice isimli yeni bir yaklaşım tanıtılmıştır.

Önerilen hacimsel görüntü çakıştırma yöntemi, normalize ortak bilgi ölçütünü, SPSA olarak bilinen çok hızlı bir algoritma kullanarak optimize etmektedir. SPSA ilk kez bu çalışmada biyomedikal hacimsel görüntülerin çakıştırılmasında kullanılmıştır. Önerilen yöntem, birleşik histogramı oluşturmak için tüm hacimi

kullanmak yerine, sadece hacimdeki birbirine dik üç merkezi görüntü dilimini kullanmaktadır. Performans sonuçları göstermektedir ki, yöntemimiz klasik ortak bilgi hesaplamaları ve diğer optmizasyon algoritmalarına göre kayda değer bir hız artışı sağlamıştır. Bunlara ek olarak, SPSA kazanç parametrelerinin deneysel çakıştırma öncesi denemeleriyle hesaplanması yerine, bu parametrelerin, bır algoritma tarafindan otomatik olarak hesaplanmasına yönelik bir yaklaşım sunulmuştur.

Anahtar sözcükler: Hacimsel medikal görüntü çakıştırması, çok-modaliteli görüntü çakıştırması, normalize ortak bilgi, SPSA optimizasyonu.

CONTENTS

Page

CHAPTER ONE
INTRODUCTION

Image registration is the process of aligning images so that corresponding features can easily be related. The term is also used to mean aligning images with a computer model or aligning features in an image with locations in physical space. The images might accured with different sensor or the same sensor at different times. Image registration has applications in many fields, especially in medical imaging.

The past decades have seen remarkable developments in medical imaging technology. Universities and medical industry have made huge investments in inventing and developing the technology needed to acquire images from multiple imaging modalities. Medical images are widely used in healtcare and biomedical research. X-ray computed tomography (CT) images are sensitive to tissue density and atomic composition, and the x-ray attenuation coefficient and magnetic resonance images (MR) are related to proton density, relaxation times, flow, and other parameters. The introduction of contrast agents provides information on the patency and function of tabular structures such as blood vessels, as well as the state of the blood-brain barrier. In nuclear imaging with single-photon emission computed tomography (SPECT) and positron emission tomography (PET), radiopharmaceuticals introduced into the body allow delineation of functioning tissue and measurement of metabolic and pathophysiological processes. These and other imaging technoligies now provide rich sources of data on the physical properties and biological function of tissues.

Medical imaging about establishing shape, structure, size, and spatial relationships of anatomical structures within the patients, together, with spatial information about function and any pathology or the other abnormality. Establishing the correspondance of spatial information in medical images and equivalent structures in the body is fundamental to medical image interpretation and analysis.

In medical applications, aligning images of similar or different modalities is frequently required as a pre-processing step for many diagnosis, therapy planning, monitoring change, data-fusion and visualization tasks. Neurosurgery and orthopedic surgery are also other important application fields of image registration. If the registration process involves images emerged from different modalities it is defined as multimodal registration. Otherwise, the registration of images produced by the same modality is called unimodal. Once two images have been aligned on computer after determining the true transformation that relates coordinates of one image to those of the other, they can be fused into a single display by various techniques. This process of fusion provides the physician with important additional information often not available from a single image (or modality) alone. Registration of both 2D and 3D images, whether it is unimodal or multimodal, has been broadly researched in the medical imaging field.

The classes of registration methods, referred to as intrinsic, only use natural features that are extracted from the subject's anatomy. In contrast, the methods termed extrinsic require artificially introduced features for registration purposes. Extrinsic methods are usually invasive and involve fixing unnatural objects such as stereotactic frames or fiducial markers, while intrinsic methods are non-invasive and an be used retrospectively. Intrinsic methods match a set of corresponding anatomical landmarks (Maintz & Viergever, 1998), employ a set of structures extracted by segmentation (Banarjee, Mukherjee, & Majumdar, 1995), or be based on the entire content of images such as voxel intensities. Content-based methods are particularly draw concern since they can be fully automated but have generally a high computational load.

Registration parameters are usually determined using information from entire images in content-based volume image registration. One question that can arise is whether an optimization algorithm has to use entire volume data in the 3D-to-3D registration to obtain good results. Actually, we know from the 3D-to-2D registration that a single slice from one modality can be successfully registered against a volume from other modality.

In many clinical scenarios, image from several modalities may be acquired and the diagnostician's task is to mentally combine or fuse this information to draw useful clinical conclusions. This generally requires mental compansation for changes in subject position. Image registration aligns the images and so establishes the correspondance between different features seen on different imaging modalities, allows monitoring of subtle changes in size or intensity over time or across a population, and establishes correspondances between images and physical space in image guided interventions. Registration of an atlas or computer model aids in the delineation of anatomical and pathological structures in medical iamges and important precursor to detailed analysis.

It is now common for patients to be imaged multiple times, either by repeated imaging with a single modality, or by imaging with different modalities. It is also common for patients to be imaged dynamically, that is to have sequences of images acquired, often at many frames per second. The ever increasing amount of image data acquired makes it more and more desirable to relate one image to another to assist in exracting relevant clinical information. Image registration can help in this task: intermodality registration enables the combination of complemantary information from different modalities, and intramodality registration enables accurate comparisons between images from same modality.

International concern about escalating healtcare costs drives development of methods that make the best possible use of medical images and, once again, image registration can help. However, medical image registration does not just enable better use of images that would be acquired anyway, it also opens up new applications for medical images. These include serial imaging to monitor subtle changes due to disease progression or treatment; perfusion or other functional studies when the subject can not be relied upon to remain in a fixed position during the dynamic acqusition. Therefore image-guided interventions, in which images acquired prior to the intervention are registered with treatment device, enabling the surgeon or interventionalist to use the preintervention images to guide his/her work. Image registration has also become a valuable technique for biomedical research, especially

in neuroscience, where imaging studies are making substantial contributions to understanding of the way the brain works.

Subjective judgements of the relative size, shape, and spatial relationships of visible structures and physiology inferred from intensity distributions are used for developing a diagnosis, planning therapy, and monitoring disease progression or response to therapy. A key process when interpreting these images together is the explicit or implicit establishment of correspondance between different points in the images. The spatial integrity of the images can allow very accurate correspondance to be determined. Once correspondance has been established in a verifiable way, multiple images can be interpreted as single unified data sets and conclusion drawn with increased confidence. Creating this single unified dataset is the process of fusion. In many instances, new information becomes available that could not have been deduced from inspection of individual images in loose association with one another.

In this thesis, considering the needs and trends in biomedical image registration field, multimodal medical images, optimization methods and normalized mutual information similarity measure, are studied to propose automated robust, fast and accurate approximations for 3D medical images of different modalities.

1.1 Organization of Thesis

The thesis consists of six chapters. Chapter 1 states the importance and necessity of image registration; outlines the motivation and the objective of the thesis.

Chapter 2 introduces image registration methods in detail with categorization and gives theoretical background of these categories.

Chapter 3 describes the information theory in detail and gives previous works on mutual information based registration.

In Chapter 4, the methods, which are used in the thesis, are described in four categories: preprocess, transformer, measure and optimizer.

Chapter 5 introduces the registration results which consist 34 data sets including phantom MR and clinical PET, CT and MR volume images and gives discussions on results that supported with figures and tables.

Finally, Chapter 6 gives conclusions and contributions of the thesis and recommendatitons for future works.

CHAPTER TWO
IMAGE REGISTRATION

Registration is the determination of a geometrical transformation that aligns points in one view of an object with corresponding points in another view of that object or another object. We use the term "view" generically to include a three-dimensional image, a two-dimensional image or the physical arrangement of an object in space. Three-dimensional images are acquired by tomographic modalities, such as computed tomography (CT), magnetic resonance (MR) imaging, single-photon emission computed tomography (SPECT) and positron emission tomography (PET). In each of these modalities, a contiguous set of two-dimensional slices provides a three-dimensional array of image intensity values. Typical two-dimensional images may be x-ray projections captured on film or as a digital radiograph or projections of visible light captured as a photograph or a video frame. In all cases, we are concerned primarily with digital images stored as discrete arrays of intensity values. In medical applications, which are our focus, the object in each view will be head region of the body. The two views are typically acquired from the same patient, in which case the problem is that of intrapatient registration, but interpatient registration has application as well.

2.1 Operational Goal of Registration

From an operational view, the inputs of registration are the two views to be registered; the output is a geometrical transformation, which is merely a mathematical mapping from points in one view to points in the second. To the extent that corresponding points are mapped together, the registration is successful. The determination of the correspondence is a problem specific to the domain of objects being imaged, which is, in our case, the human anatomy. To make the registration beneficial in medical diagnosis or treatment, the mapping that it produces must be applied in some clinically meaningful way by a system

that may will typically include registration as a subsystem. The larger system may combine the two registered images by producing a reoriented version of one view that can be "fused" with the other. This fusing of two views into one may be accomplished by simply summing intensity values in two images, by imposing outlines from one view over the gray levels of the other or by encoding one image in hue and the other in brightness in a color image. Regardless of the method employed, image fusion should be distinguished from image registration, which is a necessary first step before fusion can be successful.

The larger system may alternatively use the registration simply to provide a pair of movable cursors on two electronically displayed views linked via the registering transformation so that the cursors are constrained to visit corresponding points. This latter method generalizes easily to the case in which one view is the physical patient and one of the movable "cursors" is a physical pointer held by the surgeon. The registration system may be part of a robotically controlled treatment system whose guidance is based on registration between an image and the physical anatomy. Drills, for example, may be driven robotically through bone by following a path determined in CT and registered to the physical bone. Gamma rays produced by a linear accelerator or by radioactive isotopes may be aimed at tissue that is visible in MR but hidden from view during treatment with the aiming being accomplished via automatic calculations based on a registering transformation. Registration also serves as a first step in multimodal segmentation algorithms that incorporate information from two or more images in determining tissue types. Fusion, linked cursors, robotic controls, and multimodal segmentation algorithms exploit knowledge of a geometrical relationship between two registered views in order to assist in diagnosis or treatment. Registration is merely the determination of that relationship. The goal of registration is thus simply to produce as output a geometrical transformation that aligns corresponding points and can serve as input to a system further along in the chain from image acquisition to patient benefit.

2.2 Classification of Registration Methods

There are many image registration methods and they may be classified in many ways (Fitzpatrick, Hill & Maurer, 2000; Maintz & Viergever, 1998; Maurer & Fitzpatrick, 1993; Pluim, Maintz & Viergever, 2003; Shams, Sadeghi, Kennedy & Hartley, 2010; Van den Elsen, 1993; Zitová & Flusser, 2003;). Maintz has suggested a nine-dimensional scheme that provides an excellent categorization (Maintz et al., 1998). Fitzpatrick et al. (2000) condense it slightly into the following eight categories: image dimensionality, registration basis, geometrical transformation, degree of interaction, optimization procedure, modalities, subject, and object.

"Image dimensionality" refers to the number of geometrical dimensions of the image spaces involved, which in medical applications are typically three-dimensional but sometimes two-dimensional. The "registration basis" is the aspect of the two views used to effect the registration. For example, the registration might be based on a given set of point pairs that are known to correspond or the basis might be a set of corresponding surface pairs. Other loci might be used as well, including lines or planes (a special case of surfaces). In some cases, these correspondences are derived from objects that have been attached to the anatomy expressly to facilitate registration. Such objects include, for example, the stereotactic frame and point-like markers, each of which have components designed to be clearly visible in specific imaging modalities. Registration methods that are based on such attachments are termed "prospective" or "extrinsic" methods and are in contrast with the so-called "retrospective" or "intrinsic" methods, which rely on anatomic features only. Alternatively, there may be no known correspondences as input. In that case, intensity patterns in the two views will be matched, a basis that we call "intensity". The category, "geometrical transformation", is a combination of two of Maintz's categories, the "nature of transformation" and the "domain of transformation". It refers to the mathematical form of the geometrical mapping used to align points in one space with those in the other. "Degree of interaction" refers to the control exerted by a human operator over the registration algorithm. The interaction may consist simply of the initialization of certain parameters, or it may involve

adjustments throughout the registration process in response to visual assessment of the alignment or to other measures of intermediate registration success. The ideal situation is the fully automatic algorithm, which requires no interaction. “Optimization procedure” refers to the standard approach in algorithmic registration in which the quality of the registration is estimated continually during the registration procedure in terms of some function of the images and the mapping between them. The optimization procedure is the method, possibly including some degree of interaction, by which that function is maximized or minimized. The ideal situation here is a closed-form solution which is guaranteed to produce the global extremum. The more common situation is that in which a global extremum is sought among many local ones by means of iterative search.

“Modalities” refers to the means by which the images to be registered are acquired. Registration methods designed for like modalities are typically distinct from those appropriate for differing modalities. Registration between like modalities, such as MR-MR, is called “intramodal” or “monomodal” registration; registration between differing modalities, such as MR-PET, is called “intermodal” or “multimodal” registration. “Subject” refers to patient involvement and comprises three subcategories: intrapatient, interpatient, and atlas, the latter category comprising registrations between patients and atlases, which are themselves typically derived from patient images (Collins et al., 1998). “Object” refers to the particular region of anatomy to be registered (e.g., head, liver, vertebra).

To build a registration hierarchy based on these eight categorizations, one categorization must be placed at the top level, which, in the organization of this chapter, is the registration basis. The three categories of registration basis mentioned above are examined in three major sections below: Point-based methods, Surface-based methods, and Intensity-based methods. Before the discussions of the basis for registration, in the next section the category of geometrical transformation will be presented.

2.2.1 Geometrical Transformations

Each view that is involved in a registration will be referred to a coordinate system, which defines a space for that view. Our definition of registration is based on geometrical transformations, which are mappings of points from the space X of one view to the space Y of a second view. The transformation T applied to a point in X represented by the column vector x produces a transformed point $\boldsymbol{x}'$,

$$\boldsymbol{x}' = T(\boldsymbol{x}) \tag{2.1}$$

If the point y in Y corresponds to x, then a successful registration will make x' equal, or approximately equal, to y. Any nonzero displacement $T(x)$-y is a registration error. The set of all possible T may be partitioned into rigid and nonrigid transformations with the latter transformations further divided into many subsets. This top-level division makes sense in general because of the ubiquity of rigid, or approximately rigid, objects in the world. It makes sense for medical applications in particular because of the rigid behavior of many parts of the body, notably the bones and the contents of the head (not during surgery). It is also a simple class with only six parameters completely specifying a rigid transformation in three dimensions. (We note here that, while one and two-dimensional motion is possible, such limited motion is sufficiently rare that it will be ignored.)

2.2.1.1 Rigid Transformations

Rigid transformations, or rigid mappings, are defined as geometrical transformations that preserve all distances. These transformations also preserve the straightness of lines (and the planarity of surfaces) and all nonzero angles between straight lines. Registration problems that are limited to rigid transformations are called rigid registration problems. Rigid transformations are simple to specify, and there are several methods of doing so. In each method, there are two components to the specification, a translation and a rotation. The translation is a three-dimensional vector t that may be specified by giving its three coordinates t_x, t_y, t_z relative to a set

of *x, y, z* cartesian axes or by giving its length and two angles to specify its direction in polar spherical coordinates. There are many ways of specifying the rotational component, among them Euler angles, Cayley-Klein parameters, quaternions, axis and angle, and orthogonal matrices (Fu, Gonzales & Lee, 1987; Goldstein, 1950; Horn, 1986, 1987; Rosenfield, 1959). In our discussions we will utilize orthogonal matrices. With this approach, if *T* is rigid, then

$$\boldsymbol{x}' = R\boldsymbol{x} + \boldsymbol{t} \tag{2.2}$$

Where *R* is a 3x3 orthogonal matrix, meaning that $R^t R = R R^t = I$ (the identity matrix). Thus $R^{-1}=R^t$. This class of matrices includes both the proper rotations, which describe physical transformations of rigid objects, and improper rotations, which do not. These latter transformations both rotate and reflect rigid objects, so that, for example, a right-handed glove becomes a left-handed one. Improper rotations can be eliminated by requiring *det(R) = +1*.

Proper rotations can be parameterized in terms of three angles of rotation, Θ_x, Θ_y, Θ_z, about the respective cartesian axes, the so-called "Euler angles". The rotation angle about a given axis is, with rare exception, considered positive if the rotation about the axis appears clockwise as viewed from the origin while looking in the positive direction along the axis. The rotation of an object (as opposed to the coordinate system to which it is referred) about the *x*, *y* and *z* axes, in that order leads to

$$R = \begin{bmatrix} cos\theta_z & -sin\theta_z & 0 \\ sin\theta_z & cos\theta_z & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} cos\theta_y & 0 & sin\theta_y \\ 0 & 1 & 0 \\ -sin\theta_y & 0 & cos\theta_y \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & cos\theta_x & -sin\theta_x \\ 0 & sin\theta_x & cos\theta_x \end{bmatrix}$$

$$= \begin{bmatrix} cos\theta_y cos\theta_z & -cos\theta_x sin\theta_z + sin\theta_x sin\theta_y cos\theta_z & sin\theta_x sin\theta_z + cos\theta_x sin\theta_y cos\theta_z \\ cos\theta_y sin\theta_z & cos\theta_x cos\theta_z + sin\theta_x sin\theta_y sin\theta_z & -sin\theta_x cos\theta_z + cos\theta_x sin\theta_y sin\theta_z \\ -sin\theta_y & sin\theta_x cos\theta_y & cos\theta_z cos\theta_y \end{bmatrix}$$

with the three matrices in the first line representing the rotations $R_z(\Theta_z)$, $R_y(\Theta_y)$ and $R_x(\Theta_x)$ about *z*, *y*, and *x*, respectively (in reverse order because they are applied

from right to left). Other angular parameterizations are sometimes used, including all permutations of the order of R_x, R_y and R_z. General rotations can also be produced by three rotations about only two of the cartesian axes provided that successive rotations are about distinct axes. The most common of these is $R = R_z(\Theta_{z2})R_x(\Theta_x)R_z(\Theta_{z1})$.

2.2.1.2 Nonrigid Transformations

Nonrigid transformations are important not only for applications to nonrigid anatomy, but also for interpatient registration of rigid anatomy and intrapatient registration of rigid anatomy when there are nonrigid distortions in the image acquisition procedure. In all cases, it is preferable to choose transformations that have physical meaning, but in some cases, the choice is made on the basis of convenient mathematical properties.

2.2.1.2.1 Scaling Transformations. The simplest nonrigid transformations are rigid except for scaling,

$$\boldsymbol{x}' = RS\boldsymbol{x} + \boldsymbol{t} \tag{2.3}$$

and

$$\boldsymbol{x}' = SR\boldsymbol{x} + \boldsymbol{t} \tag{2.4}$$

where $S = diag(s_x, s_y, s_z)$ is a diagonal matrix whose elements represent scale factors along the three coordinate axes. Because RS is not in general equal to SR, these equations represent two different classes of transformations. Such transformations may be needed to compensate for calibration errors in image acquisition systems. They are appropriate, for example, when gradient strengths are in error in MR. The diagonal elements of S then become the respective correction factors for the x, y and z gradients. If the scaling is isotropic, the transformation is a similarity transformation,

$$\boldsymbol{x}' = sR\boldsymbol{x} + \boldsymbol{t} \tag{2.5}$$

where s is a positive scalar, sometimes known as a "dilation" (for values less than one as well). This transformation preserves the straightness of lines and the angles between them. Both Eq. (2.3) and Eq. (2.4) reduce to Eq. (2.5) when $s_x = s_y = s_z = s$. The coupling of scaling with the rigid transformation is effective when registrations must account for erroneous or unknown scales in the image acquisition process.

2.2.1.2.2 Affine Transformations. The scaling transformations are special cases of the more general affine transformation,

$$\boldsymbol{x}' = A\boldsymbol{x} + \boldsymbol{t} \tag{2.6}$$

in which there is no restriction on the elements a_{ij} of the matrix A. The affine transformation preserves the straightness of lines, and hence, the planarity of surfaces, and it preserves parallelism, but it allows angles between lines to change. It is an appropriate transformation class when the image may have been skewed during acquisition as, for example, when the CT gantry angle is incorrectly recorded.

The affine transformations and their associated special cases are sometimes represented by means of homogeneous coordinates. In this representation, both A and $\boldsymbol{t}$ are folded into one 4x4 matrix M whose elements are defined as:

$m_{ij}=a_{ij}$, $m_{i4}=t_i$, $m_{4j}=0$ and $m_{44}=1$ where $i=1, 2, 3$ and $j=1, 2, 3$.

To accomplish the transformation, augmented vectors $\boldsymbol{u}$ and $\boldsymbol{u'}$ are used for which $u_i = x_i$ and $u_i' = x_i'$ for $i=1, 2, 3$ and $u_4 = u_4' = 1$.

$$\boldsymbol{u}' = \begin{bmatrix} u_1' \\ u_2' \\ u_3' \\ 1 \end{bmatrix} = M\boldsymbol{u} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & t_1 \\ a_{21} & a_{22} & a_{23} & t_2 \\ a_{31} & a_{32} & a_{33} & t_3 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ 1 \end{bmatrix} \tag{2.7}$$

While the use of homogeneous coordinates does not produce any extra power or generality for rigid transformations, it does simplify notation, especially when rigid transformations must be combined with projective transformations.

2.2.1.2.3 Projective Transformations. The nonrigid transformations that we have considered, all of which are affine transformations, preserve parallelism. The more general nonrigid transformations include the projective transformations, which preserve the straightness of lines and planarity of surfaces, and the curved transformations, which do not. The projective transformations, which have the form,

$$\boldsymbol{x}' = (A\boldsymbol{x} + \boldsymbol{t})/(\boldsymbol{p} \cdot \boldsymbol{x} + \propto) \tag{2.8}$$

can be written simply in homogeneous coordinates,

$$\boldsymbol{u}' = \begin{bmatrix} u_1' \\ u_2' \\ u_3' \\ u_4' \end{bmatrix} = M\boldsymbol{u} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & t_1 \\ a_{21} & a_{22} & a_{23} & t_2 \\ a_{31} & a_{32} & a_{33} & t_3 \\ p_1 & p_2 & p_3 & \propto \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ 1 \end{bmatrix} \tag{2.9}$$

where, as for the affine transformation, $u_i = x_i$, $u_4=1$ and $m_{ij}=a_{ij}$ for $i=1, 2, 3$ and $j=1, 2, 3$ but u_4' is no longer necessarily equal to 1, $m_{4j}=p_j$, $m_{44}=\alpha$ and $x_i' = u_i'/u_4'$. The linearity of Eq. (2.9) can provide considerable simplification for the projective transformations and the perspective projections.

2.2.1.2.4 Perspective Transformations. Images obtained by x-ray projection, endoscopy, laparoscopy, microscopy, and direct video acquisition are all two-dimensional views of three-dimensional objects rendered by means of projecting light rays or x-rays from a three-dimensional scene onto a two-dimensional plane. The geometrical transformation, which is called a perspective projection, produced by each of these modalities is equivalent to that of photography.

These perspective projections are a subset of the projective transformations of Eqs. (2.8) and (2.9). The projective transformations, unlike the perspective

projections, do not, in general, transform $\boldsymbol{x}$ to a plane. Furthermore, the affine portion of the transformation is typically assumed to be the identity for perspective projections. Specializing now to perspective projections, we let $f = 1/|\boldsymbol{p}|$ in Eq. (2.8), and let $\boldsymbol{\acute{p}}$ be a unit vector in the direction of the projection axis, $\boldsymbol{p}$. These substitutions lead to

$$\boldsymbol{x}' = f\boldsymbol{x}/(\boldsymbol{x} \cdot \boldsymbol{p}' + \alpha f) \tag{2.10}$$

If α is nonzero, then Eq. (2.10) does not, in fact, transform $\boldsymbol{x}$ to a plane and, hence, is not a perspective projection. Perspective projection can be produced, however, by zeroing the component of $\boldsymbol{x}'$ in the direction of $\boldsymbol{p}$:

$$\boldsymbol{x}' \rightarrow \boldsymbol{x}' - (\boldsymbol{x}' \cdot \boldsymbol{p}')\boldsymbol{p}' \tag{2.11}$$

Equation (2.10) and substitution (2.11) give the general form of the transformation produced when a photograph of a three-dimensional scene is acquired with a "pinhole camera," which is a camera in which a small hole substitutes for the lens system. A ray from a point $\boldsymbol{x}$ in the scene is projected through the pinhole onto a film screen, which is perpendicular to the axis of projection $\boldsymbol{p}$ and located at a distance $\boldsymbol{f}$ from the pinhole. Fortunately, all the systems mentioned above can be approximated by the pinhole camera system by identifying the unique point relative to the lens system through which light rays travel undeflected, or the point from which the x rays emanate. That point, also known as the "center of perspectivity" (Haralick & Shapiro, 1993), plays the role of the pinhole. Because the film image is inverted, it is convenient to treat instead an equivalent upright "image" located in front of the camera, which is also perpendicular to p and located the same distance $\boldsymbol{f}$ from the pinhole. The transformed point, $\boldsymbol{x}'$, of Eq. (2.10), followed by substitution Eq. (2.11) lies in that plane.

The parameter $\boldsymbol{f}$ is called the focal length, or, alternatively, the camera constant or principal distance. The name focal length is derived from lens systems. It is meant to imply that the lens is adjusted so that all light emanating from a given point in front

of the camera and passing through its lens will be focused to a single point on a screen located at that distance from the effective pinhole. The focusing is only approximate and varies in quality with the distance of the anatomy from the lens. The value of α in Eq. (2.10) is determined by the placement of the origin. Typically, the origin is placed at the pinhole, for which $\alpha=0$, or at the intersection of $\boldsymbol{p}$ and the image plane, for which $\alpha=1$.

2.2.1.2.5 Curved Transformations. Curved transformations are those that do not preserve the straightness of lines. In curved transformations, the simplest functional form for T is a polynomial in the components of $\boldsymbol{x}$ (Goshtasby,1986; Maguire et al., 1991),

$$\boldsymbol{x}' = \sum_{ijk}^{IJK} \boldsymbol{c}_{ijk} x^i y^j z^k \tag{2.12}$$

where $\boldsymbol{c}_{ijk}$ is the three-element vector of coefficients for the i, j, k term in the polynomial expression for the three components x', y', z' of $\boldsymbol{x}'$.

Modifications may be employed that include all terms for which $i+j+k \leq M$. These transformations are rarely used with values of I, J and K greater than 2 or M greater than 5 because of spurious oscillations associated with high-order polynomials. These oscillations can be reduced or eliminated by employing piecewise polynomials. The resulting transformations are defined by first partitioning the space into a set of three-dimensional rectangles by means of three sets of cut planes, each perpendicular to one of the cartesian axes.

A transformation that has been heavily used for two-dimensional problems is the thin-plate spline, which was originally called the surface spline. The thin-plate splines were first employed to describe deformations within the two-dimensional plane by Goshtasby (1988). Goshtasby's formulation, which is now widely employed in the image processing literature, is as follows:

$$\boldsymbol{x}' = A\boldsymbol{x} + \boldsymbol{t} + \sum_i^N \boldsymbol{c}_i r_i^2 \ln r_i^2 \tag{2.13}$$

where $r_i = |\boldsymbol{x}\text{-}\boldsymbol{x_i}|$ and $\boldsymbol{x_i}$ is a control point. Unlike the rectangular grid of knots required for the cubic-splines, the control points can be placed arbitrarily, a feature that is of considerable advantage in the registration of medical images. For three-dimensional transformations, the thin-plate spline has a simpler form in which $\boldsymbol{c}_i r_i^2 \ln r_i^2$ in Eq. (2.13) is replaced by r_i. For both the two-dimensional and the three-dimensional forms, the affine portion of Eq. (2.13) is a necessary part of the transformation. Without this component there may be no set of $\boldsymbol{c_i}$ that satisfies the equation at all $\boldsymbol{N}$ points. With the affine part included, it is always possible, by means of the imposition of a set of side conditions on the $\boldsymbol{c_i}$, to ensure that a solution exists for any arrangement of points.

Other curved transformations have been employed, including solutions to the equations of continuum mechanics describing elastic and fluid properties attributed to the anatomy being registered (Bajcsy & Kovacic, 1989; Broit, 1981; Chirstensen, Rabbitt & Miller, 1996; Davis, Khotanzad, Flamig & Harms, 1997; Gee, Reivic & Bajcsy, 1993; Miga et al., 1999; Thompson & Toga, 1996). These equations, which are derived from conservation of mass, momentum, and energy and from experimentally measured material properties, involve the displacement vector, $\boldsymbol{x'} - \boldsymbol{x}$, and the first and second spatial derivatives of its components. The nonrigid transformations that result from the numerical solution of these partial differential equations are appropriate for intrapatient registration when the anatomy is nonrigid, especially when surgical resection has changed its shape. In these cases, the major problem in registration is the determination of the material properties and forces that act on the tissue. With that information available, the solution of the differential equations may be carried out numerically by means of finite-difference or finite-element methods. These transformations have been also used for interpatient registration and for the closely related problem of mapping an atlas to a patient.

2.2.2 Point-Based Methods

If some set of corresponding point pairs can be identified a priori for a given pair of views, then registration can be effected by selecting a transformation that aligns the points. Because such points are taken as being reliable for the purposes of registration, they are called fiducial points, or fiducials. To be reliable, they must lie in clearly discernible features, which we will call fiducial features. The determination of a precise point within a feature is called fiducial localization. The transformation that aligns the corresponding fiducial points will then interpolate the mapping from these points to other points in the views.

The fiducial localization process may be based on interactive visual identification of anatomical landmarks, such as the junction of two linear structures, e.g., the central sulcus with the midline of the brain or the intersection of a linear structure with a surface, e.g., the junction of septa in an air sinus, etc. (Hill et al, 1995). Alternatively, the feature may be a marker attached to the anatomy and designed to be accurately localizable by means of automatic algorithms. In either case, the chosen point will inevitably be erroneously displaced somewhat from its correct location. This displacement in the determination of the fiducial point associated with a fiducial feature is commonly called the fiducial localization error (FLE). Such errors will occur in both image spaces. They cannot ordinarily be observed directly, but they can be observed indirectly through the registration errors that they cause.

Marker-based registration has the considerable advantage over landmark-based registration that the fiducial feature is independent of anatomy. Automatic algorithms for locating fiducial markers can take advantage of knowledge of the marker's size and shape in order to produce a consistent fiducial point within it (Wang, Fitzpatrick & Maurer, 1995). Typically, the fiducial point chosen by a localization algorithm will lie near its center. Hence, the point is often referred to as the fiducial "centroid". However, registration accuracy depends only on the degree to which the chosen points correspond in the two views. Random errors in the localized position will be caused by noise in the image and by the varying positioning of the marker relative to

the voxel grid. For reasonable localization algorithms, the mean of the fiducial points chosen should be the same in the two views relative to a coordinate system fixed in the marker. Because of this consistency, the effective mean displacement, $\langle \boldsymbol{FLE} \rangle$, in each view is zero ($\langle x \rangle$ indicates the expected value of x). The variance $\langle FLE^2 \rangle$ may be appreciable, however. (Note that bold font indicates the vector displacement **FLE** and a normal font indicates the magnitude FLE.)

The goal of fiducial design and of the design of the accompanying fiducial localization algorithm is to produce a small variance. In general, as the marker volume becomes larger, and as the signal per volume produced in the scanner by its contents becomes larger, FLE will become smaller. Thus, larger markers tend to exhibit smaller FLEs. Brighter markers also have smaller FLEs because of the smaller contribution of image noise relative to marker intensity.

Any nonzero displacement $T(\boldsymbol{x})-\boldsymbol{y}$ between a transformed point $T(\boldsymbol{x})$ and its corresponding point $\boldsymbol{y}$ is a registration error. To the extent that FLE is small and that the form of the transformation correctly describes the motion of the object, the alignment of the fiducial points in the two views will lead to small registration errors for all points. If the transformation is selected from some constrained set (as, for example, the rigid transformations), then, because of FLE, it will ordinarily not be possible to achieve a perfect alignment of fiducials. The resultant misalignment may, in some cases, be used as feedback to assess whether or not the registration is successful. A common measure of overall fiducial misalignment, is the root-mean-square (RMS) error. This error, which called the fiducial registration error, or FRE, is defined as follows. First, an individual fiducial registration error ($\mathbf{FRE}_i$) should be defined,

$$\boldsymbol{FRE}_i = T(\boldsymbol{x}_i) - \boldsymbol{y}_i \tag{2.14}$$

where $\boldsymbol{x}_i$ and $\boldsymbol{y}_i$ are the corresponding fiducial points in views X and Y, respectively, belonging to feature i. Then FRE can be defined in terms of the magnitudes of the $\mathbf{FRE_i}$,

$$FRE^2 = (\frac{1}{N}) \sum_i^N w_i^2 FRE_i^2 \quad (2.15)$$

Where *N* is the number of fiducial features used in the registration and w_i^2 is a non-negative weighting factor, which may be used to decrease the influence of less reliable fiducials. For example, if $\langle FLE_i^2 \rangle$ is the expected squared fiducial localization error for fiducial *i*, then it may be choosen to set $w_i^2 = 1/\langle FLE_i^2 \rangle$, where FLE_i is the fiducial localization error for fiducial *i*.

2.2.3 Surface-Based Methods

The 3-D boundary surface of an anatomic object or structure is an intuitive and easily characterized geometrical feature that can be used for medical image registration. Surface-based image registration methods involve determining corresponding surfaces in different images (and/or physical space) and computing the transformation that best aligns these surfaces.

The skin boundary surface (air-skin interface) and the outer cranial surface are obvious choices that have frequently been used for both image-to-image (e.g., CT-MR, serial MR) and image-to-physical registration of head images. The surface representation can be simply a point set (i.e., a collection of points on the surface), a faceted surface (e.g., triangle set), an implicit surface, or a parametric surface (e.g., B-spline surface). Extraction of a surface such as the skin or bone is relatively easy and fairly automatic for head CT and MR images. Extraction of many soft tissue boundary surfaces is generally more difficult and less automatic. Image segmentation algorithms can generate 2-D contours in contiguous image slices that are linked together to form a 3-D surface, or they can generate 3-D surfaces directly from the image volume. In physical space, skin surface points can be easily determined using laser range finders; stereo video systems; and articulated mechanical, magnetic, active and passive optical, and ultrasonic 3-D localizers. Bone surface points can be found using tracked A-mode (Maurer et al, 1999) and B-mode (Lavallee et al, 1996) ultrasound probes. The computer vision sensors, 3-D localizers, and tracked A-

mode ultrasound probes produce surface point sets. Tracked B-mode probes produce a set of 2-D images (or a single compounded 3-D image) from which bone surface points need to be segmented.

Surfaces can provide basic features for both rigid-body and nonrigid registration. A central and difficult question that must be addressed by any nonrigid surface-based registration algorithm is how deformation of the contents of an object is related to deformation of the surface of the object. Most of the surface-based registration algorithms that have been reported are concerned with rigid-body transformation, occasionally with isotropic or nonisotropic scaling.

2.2.3.1 Disparity Functions

The approach for solving the surface-based registration problem that is frequently used in more recent computer vision literature (where it is often called the free-form surface matching problem). That is normally used in the medical image processing community, is to search for the transformation that minimizes some disparity function or metric between the two surfaces *X* and *Y*. The disparity function is generally a distance. The disparity function normally used for surface-based image registration is an average, and optionally weighted, distance between points on one surface and corresponding points on the other surface. Let$\{\boldsymbol{x}_j\}$ for j=1,…,N_x be a set of N_x points on the surface *X*. The general approach is to search for the transformation that minimizes the disparity function,

$$d(T(X),Y)=\sqrt{\sum_{j=1}^{N_x} w_j^2 d^2\left(T(\boldsymbol{x}_j),Y\right)}=\sqrt{\sum_{j=1}^{N_x} w_j^2 \left\|\left(T(\boldsymbol{x}_j)-\boldsymbol{y}_j\right)\right\|^2} \quad (2.16)$$

where

$$\boldsymbol{y_j}=C\left(T(\boldsymbol{x_j}),Y\right) \quad (2.17)$$

is a point on the surface *Y* corresponding to the point $\boldsymbol{x}_j$, *C* is a correspondence

function (e.g., closest point operator), and $\{w_j\}$ is a set of weights associated with $\{\boldsymbol{x}_j\}$. The principal difference between point-based registration and surface-based registration is in the availability of point correspondence information. It is the lack of exact point correspondence information that causes surface-based registration algorithms to be based on iterative search. Eq.(2.17) merely provides approximate point correspondence information for a particular *T* during an iterative search.

The point set $\{\boldsymbol{x}_j\}$ and the surface *Y* have been called, respectively, the hat and head (Pelizzari et al, 1989), the dynamic and static feature sets (Zuk, Atkins & Booth, 1994) and the data point set and model surface shape (Besl & McKay, 1992). Typically one surface contains more information than the other. The surface from the image that covers the larger volume of the patient and/or has the highest resolution is generally picked as the model shape.

In surface-based registration (Eq. 2.16), statistical independence of errors is unlikely to point-based. The skin is a movable and deformable structure, and local deformations tend to be highly correlated. Physical space surface points acquired with a sensor can have biased error due to miscalibration. Nonetheless, weights can be useful to reduce the influence of less reliable surface points. Many sensors and tracking devices have less accuracy at the edges of the working volume. Weights could potentially be used to account for the sensitivity of the registration to the perturbation of individual surface points (Simon & Kanade, 1997). Weights can be used to account for nonuniform sampling density (Maurer et al., 1996). Finally, weights can also be used to deal with outliers that can arise from nonoverlapping sections of surfaces, poor segmentation, and erroneous sensor data (Maurer et al., 1996; Maurer, Maciunas & Fitzpatrick, 1998).

2.2.3.2 Head and Hat Algorithm

The first investigators to apply surface-based registration to a medical problem were Pelizzari, Chen, and colleagues (Pelizzari et al., 1989). They used their head and hat algorithm to register CT, MR, and PET images of the head. The hat is a skin

surface point set $\{x_j\}$.

The head is a polygon set model of the skin surface Y created by segmenting contours in contiguous transverse image slices. They define y_j as the intersection with the head Y of a line joining the transformed hat point $T(x_j)$ and the centroid of the head Y. The intersection is efficiently calculated by reducing the 3-D line-polyhedron intersection problem to a 2-D line-polygon intersection problem. The transformation T that minimizes Eq. (2.16) is found using a standard gradient descent technique. The major limitations of this technique are due to the particular distance used, the distance from the surface point to the surface intersection along a line passing through the surface centroid. This definition of distance requires that the surface be approximately spherical. It also requires that a good initial transformation be supplied as input to the transformation parameter search. Finally, it is probably related to the observation by the authors and others that the search frequently terminates in local minima and thus requires substantial user interaction.

2.2.3.3 Distance Transform Approach

The calculation of point-to-surface distance is computationally intensive, even when using special data structures and other optimizations. A computationally efficient alternative is to use a distance transform (DT). A DT of a binary image I is an assignment to each voxel v of the distance between v and the closest feature voxel in I. A DT of a binary image where the feature voxels are surface voxels is a gray-level image in which each voxel v has a value that is the distance from the center of v to the center of the nearest surface voxel. Thus a DT provides a method for precomputing and storing point-to-surface distance. Normally squared distance is stored. Then, at each step of an iterative transformation parameter search, the value of the disparity function in Eq. (2.16) is computed simply by summing the values of the voxels in the squared distance image that contain the transformed points $\{T(x_j)\}$. One limitation of this approach is that a DT is spatially quantized, i.e., a DT image contains exact point-to-surface distance only at regularly spaced lattice points (centers of voxels). A slight improvement over using the distance at the nearest

lattice point can be achieved by using a trilinear interpolation of the distances at the nearest eight lattice points. Nonetheless, the surface is fundamentally represented by the point set consisting of the centers of all feature (surface) voxels, and thus subvoxel surface position information is lost. Spatial quantization might be the reason that registrations produced by surface-based methods using DTs have been reported to be considerably less accurate than registrations produced by surface-based methods not using DTs (West et al., 1997).

2.2.3.4 Iterative Closest Point Algorithm

All surface-based registration algorithms must search for the transformation T that minimizes the disparity function in Eq. (2.16) or a variation thereof. This is a general nonlinear minimization problem that is typically solved using one of the common gradient descent techniques (Press et al., 1992). The search will typically converge to, or very close to, the correct minimum of the disparity function minimum if the initial transformation is within about 20–30 degrees and 20–30 mm of the correct solution. To help minimize the possibility of the search getting stuck in a local minimum, many investigators perform the search in a hierarchical coarse-to-fine manner.

Besl & McKay (1992) presented an algorithm which reduces the general nonlinear minimization problem to an iterative point-based registration problem. Their iterative closest point (ICP) algorithm is a general-purpose, representation-independent, shape-based registration algorithm that can be used with a variety of geometrical primitives including point sets, line segment sets, triangle sets (faceted surfaces), and implicit and parametric curves and surfaces. One shape is assigned to be the data shape and the other shape to be the model shape. For surface-based registration, the shapes are surfaces. The data shape is decomposed into a point set (if it is not already in point set form). Then the data shape is registered to the model shape by iteratively finding model points closest to the data points, registering the two point sets and applying the resulting transformation to the data points.

2.2.4 Intensity-Based Methods

Image intensity is an alternative registration basis to points or surface features. It has recently become the most widely used registration basis for several important applications. In this context, the term intensity is invariably used to refer to the scalar values in image pixels or voxels. The physical meaning of the pixel or voxel value depends on the modalities being registered and is very often not a direct measure of optical power (the strict definition of intensity).

Intensity-based registration involves calculating a transformation between two images using the pixel or voxel values alone. In its purest form, the registration transformation is determined by iteratively optimizing some similarity measure calculated from all pixel or voxel values. Because of the predominance of three-dimensional images in medical imaging, we refer to these measures as voxel similarity measures. In practice, many intensity-based registration algorithms use only a subset of voxels and require some sort of preprocessing. For example, the algorithm may run faster if only a subset of voxels are used. This subset can be chosen on a regular grid, or be randomly chosen. It is normal in these circumstances to blur the images before sampling to avoid aliasing in the subsampled images, and the amount of blurring used may be application dependent. Alternatively, an algorithm may work reliably only if the similarity measure is calculated from the voxels in a defined region of interest in the image, rather than all voxels. In this case, some sort of pre-segmentation of the images is required, and this is likely to depend both on the modalities being registered and the part of the body being studied. In some other intensity-based algorithms, the similarity measures work on derived image parameters such as image gradients, rather than the original voxel values.

For retrospective registration, a major attraction of intensity-based algorithms is that the amount of preprocessing or user-interaction required is much less than for point-based or surface-based methods. As a consequence, these methods are relatively easy to automate. The need for preprocessing does, however, mean that many intensity-based algorithms are restricted to a quite limited range of images.

One of the aims of research in this area has been to devise general algorithms that will work on a wide variety of image types, without application-specific preprocessing.

Intensity-based registration algorithms can be used for a wide variety of applications: registering images with the same dimensionality, or different dimensionality; both rigid transformations and registration incorporating deformation; and both intermodality and intramodality images. Most algorithms are applicable to only a subset of these applications, but some are quite generally applicable.

Let the images to be registered are *A* and *B*. The sets of voxels in these images are $\{A(i)\}$ and $\{B(i)\}$ respectively. We will treat image *A* as a reference image, and *B* as an image that is iteratively transformed to $B' = T(B)$ by successive estimates of the registration transformation *T*. The transformation estimates will change the overlap between the images being registered. Voxel-similarity measures are invariably calculated for the set of voxels in the overlapping region of *A* and B' , i.e., within $A \cap B'$, which is a function of *T* and so changes as the algorithm iterates. For some voxel-similarity measures, information from the intensity histogram is used, so we need to refer directly to intensity values in the image, rather than index voxels. Medical images may have 10 bits (1024 values), 12 bits (4096 values), or even 16 bits (65536 values) worth of intensity information per voxel. Many algorithms that use intensity information group voxel values into a smaller number of partitions, for example 64, 128, or 256 partitions. We refer to the sets of intensity partitions in images *A* and B' as $\{a\}$ and $\{b\}$, respectively, and the number of intensity partitions used as N_a and N_b. Because the range of voxel intensities in an image is dependent on *T* and $\{b\}$ may also be a function of *T*.

2.2.4.1 Similarity Measures

Registration using voxel similarity measure involves calculating the registration transformation *T* by optimizing some measure calculated directly from the voxel

values in the images rather than from geometrical structures such as points or surfaces derived from the images. In the sections below, most common used similarity measures will be described briefly.

2.2.4.1.1 Image Subtraction. If the assumption is made that the images *A* and *B* being registered are identical, except for the misalignment, then an intuitively obvious similarity measure to use is the sum of squares of intensity differences (SSD). In this case, SSD will be zero when the images are correctly aligned and will increase with misregistration. In the slightly more realistic scenario in which the A and B differ only by Gaussian noise, then it can be shown that SSD is the optimum measure (Viola, 1995).

For images A and B with voxels i, correct transformation can be found by minimizing

$$SSD = \frac{1}{N}\sum_i^N |A(i) - B'|^2 \quad \forall i \in A \cap B' \tag{2.18}$$

Certain image registration problems are reasonably close to this ideal case. For example, in serial registration of MR images, it is expected that the images being aligned will be identical except for small changes, which might result from disease progression or response to treatment. Similarly, in functional MR experiments, only a small number of the voxels are expected to change during the study, so all the images that need to be registered to correct for patient motion during the study are very similar to each other. If only a small fraction of the voxels being aligned are likely to have changed between image acquisitions, SSD is likely to work well. This approach has been used by Hajnal et al. (Hajnal et al., 1995) and is used in the SPM software by Friston et al. (Friston et al., 1996). This approach can fail if the data diverges too much from the ideal case. For example, if a small number of voxels change intensity by a large amount, they can have a large effect on the change in squared intensity difference. For this reason, it is sometimes desirable to pre-segment parts of the image prior to registration. This preprocessing is commonly done for the scalp when carrying out serial MR brain registration, where the scalp can deform.

2.2.4.1.2 Correlation Coefficient. If the intensities in images *A* and *B* are linearly related, then the correlation coefficient (CC) can be shown to be the ideal similarity measure (Viola, 1995). Once again, few registration applications will precisely conform to this requirement, but many intramodality applications come sufficiently close for this to be a useful measure.

For images *A* and *B* with voxels *i*, correct transformation can be found by maximizing

$$CC = \frac{\sum_i (A(i)-\bar{A})\left(B'(i)-\overline{B'}\right)}{\left[\sum_i (A(i)-\bar{A})^2 \sum_i \left(B'(i)-\overline{B'}\right)^2\right]^{1/2}} \quad \forall i \in A \cap B' \tag{2.19}$$

where $\bar{A}$ and $\overline{B'}$ are the mean values of voxels in image *A* and the transformed image *B*, respectively.

2.2.4.1.3 Ratio-Image Uniformity. An alternative intramodality registration measure was proposed by Woods, Cherry & Mazziotta (1992). This algorithm was initially devised for registration of multiple PET images of the same subject and has subsequently been widely used for serial MR registration of the brain (Freeborough & Fox, 1997). The name Ratio-Image Uniformity (RIU) is also known as Variation of Intensity Ratios (VIR). The RIU name does, however, explain well what the algorithm does.

For each estimate of the registration transformation, a ratio image *R* is calculated by dividing each voxel in *A* by each voxel in *B'*. The uniformity of *R* is then determined by calculating the normalized standard deviation of *R*. The algorithm iteratively determines the transformation *T* that minimizes the normalized standard deviation, i.e., maximizes uniformity.

In some cases, in order to get good results, it may be necessary to preprocess the images to remove some anatomy. In the original PET-PET application, for example, it was necessary to segment the brain from the images, removing all extradural

material so that both A and B include only brain voxels. This step is not necessary, however, for serial MR registration. Furthermore, an alternative registration will be obtained in almost all cases by reversing the roles of images A and B (i.e., the result will not be the inverse of the forward transformation). Woods proposed that possible bias of the registration solution produced by this asymmetry can be reduced by averaging the forward transformation with the inverse of the reverse transformation (Woods et al., 1992).

2.2.4.1.4 Partitioned Intensity Uniformity. Shortly after proposing the RIU method for intramodality registration, Woods proposed a modified version of the algorithm for MR-PET registration (Woods, Mazziotta & Cherry, 1993). This second algorithm called Partitioned Intensity Uniformity (PIU), makes an idealized assumption that all pixels with a particular MR pixel value represent the same tissue type so that values of corresponding PET pixels should also be similar to each other. The algorithm therefore minimizes the normalized standard deviation of PET voxel values for each MR intensity value (or partitioned intensity value). The PIU algorithm can usefully be thought of in terms of intensity histograms. For each intensity partition a in the MR image, there will be $n_A(a)$ voxels distributed throughout the MR image volume that have an intensity within this partition. For a given estimate T of the registration transformation, there will be $n_A(a)$ PET voxels that are co-located with these MR voxels. The intensities of these PET voxels could be plotted as a histogram. The PIU algorithm aims to find the transformation T that minimizes the spread of that histogram for all partitions. The algorithm can fail unless the histogram for each partition is unimodal.

For MR-PET registration, successful registrations are possible even with a relatively crude segmentation process (West et al., 1997). As with the RIU method, the registration obtained when A and B are interchanged will rarely be consistent with the original registration. The choice of PIU measure will depend on the application.

For images A and B, correct transformation can be found by minimizing

$$PIU = \sum_{a\in\{a\}} \frac{n_A(a)}{N} \frac{\sigma_{B'}(a)}{\mu_{B'}(a)} \tag{2.20}$$

Where B' has been interpolated to have voxels of the same size and shape as *A*, the sum is carried out over the *N* voxels in $A \cap B'$, $n_A(a)$ is the number of voxels in image *A* with intensity *a*, and $\mu_{B'}(a)$and $\sigma_{B'}(a)$ are the mean and standard deviation of the voxels in image B' that co-occur with voxels whose intensities lie in partition *a* in *A* image.

2.2.4.1.5 Joint Histograms and Joint Probability Distributions. Statistical classifiers have been widely used in MR image analysis for segmentation of multi-spectral data for many years. In these approaches, a joint histogram is constructed from two images that are correctly aligned. A joint histogram is *n*-dimensional where *n* is the number of images used to generate it. The axes of the histogram are the intensities (or intensity partitions) in each image, and the value at each point in the histogram is the number of voxels with a particular combination of intensities in the different spectral components.

If the joint histogram is normalized, it becomes an estimate of the joint probability distribution function (PDF) of intensities in the *n* images:

$$PDF\ [j,k] = \frac{\text{HIST}\ [j,k]}{\sum_{j,k} \text{HIST}\ [j,k]} \tag{2.21}$$

For many image modality combinations, the PDF changes with *T*. This observation lead several research groups to investigate similarity measures calculated from the PDF, which could be optimized to register the images. The attraction of this approach is that the observed change in PDF with *T* is qualitatively similar for many modality combinations; the signal in the PDF is "clustered" at registration but diffuses with misregistration.

2.2.4.1.6 Joint Entropy. The Shannon entropy H is widely used as a measure of information in many branches of engineering. It was originally developed as part of information theory in the 1940s and describes the average information supplied by a set of symbols $\{s\}$ whose probabilities are given by { *p(s)* }.

$$H = -\sum_s p(s) \log p(s) \tag{2.22}$$

If all symbols *s* have equal probability, then entropy will be at maximum. If one symbol has a probability of 1 and all others have a probability of zero, then entropy will have a minimum value.

The use of entropy and other information-theoretic measures for image registration came about after inspection of joint histograms and PDFs. When the images are correctly aligned, the joint histograms have tight clusters, surrounded by large dark regions. These clusters disperse as the images become less well registered. The tight clusters in the histograms at registration represent a small number of symbols having high probabilities p(s). As the clusters disperse, the high intensity regions of the joint histogram become less intense, and previously dark regions in the histograms become brighter. Misregistration, therefore, results in an increase in histogram entropy.

As a consequence of this observation, it was proposed that the entropy of the PDF calculated from images A and B' should be iteratively minimized to register these images. For the PDF $[j, k]$ calculated from images *A* and *B*, transformation *T* can be found by minimizing:

$$H = -\sum_{j,k} \text{PDF}\,[j,k] \log \text{PDF}\,[j,k] \tag{2.23}$$

Minimizing joint histogram entropy to register images can be thought of as an extension of PIU minimization, PIU minimization tends to minimize the spread of the histogram of voxels in image B' for each intensity partition in image *A*. Provided this histogram is unimodal, then minimizing the spread will also minimize the

entropy. The joint entropy measure has two advantages over PIU. Firstly, it minimizes the spread of clusters in two dimensions rather than just one. Secondly, minimizing entropy does not require that the histograms are unimodal in the way that minimizing variance does. For this reason, joint entropy would seem more generally applicable to multimodality registration than PIU, and also obviates the need to segment some parts of the image in order to avoid bi-modal histograms

2.2.4.1.7 Mutual Information. Entropy minimization is not a robust voxel similarity measure for all types of image registration. The problem is that the PDF from which the joint entropy is calculated is defined only for the region of overlap between the two images, i.e., within $A \cap B'$. The range and distribution of intensity values in the portion of either image that overlaps with the other is a function of *T*. The change in overlap with *T* can lead to histogram changes that mask the clustering effects described above.

The solution to this difficulty, proposed by researchers is to use the information-theoretic measure mutual information (MI) instead of entropy *H*. MI normalizes the joint entropy with respect to the partial entropies of the contributing signals. In terms of image registration, this measure takes account of the change in the intensity histogram of images *A* and B' with *T*.

$$MI(A, B') = H(A) + H(B') - H(A, B') \tag{2.24}$$

2.2.4.1.8 Normalization of Mutual Information. Mutual information overcomes many of the shortcomings of joint entropy but can still fail for some types of clinical image, particularly those which contain large amounts of air (noise) around the outside of the subject. Improved performance of mutual information can be obtained by various normalization schemes. These algorithms are not taken from the information theory literature but have been arrived at through experiments on medical images. Some studies have shown that it works at least as well as mutual information and in some cases performs better (Studholme, Hill & Hawkes, 1999).

$$NMI(A, B') = \frac{H(A)+H(B')}{H(A,B')} \quad (2.25)$$

2.2.4.2 Optimization and Capture Ranges

All registration algorithms require an iterative approach in which an initial estimate of the transformation is gradually refined by trial and error. In each iteration, the current estimate of the transformation is used to calculate a similarity measure. The optimization algorithm then makes another estimate of the transformation, evaluates the similarity measure again, and continues until the algorithm converges, at which point no transformation can be found that results in a better value of the similarity measure, to within a preset tolerance. A review of optimization algorithms can be found in (Press et al., 1992).

One of the difficulties with optimization algorithms is that they can converge to an incorrect solution called local optimum. It is sometimes useful to consider the parameter space of values of the similarity measure. For rigid registration, there are six degrees of freedom, giving a six-dimesional parameter space, and for an affine transformation with twelve degrees of freedom, the parameter space has twelve dimensions. Each point in the parameter space corresponds to a different estimate of transformation. Nonaffine transformation algorithms have more degrees of freedom, in which case the parameter space has correspondingly more dimensions. The parameter space can be tought of as a high dimensionality image in which the intensity at each location corresponds to the value of the similarity measure for that transformation estimate. If we consider dark intensities as good values of similarity and high intensities as poor ones, an ideal parameter space image would contain a sharp low intensitiy optimum with monotonically increasing intensity with distance away from the optimum position. The job of the optimization would then be to find the optimum location given any possible starting estimate.

Unfortunately, parameter spaces for image registration are frequently not this simple. There are often multiple optima within the parameter space, and registration

can fail if the optimization algorithm converges to the wrong optimum. Some of these optima may be very small caused either by interpolation artifacts, or a local good match between features or intensities. These small optima can often be removed from parameter space by blurring the images prior to registration. In fact, a hierarchical approach to registration is common: the images are first registered at low resolution, then transformation solution obtained at this resolution is used as the starting estimate for registration at a higher resolution, and so on (Studholme, Hill & Hawkes, 1997).

Multiresolution approaches do not entirely solve the problem of multiple optima in the parameter space. It might be thought that the optimization problem involves finding globally optimal solution within the parameter space, and that solution to the problem of multiple optima is to start the optimization algortihm with multiple starting estimates, resulting in multiple solution, and choose the solution which has the best value of the similarity measure. This sort of approach, called multistart optimization, can be effective for surface-based algorithms. For voxel similarity measures, however, the problem is more complicated (Hajnal, Hill & Hawkes, 2001). The desired optimum when registering images using voxel similarity measures is frequently not the global optimum, but is one of the local optima. When registering images using joint entropy, an extremely good value of the similarity measure can be found by transforming the images such that only air in the images overlaps. This will give a few pixels in the joint histogram with very high probabilities, surrounded by pixels with zero probability. This is a very low entropy situation and will tend to have lower entropy than the correct alignment. The global optimum in parameter space will, therefore, tend to correspond to an obvious incorrect transformation. The solution to this problem is to start the algorithm within the capture range of the correct optimum; that is, within the portion of the parameter space in which the algorithm is more likely to converge to the correct optimum than incorrect global one. In practical terms, this requires that the starting estimate of the registration transformation is reasonably close to the correct solution. The size of the capture range depends on the features in the images and can not be known a priori, so it is difficult to know in advance whether the starting estimate is sufficiently good. Visual

inspection of the registered images can easily show convergence outside the capture range. In this case, the solution is clearly and obivious wrong (e.g., relevant features in the image do not overlap at all). If this sort of failure of the algorithm detected, the registration can be restarted with a better starting estimate obtained, for example, by interactively transforming one image until it is approximately aligned with the other (Hajnal et al., 2001).

2.2.4.2.1 Brief Survey of Optimization Techniques. The optimizer applies an efficient and often non-exhaustive strategy to search the transformation parameter space for the best match between the images. In image registration, optimizers can be categorized as gradient-based or gradient-free. Gradient-based algorithms require direct measurements of the gradient of the objective function, but these measurements are generally a gradient estimate because the data involved in the computations are usually noisy. Gradient-free algorithms have been developed in a stochastic setting introducing randomness either in the noisy measurements of the objective function and its gradient, or in the computation of the gradient approximation. These methods can be useful for a broad search in the parameter space being optimized, and can provide initialization for a more powerful local search algorithm. Some gradient-free optimization methods are Simultaneous Perturbation Stochastic Approximation, Simulated Annealing, Hooke-Jeeves and the Nelder-Mead Simplex methods which attempt to minimize a scalar-valued nonlinear function using only function evaluations. Approaches based on the use of gradient estimations tend to be fast, but are sensitive to the presence of local optima. Discussion of these methods can be found in (Spall, 2003).

2.2.4.3 Interpolation

The most widely used image interpolation function is trilinear interpolation, in which a voxel value in the transformed coordinates is estimated by taking weighted average of the nearest eight neighbors in the original dataset. The weightings, which add up to one, are inversely proportional to the distance of each neighbor from the new sample point. For accurate comparison of registered images, for example by

subracting one image from another, the errors introduced by trilinear interpolation become important. It can be shown that trilinear interpolation applies a low-pass filter to the image and introduces aliasing (Parker, Kenyon & Troxel, 1983). For transformations that contain rotations, the amount of low-pass filtering varies with position in the image. If subtracting one image from another to detect small change, for example in serial MR imaging, the low-pass filtering in this process can lead to substantial artifacts. Subtracting a low-pass filtered version of an image from original is a well known edge enhancement method. So even in the case of identical images differing only by a rigid-body transformation, using linear interpolation followed by subtraction does not result in the expected null result, but instead results in an edge-enhanced versions of the original.

Hajnal et al. (2001, 1995) proposed that the solution to this issue, is to interpolate using a sinc function truncated with a suitable window function such as a Hamming window. Later on, some modifications to sinc interpolation have been proposed with the use of sinc functions with a various radii truncated with various window functions (Lehmann, Gonner & Spitzer, 1999); approximations to windowed sinc function such as cubic or B-spline interpolants (Unser, 1999); and the shear transform, which involves transforming the image using a combination of shears (Cox & Jesmanowicz, 1999; Eddy, Fitzgerald & Noll, 1996).

Many registration algorithms involve iteratively transforming image *B* with respect to image *A* while optimizing a similarity measure calculated from the voxel values. Interpolation errors can introduce modulations in the similarity measure with *T*. This is most obvious for transformations involving pure translations of datasets with equal sample spacing, where the period of the modulation is the same as the sample spacing (Pluim, Maintz & Viergever, 2000). This periodic modulation of the similarity measure introduces local optima that can lead to the incorrect registration solution being determined.

The computational cost of correct interpolation is far too high for his approach to be used in each iteration, so lower cost interpolation techniques must be used. There

are several possible approaches. The first is to use low-cost interpolation, such as trilinear, bilinear or nearest neighbor, until the transormation is close to the desired solution, then carry out the final few iterations using more expensive interpolations. An alternative strategy is to take advantage of the spatial-frequency dependence of interpolation errors. Trilinear interpolation low-pass filters the data, and therefore, if the images are blurred prior to registration (high spatial frequency components are removed), the interpolation errors are smaller, so errors in the registration are less. Although the loss of resolution that results from blurring is a disadvantage, registration errors caused by interpolation errors can be greater than the loss of precision resulting from blurring.

2.2.4.4 Applications of Intensity-Based Methods

2.2.4.4.1 Types of Transformation. Many of similarity measures have been applied to a wide variety of registration problems, including:

- registration of 2-D and 3-D images of the same subject taken at different times using the same modality
- registration of 2-D and 3-D images of the same subject taken at different times using different modalities
- registration of 2-D projection images and 3-D images of the same subject
- inter-subject registration of the 3-D images of the head

Mutual information and normalized mutual information have been demonstrated to provide good quality image registration for all these types of applications. The validation of the registration quality is a particularly difficult problem when calculating nonrigid transformations and is a rather poorly defined problem when aligning images from different subjects. In the sections below, will be described the application of voxel-similarity measures to a selection of medical image registration tasks.

2.2.4.4.2 Serial MR Registration. It is increasingly common for a patient to have multiple magnetic resonance images acquired at different times, in order to monitor disease progression or response to treatment. Traditional radiological viewing of

these images involves printing the images onto film and viewing them side-by-side on light boxes. This approach can make it difficult to identify small changes in the images between scans. A more sensitive technique is to subtract the images, generating difference images, and to view the difference images to identify regions of the images that have changed. The generation of difference images requires that the images acquired at the different times have identical contrast characteristics and that any drift in the scanner scaling parameters is corrected using suitable quality-assurance procedures, such as phantom scanning (Hill et al., 1998; Lemieux & Barker, 1998). Serial MR registration also requires accurate alignment, because registration errors of less than a voxel can lead to artifacts that can mask the changes of interest in the patient. The required registration accuracy will be a function of the spatial frequency content of the images and of the contrast-to-noise ratio. The required accuracy is often stated as of the order of tens or hundreds of microns.

Serial MR registration is almost invariably applied to the brain and makes use of a rigid transformation. Indeed, for visual inspection of difference images, it is undesirable to use nonrigid registration transformations because they may produce a transformation that removes the changes of interest.

With serial MR registration the images are likely to be very similar, and so the measures SSD and CC are widely used. RIU, MI, and NMI are also frequently used. The advantage of the information-theoretic measures in this case is that they are less sensitive to small numbers of voxels that change by large amounts, than are the SSD and CC measures. When using SSD or CC, data preparation by segmentation of extradural tissue, or small lesions that have changed substantially in intensity, is often essential. The information-theoretic approaches should be less sensitive to these small numbers of outlying voxels, but where the changes are large, data pre-processing may still be required. For many applications, the choice of an optimization algorithm that will iteratively find the optimum transformation to subvoxel accuracy is as important as the choice of similarity measure.

In serial MR registration, the choice of interpolation algorithm is also important. The use of trilinear interpolation is equivalent to convolving the data with a triangular kernel of variable width, resulting in a spatially varying blurring of the data. When subtracting image A and B' edge artifacts result. Higher order interpolation methods, such as those that incorporate truncated sinc kernels (Hajnal et al, 1995) or skew transformations (Eddy, Fitzgerald & Noll, 1996), are more effective.

2.2.4.4.3 MR and CT Registration. Registration of MR and CT images is most often applied to images of the head, and is used for surgery and radiotherapy planning. The registration transformation determined is usually a rigid one. On some occasions, however, it may be beneficial to search for an affine transformation because it can generate more accurate registration in cases when there are scaling or skew errors in the data. When the images being registered are subsequently going to be used for guiding treatment (image-guided surgery or radiation therapy), then it is not advisable to use an affine transformation unless one modality is known to have satisfactory geometrical properties and is treated as the reference image. Otherwise, the images may be well registered but measurements made in the registered images may be inaccurate. It is preferable in these cases to calibrate the imaging devices as part of a routine quality assurance procedure and to rectify the images prior to registration to remove any residual geometric distortion. The most problematic aspects of imaging geometry are geometric distortion in MR and errors in CT bed speed and gantry tilt angle. For 3-D gradient echo MR images, object dependent distortion is predominantly in the readout direction and can be minimized by selecting a high bandwidth-per-pixel in that direction (at a slight cost in signal to noise ratio). The residual scaling errors in MR due to gradient strength errors, and the CT errors, can be corrected using phantom experiments.

Since MR and CT images have very different intensity distributions, image subtraction is never used for viewing the aligned images. As a result, the very high accuracy requirements of serial MR registration do not apply here. The accuracy requirement for MR-CT registration is usually governed by the accuracy of the

treatment system that will subsequently be used. Since neurosurgery and radiosurgery systems are accurate to a small number of millimeters, registration accuracy of about 1 mm is normally sufficient (Fitzpatrick et al., 2000).

A study by West et al. (1997) established that, if the images are free from motion artifacts and have been rectified to remove scaling errors and distortion, registration accuracy of better than 1 mm is possible when information-theoretic measures are used. It is important to note that the images used in that validation study had similar fields of view in both modalities, covering most of the head, and had very similar slice orientations. It is possible to register MR and CT images that have very different slice orientations and fields of views. For example, it is possible to register coronal CT images with axial MR images or to register CT images with a field of view of about 20 mm around the temporal bone to MR images of the entire head. In both these cases, the rigid transformation solution may include translations of several tens of millimeters and rotations of many tens of degrees. In such cases, it is advisable to provide the algorithm with an initial estimate of T to increase the likelihood that it is starting within its capture range. Also, since the overlap between images at the correct registration transformation includes only a small portion of one or both modalities, the registration accuracy in these cases may be lower than in the validation study by West et al. (West et al., 1997).

2.2.4.4.4 MR or CT and PET Registration. Registration of MR or CT images with PET images of the head was one of the first applications of medical image registration. Because of the relatively low resolution of PET images, it is frequently desirable to make use of anatomical detail from MR or CT images to assist in their interpretation. The resolution of PET images is of the order of 6–10 mm, and a registration accuracy of 2–4 mm is considered desirable (Fitzpatrick et al, 2000). Although the low resolution of PET images reduces the required registration accuracy compared to MR-CT registration, there are several difficulties in this registration application. Firstly, some PET tracers are very specific to certain parts of the brain. Although several algorithms, including Partitioned Intensity Uniformity and mutual information, have been shown to work well when registering MR images

to PET fluoro-deoxyglucose images (West et al., 1997), the accuracy of registration of the more specific tracers will be tracer dependent and may be much lower. Secondly, for patients with large lesions, the PET images can have very large intensity abnormalities. A lesion that makes a small difference to image intensities in an MR scan may make a very large difference to intensities in the PET image. Many of the images used in the validation study by West et al. (1997) had large lesions of this sort, and the median registration accuracy of both PIU and MI was better than 4 mm, suggesting that these measures are reasonably robust to these intensity abnormalities.

2.2.4.4.5 Nonrigid 3D Registration. It is possible to apply voxel-similarity measures to nonrigid registration by placing a regular grid of *N* nodes across images *A* and *B* being registered. The position of these nodes in B is then iteratively altered while optimizing a voxel similarity measure. One approach is to calculate a rigid or affine transformation for a sub-image around each node, and then to smooth the resulting set of discrete transformations (Collins, Holmes, Peters & Evans, 1995). Alternatively, each node can be treated as a spline knot, and the nodes iteratively adjusted, with the entire image being transformed using the spline transformation at each iteration. The approximating B-spline, has the advantage that change in position at a node has only a local effect on the deformation (extending about two node separations), and this approach is used in an implementation by Rueckert et al. (Rueckert et al., 1998).

The first approach involves multiple registrations of sub-images. The accuracy of registration at each sub-image will be dependent on the number of voxels in the sub-images and the intensity information contained. Each sub-image registration involves 6 or 12 degrees of freedom (for rigid and affine transformations respectively), and this is carried out *N* times. Because each sub-image contains only a small number of voxels, the information-theoretic approaches may suffer from a poor estimate of the PDF, and hence noisy values of the voxel-similarity measures. The SSD and CC measures will be less sensitive to this problem. The second approach involves three translational degrees of freedom per node, producing *3N* degrees of freedom in total,

with the similarity measure being calculated for the entire images, rather than just for sub-images. This approach has the advantage that large numbers of voxels are used in each iteration, making the information-theoretic measures practicable.

These algorithms have been applied to both intrapatient registration in which there is tissue deformation and to interpatient registration.

2.2.4.4.6 2-D to 3-D Registration. Many medical imaging modalities generate images by perspective projection. The most common of these is conventional radiography or fluoroscopy. However, with video endoscopy and microscopy becoming more widely used, video images are increasingly common. Aligning projection images with 3-D tomographic modalities is a type of 2-D-to-3-D registration and requires an estimate of both the 3-D orientation of the objects being registered and the projection transformation. For rigid structures, the 3-D orientation information is a rigid transformation, so T is the composition of a rigid transformation and the projection transformation.

In many medical imaging applications, the projection transformation can be measured as part of a calibration procedure, so the registration problem has just six degrees of freedom.

The approach used for registration of either video images or radiographs to tomographic images is to optimize a similarity measure while iteratively transforming the tomographic image. For each iteration, a simulated projection image is generated from the tomographic image, and the similarity measure is calculated from these. In this case, image A is the projection image, and B the tomographic image. For each estimation of the transformation T, image $B'$is a simulated projection generated from the tomographic image B. For example, when registering a radiograph A to a CT scan B, a digitally reconstructed radiograph B' is produced using T and a model of the x-ray image formation process to simulate an x-ray projection by integrating the CT voxel intensities. This approach has been used by several authors (Lemieux, Kitchen, Hughes & Thomas, 1994; Penny et al, 1998). When registering a video image A to a tomographic image B, B' can be a perspective

surface rendering of the image structures in *B* that are visible in *A* (Viola & Wells, 1995).

CHAPTER THREE
INFORMATION THEORY AND MUTUAL INFORMATION

Image registration can be described as trying to maximize the amount of shared information in two images or volumes. In a very qualitative sense might be said that if two images are correctly aligned, then corresponding structures will overlap. When the images are out of alignment, however, combined image will have duplicate versions of these structures from images.

Using this concept, registration can be thought of as reducing the amount of information in the combined image, which suggests the use of measure of information as a registration metric. Because of its robustness and accuracy in medical image registration, mutual information has become a common metric in many clinical applications. There is a wealth of papers mentioning the use of the method as a step in a larger method or in an application.

This chapter starts with the definition of entropy and its interpretation. Then presents mutual information history in image registration, its multiple forms of definition and its properties. It continues with general introduction to normalized mutual information and previous works of medical image registration, including references to mutual information based methods.

Classification has two main categories: methodological aspects and matters of application. The aspects of the method are subdivided into preprocessing, measure, transformation and implementation, most of which have a further subclassification. Aspects of the application entail the image modalities, the subject of registration (a single person or different persons) and the object of registration (the imaged anatomy). It is also classified according to the image dimensionality and the number of images involved in registration.

3.1 Entropy

The desire for a measure of information (commonly termed entropy) of a message stems from communication theory. This field concerns the broadcast of a message from a sender to a receiver. For the first attempt Hartley (1928) defined a measure of information of a message that forms the basis of many present day measures. He considered a message of a string of symbols, with s diferent possibilities for each symbol. If the message consists of n symbols, there are s^n diferent messages possible. The measure s^n complies, but the amount of information would increase exponentially with the length of the message and that is not realistic. Hartley wanted a measure H that increases linearly with n, i.e. $H = Kn$, where K is a constant depending on the number of symbols s. He further assumed that, given messages of length n_1 and n_2 from s_1 and s_2 i.e. the numbers of symbols, respectively, if ${s_1}^{n_1} = {s_2}^{n_2}$, i.e. number of possible messages is equal, then the amount of information per message is also equal. These two restrictions led him to define the following measure of information

$$H=n \log s = \log s^n \tag{3.1}$$

Hartley's information measure depends on the number of possible outcomes: the larger the number of possible messages, the larger the amount of information you get from a certain message.

A drawback of Hartley's measure is that it assumes all symbols (and hence all messages of a given length) are equally likely to occur. Clearly, this will often not be the case. Shannon introduced an adapted measure in 1948, which weights the information per outcome by the probability of that outcome occurring. Given events $e_1, \ldots, e_m$ occurring with probabilities $p_1, \ldots, p_m$, the Shannon entropy is defined as

$$H= \sum_i p_i \log \frac{1}{p_i} = - \sum_i p_i \log p_i \tag{3.2}$$

If we apply to Shannon's entropy the assumption that all outcomes are equally likely to occur, we get

$$H=-\sum\frac{1}{s_n}log\frac{1}{s_n}=\sum s_n\,log\,s^n=log\,s^n \qquad (3.3)$$

which is exactly Hartley's entropy.

Although the second definition of the Shannon entropy in equation (3.2) is more commonly used, the first one more clearly explains the meaning. The term $\log\frac{1}{p_i}$ signifies that the amount of information gained from an event with probability p_i is inversely related to the probability that the event takes place. The more rare an event, the more meaning is assigned to occurrence of the event. The information per event is weighted by the probability of occurrence. The resulting entropy term is the average amount of information to be gained from a certain set of events.

The Shannon entropy can also be computed for an image, in which case we do not focus on the probabilities of letters or words occurring, but on the distribution of the grey values of the image. A probability distribution of grey values can be estimated by counting the number of times each grey value occurs in the image and dividing those numbers by the total number of occurrences. An image consisting of almost a single intensity will have a low entropy value; it contains very little information. A high entropy value will be yielded by an image with more or less equal quantities of many different intensities, which is an image containing a lot of information.

In this manner, the Shannon entropy is also a measure of dispersion of a probability distribution. A distribution with a single sharp peak corresponds to a low entropy value, whereas a dispersed distribution yields a high entropy value.

3.2 Joint Entropy

In image registration we have two images, *A* and *B*, to align. We therefore have two symbols at each voxel location for any estimate of transformation *T*. Joint entropy measures the amount of information we have in combined images (Shannon, 1948). If *A* and *B* are totally unrelated, then the joint entropy will be the sum of the entropies of individual images. The more similar (i.e., less independent) the images are, the lower the joint entropy compared to the sum of the individual entropies.

$$H(A,B) \leq H(A)+H(B) \tag{3.4}$$

The concept of joint entropy can be visualized using a joint histogram calculated from images *A* and *B*, examples of which are shown in Figure 3.1. The joint histogram can be normilized by dividing it by the total number of voxels, and regarded as a joint probability distribution function (PDF) p_{AB} of images *A* and *B*. Due to quantization of image intensity values, the PDF is discrete, and the values in each element represent the probability of pairs of image values occuring together.

The joint entropy *H(A,B)* is therefore given by:

$$H(A,B) = -\sum_a \sum_b p_{AB} log p_{AB}(a,b) \tag{3.5}$$

The number of elements in the PDF can either be determined by the range of intensity values in two images or from a reduced number of intensity bins. For example, MR and CT images could have up to 4096 (12bits) intensity values, leading to very sparse PDF with 4096 by 4096 elements. The use of 32 to 256 bins are common. In the above equation, *a* and *b* represent either the original image intensities or the selected intensity bins.

As seen in Figure 3.1, as misregistration increases the brightest regions of the histogram increasingly disperse, and size of surrounding dark regions reduce. If we interpret the joint histogram as a joint probability distribution, then misregistration results in reductions of the highest values in the PDF and reducing the size of zero

regions in the PDF; this will obviously increase the entropy. Conversely, when registering images we desire to find a transformation that will produce a small number of PDF elements with very high probabilities and give us as many zero probability elements in the PDF as possible, which will minimize the joint entropy.

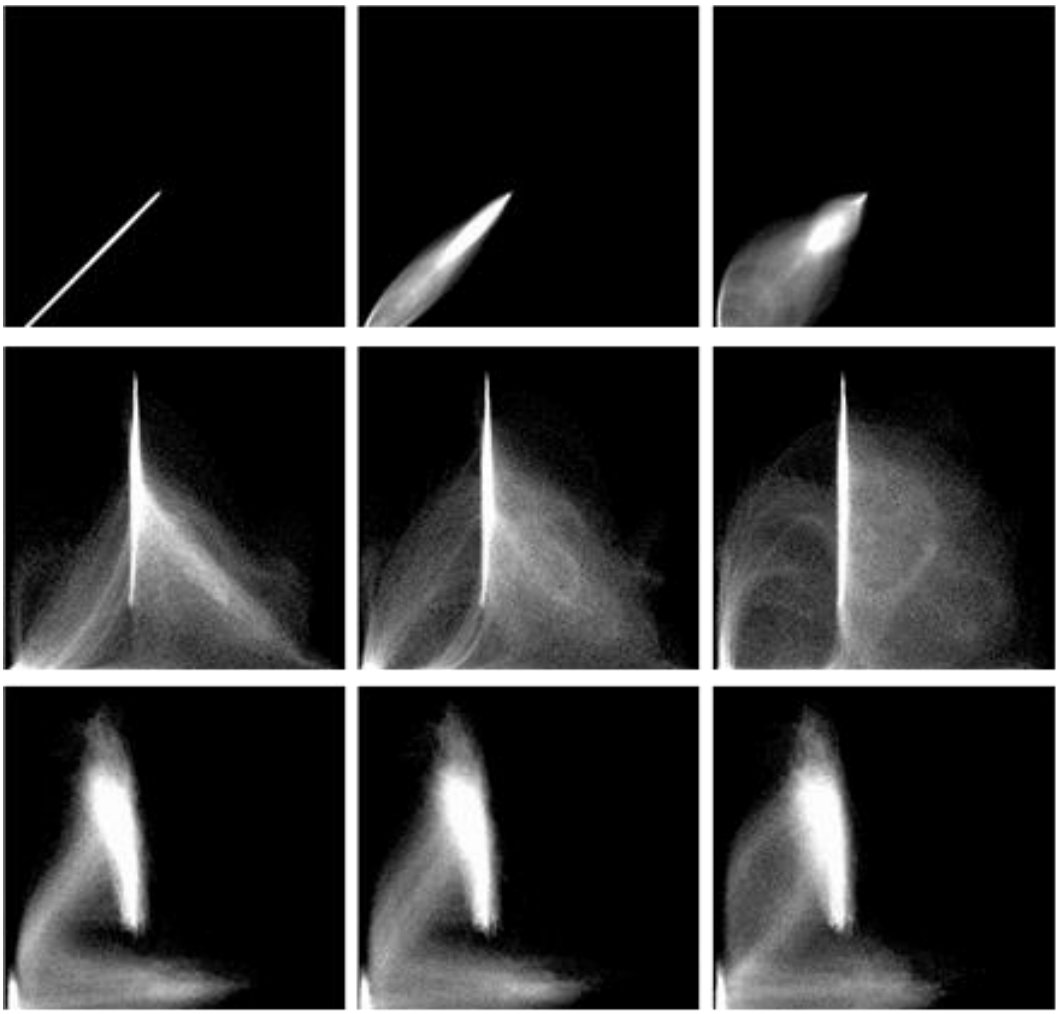

Figure 3.1 Example 2D joint histograms (Hill et al., 1994). First row identical MR images of the head, second row is a MR and CT images of the head and third row MR and PET images of the head. For all modality combinations, the left panel is generated from the images when aligned, the middle panel when translated by 2mm, and the right panel when translated by 5mm.

The simple form of the equation for joint entropy (Eq. 3.5) can hide an important limitation of this measure: joint entropy is dependent on T. In particular, p_{AB} is very dependent on overlapping regions of the images which is undesirable, and also on the interpolation algorithm used to transform the image B at each iteration. A registration

algorithm that seeks to minimize joint entropy will tend, therefore, to maximize the amount of air in overlapping regions, whic may result in an incorrect solution. The interpolation dependence of p_{AB} is clear if we remember that interpolation algorithms will tend to blur images, which sharpens the corresponding image histogram, changing the joint histogram and consequently joint probability distribution p_{AB}.

3.3 Mutual Information

3.3.1 History

The research that eventually led to the introduction of mutual information as a registration measure dates back to the early 1990's. Woods et al. (1992) first introduced a registration measure for multimodality images based on the assumption that regions of similar tissue (and hence similar grey values) in one image would correspond to regions in the other image that also consist of similar grey values (though probably different values to those of the first image). Ideally, the ratio of the grey values for all corresponding points in a certain region in either image varies little. Consequently, the average variance of this ratio for all regions is minimized to achieve registration.

Hill et al. (1993) proposed an adaption of Woods' measure. They constructed a feature space, which is a two dimensional plot showing the combinations of grey values in each of the two images for all corresponding points. The difference with Woods' method is that instead of defining regions of similar tissue in the images, regions are defined in the feature space. These regions are based on the clustering one finds in the feature space for registered images.

The feature space (or joint histogram) changes as the alignment of the images changes. When the images are correctly registered, corresponding anatomical structures overlap and the joint histogram will show certain clusters for the grey values of those structures. As the images become misaligned, structures will also start overlapping structures that are not their anatomical counterparts in the other

image. Consequently, the intensity of the clusters for corresponding anatomical structures will decrease and new combinations of grey values emerge, such as skull and brain or skin and background. This will manifest itself in the joint histogram by a dispersion of the clustering. Clearly, the histogram shows increasing dispersion as the misregistration increases.

Using this characteristic of the joint histogram of two images, measures of dispersion emerged, were used for image registration. Hill, Studholme & Hawkes (1994) proposed the third order moment of the joint histogram, which measures the skewness of a distribution. Both Collignon, Vandermeulen, Suetens & Marchal (1995) and Studholme, Hill & Hawkes (1995) suggested to use entropy as a measure of registration. As explained before, entropy measures the dispersion of a probability distribution. It is low when a distribution has a few sharply defined, dominant peaks and it is maximal when all outcomes have an equal chance of occurring. A joint histogram of two images can be used to estimate a joint probability distribution of their grey values by dividing each entry in the histogram by the total number of entries. By finding the transformation that minimizes their joint entropy, images should be registered.

Once entropy, a measure from information theory, had been introduced for the registration of multimodality medical images, another such measure quickly appeared: mutual information. It was pioneered by Collignon et al. (1995), Collignon (1998), Viola (1995), Viola & Wells (1995) and Wells, Viola & Kikinis (1995). Applied to rigid registration of multimodality images, mutual information showed great promise and within a few years it became the most investigated measure for medical image registration.

3.3.2 Definition

Most books on information theory (Aczél & Daroczy, 1975; Cover & Thomas, 1991; Vajda, 1989) discuss the notion of mutual information. The definition of the term, however, can be presented in various ways. We will next treat three frequently

used forms of the definition, because more than one is used in the literature. All three forms are identical; each can be rewritten into the other two. Each form of definition, however, explains the relation to registration in a different way. We will describe mutual information for two images, as used in image registration, and not in a general sense.

The first form of definition we discuss is the one that best explains the term mutual information. For two images *A* and *B*, mutual information *I* can be defined as

$$I(A,B)=H(B)-H(B|A) \tag{3.6}$$

where $H(B)$ is the Shannon entropy of image *B*, computed on the probability distribution of the grey values. $H(B|A)$ denotes the conditional entropy, which is based on the conditional probabilities $p(b|a)$, the chance of grey value b in image *B* given that the corresponding voxel in *A* has grey value a. When interpreting entropy as a measure of uncertainty, equation (3.5) translates to the amount of uncertainty about image *B* minus the uncertainty about *B* when *A* is known. In other words, mutual information is the amount by which the uncertainty about *B* decreases when *A* is given: the amount of information *A* contains about *B*. Because *A* and *B* can be interchanged, $I(A,B)$ is also the amount of information *B* contains about *A*. Hence, it is mutual information. Registration is assumed to correspond to maximizing mutual information: the images have to be aligned in such a manner that the amount of information they contain about each other is maximal.

The second form of definition is most closely related to joint entropy. It is

$$I(A,B)=H(A)+H(B)-H(A,B) \tag{3.7}$$

This form contains the term $-H(A,B)$, which means that maximizing mutual information is related to minimizing joint entropy. We have described above how the joint histogram of two images' grey values disperses with misregistration and that joint entropy is a measure of dispersion. The advantage of mutual information over

joint entropy, is that it includes the entropies of the separate images. Mutual information and joint entropy are computed for the overlapping parts of the images and the measures are therefore sensitive to the size and the contents of overlap. A problem that can occur when using joint entropy on its own, is that low values (normally associated with a high degree of alignment) can be found for complete misregistrations. For example, when transforming one image to such an extent that only an area of background overlaps for the two images, the joint histogram will be very sharp. There is only one peak, that of background. Mutual information is better equipped to avoid such problems, because it includes the marginal entropies $H(A)$ and $H(B)$. These will have low values when the overlapping part of the images contains only background and high values when it contains anatomical structure. The marginal entropies will thus balance the measure somewhat by penalizing for transformations that decrease the amount of information in the separate images. Consequently, mutual information is less sensitive to overlap than joint entropy, although not completely immune.

The final form of definition we discuss is related to the Kullback-Leibler distance, which is defined as $\sum_i p(i) \log \frac{p(i)}{q(i)}$ for two distributions p and q. It is a measure of the distance between two distributions. Analogous to the Kullback-Leibler measure, the mutual information of images A and B is defined as

$$I(A,B) = \sum_{a,b} p(a,b) \log \frac{p(a,b)}{p(a)p(b)} \qquad (3.8)$$

The interpretation of this form is that it measures the distance between the joint distribution of the images' grey values $p(a,b)$ and the joint distribution in case of independence of the images, $p(a)p(b)$. It is a measure of dependence between two images. The assumption is that there is maximal dependence between the grey values of the images when they are correctly aligned. Misregistration will result in a decrease in the measure.

3.3.3 Properties

Mutual information has the following properties (Aczél et al., 1975):

i. $I(A, B) = I(B, A)$

It is symmetric; otherwise it would not be *mutual* information. However, although it is a logical property in theory, mutual information is not symmetric in practice. Implementational aspects of a registration method, such as interpolation and number of samples, can result in diferences in outcome when registering A to B or B to A.

ii. $I(A, A) = H(A)$

The information image A contains about itself is equal to the information (entropy) of image A.

iii. $I(A, B) \leq H(A),\ I(A, B) \leq H(B)$

The information the images contain about each other can never be greater than the information in the images themselves.

iv. $I(A, B) \geq 0$

The uncertainty about A cannot be increased by learning about B.

v. $I(A, B) = 0$ if and only if A and B are independent.

When A and B are not in any way related, no knowledge is gained about one image when the other is given.

3.4 Normalized Mutual Information

It is suggested by the information theoretic approach that, at alignment, the value of a voxel in A is a good predictor of the value at the corresponding location in B. As misregistration increases, one image becomes a worse predictor of the second. In practice, two structures which have the same intensity in image A may have very

different intensities in image *B*. For example, in an MR image, cortical bone and air will both have very low intensities, whereas in CT, air will have a very low intensity but cortical bone a high intensity. If we have a low-intensity voxel in an MR image, then at correct alignment we know that this should either be air or bone in CT image. The histogram of CT intensity values corresponding to low intensities in MR will, therefore, have sharp peaks at both low-intensity values and high-intensity values, resulting in low entropy (Hajnal et al., 2001). MI can not fully solve this overlap problem. Changes in overlap of very low-intensity regions of the image (especially air around the patient) can excessively contribute to the mutual information. To alleviate this dependence on volume of overlap MI should be normalized to the combined information in the overlapping volume. Studholme et al. (1999) have proposed a normalization devised to overcome the sensitivity of mutual information to change in image overlap. Their normalized measure is given by:

$$NMI(A,B) = \frac{H(A)+H(B)}{H(A,B)} \qquad (3.9)$$

This formula of NMI has been shown to be evidently more robust than standard MI for intermodality registration in which the overlap volume changes significantly (Studholme et al., 1999).

3.5 Previous Works On Mutual Information Based Registration

In this section, a classification scheme which builds on earlier schemes for medical image registration has been defined in general which is presented in Table 3.1 (Maintz et al., 1998; Van den Elsen et al., 1993; Pluim et al., 2003).

The main subdivision of the classification is in aspects concerning the method and those concerning the application. In addition, the classes image dimensionality and number of images are defined. The elements of these two subclasses can concern purely the application, but they can also necessitate an adaptation of the method. They are therefore treated separately. The class method can be further subdivided into preprocessing, measure, transformation and implementation. Preprocessing

entails any image processing to prepare or improve the images for registration. Typical examples are filtering to remove noise, extraction of regions of interest and isotropic resampling. The measure will obviously be based on mutual information, but differences are possible, characterized by the choice of entropy, by normalization and by adaptations to incorporate spatial information. The transformation can be classified as either rigid (rotations and translations only), affine (rigid plus scaling and shearing), perspective (affine without preservation of parallelism of lines) or curved. Implementation is an important category, because the choice of method for interpolation, probability distribution function (pdf) estimation, optimization and acceleration can have a substantial influence on the registration results.

The first aspect of the class application is the type of modalities it concerns. The images can be of the same kind (monomodality), acquired by different techniques (multimodality), a type of model can be involved (a phantom or atlas) or images are registered to physical space. By the latter we mean registration of previously acquired images to a person, as is used for image guided surgery or radiotherapy treatment. Subject denotes whether images of a single person are involved, which is called intrasubject registration, or images of different persons, inter subject registration, or whether images of a person are matched to a model. Lastly, the anatomy that the registration focuses on is what we term object.

Of image dimensionality there are instances of 2D/2D, 3D/3D and 2D/3D registration in the literature. The number of images involved in the registration is usually 2, but registration of more than 2 images has been described in a number of publications. The latter aspect can be further divided into registration problems where the transformations between several images are known and only a single transformation is to be found or problems that require several transformations.

Most categories of the scheme will be discussed in the following. Some aspects are so dominant in the literature (for example, 3D/3D registration), that will only review the exceptions. After the classification scheme, will be considered a number of comparison studies.

Table 3. 1 Classification scheme for mutual information based registration methods.

Mutual Information Based Registration Methods			
		Preprocessing	
	Method	Measure	- Entropy - Normalization - Spatial Information
		Transformation	- Rigid - Affine - Perspective - Curved
		Implementation	- Interpolation - PDF Estimation - Optimization - Accelaration
	Application	Modalities	- Monomodality - Multimodality - Modality to Modal - Modality to Physical Space
		Subject	- Intrasubject - Intersubject - Modal
		Object	- Brain - Thorax - Lungs - Liver - Heart - Breast

3.5.1 Preprocessing

Several techniques of processing images before registration have been described. The most common preprocessing step is defining a region (Fei, Wheton, Lee, Duerk, & Wilson, 2002; Slomka, Mandel, Downey, & Fenster, 2001) or structures (Erdi et al., 2002; Kagadis et al., 2002; Shekhar & Zagrodsky, 2002) of interest in the images to exclude structures that may negatively influence the registration results. Other

processing techniques reported include low-pass filtering to remove speckle in ultrasound images (Shekhar et al., 2002; Zagrodsky, Shekhar, & Cornhill, 2000) and thresholding or filtering to remove noise (Ding, Goshtasby, & Satter, 2001). Blurring is also applied to correct for differences in the intrinsic resolution of the images (Studholme, 1997; Studholme et al., 1999). Intensity inhomogeneities in images are corrected in several methods, both for MR images (Hayton, Brady, Smith, & Moore, 1999) and for portal images (Plattard et al., 2000). Some methods resample the images isotropically, to achieve similar voxel sizes in all image dimensions (Fei et al., 2002; Slomka et al., 2001Studholme, 1997; Studholme et al., 1999; Studholme et al., 1996).

3.5.2 Measure

Obviously, in a literature survey on mutual information based image registration, the measure in question will always be mutual information. However, when using a definition of mutual information based on entropy, different definitions of entropy can be chosen. Furthermore, several adaptations of mutual information have been proposed: normalization with respect to the overlapping part of the images and inclusion of spatial information.

3.5.2.1 Entropy

By far the most common measure of entropy in the papers is the Shannon entropy. Rodriguez & Loew (1998) use the Jumarie entropy (Jumarie, 1997). The Jumarie entropy is defined for one-dimensional signals and resembles a normalized version of Shannon entropy, applied not to a probability distribution, but to function value differences of neighbouring samples. In Rodriguez et al. (1998), two-dimensional images are registered. The authors define the Jumarie entropy of a 2D image on the gradient magnitude of pixels. The joint Jumarie entropy is defined on the grey value difference of corresponding pixels, which presumably makes the measure less suitable to registration of multimodality images. Ioannides et al. (2000) use the Rényi entropy of order 4, although not for image registration, but for comparison of brain

activity during different tasks. The Rényi entropy of order 2 is employed by Pompe, Blidh, Hoyer, & Eiselt (1998) to measure the strength of dependence between 1D respiratory and cardiac signals.

3.5.2.2 Normalization

The size of the overlapping part of the images influences the mutual information measure in two ways. First of all, a decrease in overlap decreases the number of samples, which reduces the statistical power of the probability distribution estimation. Secondly, Studholme et al. (1999) have shown that with increasing misregistration (which usually coincides with decreasing overlap) the mutual information measure may actually *increase*. This can occur when the relative areas of object and background even out and the sum of the marginal entropies increases, faster than the joint entropy. Studholme et al. (1999), found a distinct improvement in the behaviour of the normalized measure for rigid registration of MR-CT and MR-PET images.

Collignon (1998) and Maes et al. (1997) have suggested the use of the *Entropy Correlation Coefficient* (*ECC*), another form of normalized mutual information. *NMI* and *ECC* are related in the following manner:

$$ECC = 2 - 2/NMI \tag{3.10}$$

Normalized mutual information was used in a large number of studies (Denton et al., 1999; Estevez, Tesmer, Perez, & Zurada, 2009; Hill et al., 1998; Holden, Schnabel, & Hill, 2001; Knops, Maintz, Viergevera, & Pluim, 2006;Lynch et al., 1999; Pluim et al., 2000; Pluim, Maintz, & Viergever, 2001; Radau et al., 2001; Rueckert et al., 2001; Studholme et al., 2001; Studholme, Constable, & Duncan, 2000; Veninga, Huisman, van der Maazen, & Huizenga, 2004;Yokoi, Soma, Shinohara, & Matsuda, 2004). An upper bound of mutual information was derived by Skouson, Guo, & Liang (2001).

3.5.2.3 Spatial Information

A drawback of mutual information as it is commonly used, i.e. based on the Shannon entropy, is that the dependence of the grey values of neighbouring voxels is ignored. The original Shannon entropy does include a dependence of preceding signals, but the definition used in practice is the one for independent successive signals. The assumption of independence does not generally hold for medical images. Incorporating the dependence of the grey values of neighbouring voxels, what we term the spatial information of the images, could improve registration.

As mentioned, Rodriguez & Loew (1998) employ the Jumarie entropy, which considers the grey value differences of neighbouring voxels in an image. Studholme, Hill, & Hawkes (1996) compute the mutual information of two images together with a labelling of one of the images. Voxels with identical grey values can then be differentiated when they belong to different regions. The use of a cooccurrence matrix has been put forth by Rueckert, Clarkson, Hill, & Hawkes (2000). The cooccurrence matrix of distance *d* of an image is a 2D histogram giving the frequencies of two grey values in the image being distance *d* apart. Rueckert et al. show the effect the method has on curved registration of MR images. Another method of incorporating spatial information is to combine mutual information with a measure based on the gradients at corresponding points. The measure seeks to align gradient vectors of large magnitude as well as of similar orientation (Pluim et al., 2000; Pluim, 2001).

3.5.3 Transformation

The transformation applied to register the images can be categorized according to the degrees of freedom. We define a rigid transformation as one that includes only translations and rotations. Although in the literature, rigid transformations are sometimes allowed to include scaling, we classify such transformations as affine. An affine transformation can furthermore include shearing. This type of transformation maps straight lines to straight lines and preserves the parallelism between lines. The perspective transformation differs from the affine transformation in the sense that the

parallelism of lines need not be preserved. It is usually applied in 2D/3D registration. No instances of 'true' perspective transformation were encountered. All methods using a perspective transformation limited the optimization to the rigid-body or affine parameters; the projective parameters were kept fixed. The final class consists of curved transformations, which allow the mapping of straight lines to curves.

3.5.3.1 Rigid

Translations and rotations suffice to register images of rigid objects. Examples include registration of bone or of the brain when neither skull nor dura has been opened. Rigid registration of images based on mutual information has been applied in a large number of papers (Carrillo et al., 2000; Castro-Pareja, Jagadeesh, & Shekhar R. 2004; Collignon, 1998; Ding et al., 2001; Fei et al., 2002; Hill et al., 1998; Holden et al., 2000; Kim et al., 2001; Kim, Li et al., 2007; Maes et al., 1997; Nikou, Heitz, & Armspach, 1999; Penny et al., 1998; Plattard et al., 2000; Rohlfing et al., 2000; Slomka et al., 2001; Studholme, 1997; Studholme et al., 1996; Studholme et al., 1997; Studholme et al., 1999; Thévenaz & Unser, 2000; Viola & Wells, 1997; Wu, Kim, Peters, Chung, & Samant 2009). Rigid registration is also used to approximately align images that show small changes in object shape (Denton et al., 2000; Holden et al., 1999) or small changes in object intensity, as in functional MR time series images (Freire & Mangin, 2001).

3.5.3.2 Affine

The affine transformation preserves the parallelism of lines, but not their lengths or their angles. It extends the degrees of freedom of the rigid transformation with a scaling factor for each image dimension (Erdi et al., 2002; Blackall et al., 2000;, Radau et al., 2001; Ritter et al., 1999; Van Laere et al., 2001) and additionally, a shearing in each dimension [Jenkinson & Smith, 2001; Kagadis et al., 2002; Viola & Wells, 1995, 1997; Zagrodsky, Shekhar, & Cornhill, 2001). In Hill et al. (1998) and Holden (2001) an affine registration with nine degrees of freedom is performed to correct calibration errors in the voxel dimensions. Holden (2001) furthermore

measures the *relative* scaling error between scans. Shekhar & Zagrodsky (2002) compare registration of ultrasound images using transformations of increasing complexity (rigid, rigid with uniform scaling, rigid with nonuniform scaling and fully affine).

3.5.3.3 Curved

Curved registration methods can differ on several aspects. The mutual information measure can be calculated globally, on the entire image, or locally, on a subimage. Smoothness of the deformation can be achieved in different ways and the deformation can be either free-form (any deformation is allowed) or guided by an underlying physical model of material properties, such as tissue elasticity or fluid flow. Besides these aspects, methods can also differ in smaller, implementational details, but such differences will not be discussed.

Meyer et al. (1999), Kim, Boes, Frey, & Meyer (1997), Krücker et al. (2000) and Meyer et al. (1997) compute the mutual information measure globally. The deformation is determined by thin-plate splines through a number of control points, which are initialized by the user, but are adapted automatically. The number of control points defines the elasticity of the deformation. Apart from registration of 3D multimodality images, the method was applied to warp a slice into a volume, including out-of-plane deformations by Kim, Boes, Bland, & Meyer (1999). Also computing both measure and deformation globally is proposed by Horsfield (1999), who uses a third-order polynomial to nonuniformly correct MR images for eddy current distortion.

Other methods compute the mutual information globally, but find the deformation on a local scale. A grid of control points is defined to determine the deformation, usually in a multiresolution manner. The points of the grid are moved individually, defining local deformations. Transformations in between control points are propagated by linear interpolation (Baker, Busse, & Vogt, 2000; Kjems et al., 1999), Gaussian kernels (Gaens, Maes, Vandermeulen, & Suetens, 1998) or other

symmetrical, convex kernels (Rohde, Aldroubi, & Dawant, 2001). Rueckert et al. (2001) and Studholme et al. (2001) calculate B-splines through the control points, which have a local region of influence (as opposed to thin-plate splines). A similar method is employed in Rohlfing & Maurer (2001), Rohlfing, Maurer, O'Dell, & Zhong (2001) and Mattes et al. (2001). The effect of the choice of transformation (rigid, affine or curved) on registration of MR breast images was studied by Denton et al. (1999). The method by Rueckert was adapted to allow for rigid structures within deformable tissue by Tanner et al. (2000), through fixation of intercontrol point distances. A nonuniform deformation grid of active and passive control points is described in Schnabel et al. (2001). Applications of the method include propagation of segmentations (Castellano et al., 2001; Holden et al., 2001) and the construction of a statistical deformation model (Rueckert, Frangi, & Schnabel, 2001).

Contrary to the previous methods which compute mutual information globally, some methods compute the mutual information measure for subsets of the images (Hata et al., 1998; Hayton et al., 1999; Hellier & Barillot, 2000). A problem with local computation of mutual information is that the results can suffer from the small number of samples. Usually, relatively large subimages are required, which prohibits deformations on a very small scale. Several adaptations have been proposed to overcome this problem. Likar & Pernus (2001) define local probabilities as a weighted combination of the probability distribution of a subimage and the global distribution. Maintz, Meijering, & Viergever (1998) compute a conditional probability distribution of intensities in one image given intensities in the other image, based on a global joint histogram. Using the conditional distribution, translations of subimages are computed. Finally, Rueckert et al. (2000) enhance the power of locally computed measures by including spatial information, in the form of cooccurrence matrices.

Hermosillo & Faugeras (2001) compare global and local computation of both mutual information and the correlation ratio. Schnabel et al. (2001) propose a validation method for curved registration which is demonstrated on mutual information.

Most methods ensure smoothness of the deformation field, by filtering of the vector field. Rohlfing & Maurer (2001) incorporate a regularization term that prevents compression of contrast-enhanced structures.

3.5.4 Implementation

The importance of the implementation of a mutual information based method should not be underestimated, since implementational decisions can have a large influence on the registration results. The main choices involve interpolation, estimation of the probability distributions and optimization. Additionally, one may choose to improve the speed of registration. Zhu & Cochoff (2002) study the influence of several implementation choices, optimization method, interpolation method, number of histogram bins and multiresolution approaches. The choice of implementation remains a matter of debate. An optimal implementation has not been agreed on, partly because all aspects of the implementation interact.

3.5.4.1 Interpolation

When transforming points from one image to another, interpolation is usually required to estimate the grey value of the resulting point. In this section we focus on interpolation *during* the registration process, which is applied numerous times and which, consequently, necessitates a tradeoff between accuracy and speed. In addition, interpolation is required to yield a final, registered image. Since this task is performed only once, speed is less of an issue and a different choice of interpolation method (e.g. a higher-order method) may be more appropriate.

The most popular technique of interpolation is linear interpolation, which defines the intensity of a point as the weighted combination of the intensities of its neighbours. The weights are linearly dependent on the distance between the point and its neighbours. A handful of papers report the use of nearest neighbour interpolation (assigning the grey value of the spatially closest neighbour), often for

speed (Plattard et al., 2000; Radau et al., 2001), for comparison to other interpolation methods (Maes et al., 1997; Zhu & Cochoff, 2001) or for the initial testing of a novel idea (Rodriguez et al., 1998).

An interpolation method specifically designed to create joint histograms of intensities is partial volume interpolation, introduced by Collignon et al. (1995). It uses the weights of linear interpolation, but not to compute a weighted intensity and update a single histogram entry, like linear interpolation. It uses the weights for fractional updates of the histogram entries corresponding to a transformed point and each of its neighbours. Effectively, this creates smoother changes of the joint histogram for varying transformations and hence a smoother registration function. The method has been adopted by several others (Hellier & Barillot, 2000; Maes et al., 1997; Shekhar et al., 2002; Pluim et al., 2001).

Maes (1998) introduced partial intensity interpolation. This method calculates a weighted average of the neighbouring grey values, identical to linear interpolation. Then, however, two histogram entries (those corresponding to the floor and the ceiling of the weighted average) are updated by a fractional amount.

Thévenaz & Unser (2000) are advocates of higher order interpolation methods. They suggest cubic interpolation, particularly in multiresolution methods. Cubic spline interpolation is also used in Freire et al. (1999) and Mattes et al. (2001).

A serious problem with interpolation is that it can cause patterns of artefacts in the registration function. When the grids of two images can be aligned for certain transformations, no interpolation is required for such transformations. Because interpolation influences the value of the registration measure, the absence of interpolation (at grid-aligning transformations) can cause a sudden change in the value of the measure, resulting in a pattern of local extrema. The occurrence of such patterns has been noted in several publications (Collignon, 1998; Maes, 1998). In Pluim et al. (2000), the different patterns created by linear and partial volume interpolation are extensively studied. Holden (2001) describes the existence of

artefacts for both mutual information and the ratio image uniformity (Woods et al., 1998) measures when using linear interpolation and proposes low-pass filtering as a solution. Likar & Pernus (2001) try to overcome the severe artefacts in the registration functions of subimages, either by a random resampling of the image grids or by including the probability distribution of the entire images. Chen & Varshney (2001) employ a generalized partial volume interpolation method, which is identical to partial volume interpolation using a higher-order kernel instead of a first-order one. Interpolation artefacts deserve serious attention, not only because they can cause misregistrations, but also because they prohibit subvoxel accuracy.

3.5.4.2 Probability Distribution Estimation

The most straightforward way to estimate the joint probability distribution of intensities in two images is to compute a joint histogram of intensities. Each entry $h(a, b)$ in the histogram denotes the number of times intensity a in one image coincides with b in the other image. Dividing the entries by the total number of entries yields a probability distribution. The probability distributions for each image separately are found by summing over the rows and columns respectively of the histogram. This method is chosen in the majority of papers (Blackall et al., 2000; Collignon et al., 1995; Estevez et al., 2009; Gaens et al., 1998; Hellier & Barillot, 2000; Jenkinson & Smith, 2001; Knops et al., 2006; Likar & Pernus, 2001; Maintz et al., 1998; Meyer et al., 1997; Plattard et al., 2000; Radau et al., 2001; Slomka et al., 2001; Studholme et al., 1999; Tanner et al., 2000, Yokoi et al., 2004).

Another frequently used method of distribution estimation is Parzen windowing. Given a set S of n samples, the probability $p(x)$ of x occurring is the sum of the contributions of each sample s from S to $p(x)$. The contributions are functions of the distance between s and x. This results in the following definition of the probability of x given a:

$$P(x, S) = \frac{1}{n} \sum_{s \in S} W(x - s) \qquad (3.11)$$

The weighting function *W* is a Gaussian function in most implementations described in the literature (Baker et al., 2000; Hermosillo & Faugeras, 2001; Viola & Wells, 1997). Other choices are double exponential functions (Kim et al., 2000) and splines (Mattes et al., 2001; Thévenaz et al., 2000).

Whereas the simple histogram method places a spike function of value 1 at the bin corresponding to *s* and updates only a single bin, Parzen windowing places a kernel at the bin of *s* and updates all bins falling under the kernel with the corresponding kernel value.

3.5.4.3 Optimization

The registration measure as a function of transformation defines an *n*-dimensional function, with *n* the degrees of freedom of the transformation. The optimum of this function is assumed to correspond to the transformation that correctly registers the images. Unfortunately, the registration function is generally not a smooth function, but one containing many local maxima. The local maxima can have two different causes. Some represent a local good match of the two images. Others are imperfections inherent to the implementation, for example, local maxima can occur as a result of interpolation or because of changes in the overlapping part of the images. Local maxima in the registration function can be reduced, among other things, by improving implementation choices (e.g. a higher order interpolation method), by filtering the images to reduce noise or by increasing the bin size of the intensity histogram. Because of the existence of local maxima, the choice of optimization routine has a large influence on the results of the registration method, particularly on the robustness of the method with respect to the initial transformation.

A second important property of the registration function that influences the choice of optimization method is the capture range of the optimum (Hajnal et al., 2001; Hill et al., 2001; Studholme et al., 1997; Studholme et al., 1999). For intensity-based registration measures it is possible that a large misregistration of two

images results in a higher value of the measure than the correct transformation. The desired maximum may not be the global maximum of the search space and only part of the search space leads to the desired maximum. This has two consequences for optimization of the registration function. First of all, an optimization started outside the capture range of the desired maximum has little chance of leading to a correct registration of the images. Secondly, probabilistic optimization routines, such as some multistart methods and genetic algorithms, may prove to be less suitable for optimization of the mutual information measure, because they can move outside the capture range. The extent of the capture range depends on the registration measure and on image properties (e.g. modality, contents, field of view) and cannot be determined a priori.

A popular method is Powell's routine, which optimizes each transformation parameter in turn. It does not require function derivatives to be calculated, but is relatively sensitive to local optima in the registration function (Collignon, 1998; Fei et al., 2002; Likar & Pernus, 2001; Maes et al., 1997; Pluim et al., 2001; Rodriguez et al., 1998; Zhu & Cochoff, 2001).

Equally popular is the Simplex method, which does not require derivatives either, but, contrary to the previous method, considers all degrees of freedom simultaneously (Horsfield, 1999; Kim et al., 2001; Meyer et al., 1997; Meyer et al., 1999; Radau et al., 2001; Shekhar et al., 2002; Slomka et al., 2001; Van Laere et al., 2001; Zagrodsky et al., 2000, 2001; Zhu & Cochoff, 2001). It is not known for its speed of convergence.

Plattard et al. (2000) use a combination of the Powell and Simplex methods, whereas Kagadis et al. (2002) combine Powell and a genetic algorithm. Jenkinson & Smith (2001) propose an optimization routine that extends Powell's method with initialization and a multi-start technique.

Rodriguez & Loew (1999) combine Powell with topographical global optimization. This involves a graph structure with the nodes denoting points in the

search space and the arcs pointing in the direction of nodes with lower function values. In this manner, the influence zones of local maxima can be determined and a number of local maxima is selected based upon the graph to start optimizations from.

Although being one of the simplest optimization techniques, hill-climbing optimization was shown to produce good results in a multiresolution scheme, with the step size of the hill-climbing method decreasing as the image resolution increased (Studholme, 1997; Studholme et al., 1999).

Methods that do require function derivatives (whether mathematically derived or numerically estimated) are gradient ascent (Hermosillo & Faugeras, 2001; Rohde et al., 2001; Rohlfing et al., 2001; Rohlfing & Maurer, 2001; Rueckert et al., 2001; Studholme et al., 2001; Viola & Wells, 1997), quasi-Newton methods (Mattes et al., 2001) and the method by Levenberg-Marquardt (Thévenaz et al., 2000).

A method little used in image registration is simulated annealing, which has the seemingly paradoxical property of sometimes taking a step in the "wrong" direction (i.e. towards a smaller function value when the goal is maximization) (Nikou et al., 1999; Ritter et al., 1999). This move is allowed occasionally to make escapes from local maxima possible. Equally uncommon are genetic algorithms, which are based on the survival-of-the-fittest principle of combining current elements and selecting the best of the new elements. Another unconventional approach presented by Spall (1998, 2003) called SPSA which uses gradient estimations with simultaneous perturbation of parameters to find global optimum.

An unconventional approach of finding the optimal transformation is employed in (Ding et al., 2001). Template matching of subimages is used to define a set of corresponding points (the centre points of the subimages), based upon which a rigid transformation is determined.

To improve the chances of finding the global optimum of the registration function, Chen and Varshney (2001) compute the mutual information both of the

entire images and of four subimages, assuming that when the global mutual information is maximum, this should also hold for subimages. Zagrodsky et al. (2001) use the mutual information value of three intensity histograms of different bin widths simultaneously to find the optimal transformation.

Optimization is often performed in a multiresolution manner, as this is expected to decrease the sensitivity of the method to local maxima in the registration function. The term multiresolution can be used with respect to the images, in the sense that the images are down or up-scaled to a number of resolution levels (Fei et al., 2002; Gaens et al., 1998; Hayton et al., 1999; Hermosillo & Faugeras, 2001; Kjems et al., 1999; Lynch et al., 1999; Penny et al., 1998; Radau et al., 2001; Rohlfing et al., 2000; Studholme et al., 1997, 1999; Thévenaz et al., 2000; Zhu & Cochoff, 2001). Multiresolution can also apply to the deformation grid of curved registration methods (Hayton et al., 1999; Hellier & Barillot, 2000; Kjems et al., 1999; Mattes et al., 2001; Nikou et al., 1999; Rohde et al., 2001; Rohlfing et al., 2001; Rueckert et al., 2001; Schnabel et al., 2001; Studholme et al., 2001)

3.5.5 Image Dimensionality

The majority of papers treats registration of three-dimensional images. We will next discuss the exceptions: two 2D images or a 2D and a 3D image.

3.5.5.1 2D/2D

The difficulty with two-dimensional images is that the number of samples usually is substantially smaller than with three-dimensional images. This can result in a less reliable estimation of the probability distributions. Good results have been reported nonetheless. The choice for two-dimensional images is often guided by the application (Baker et al., 2000; Kim et al., 1997; Likar & Pernus, 2001; Plattard et al., 2000; Ritter et al., 1999). Other times two-dimensional images are chosen for initial testing of a novel idea, frequently with the intention of extension to three dimensions (Rodriguez et al., 1998).

3.5.5.2 2D/3D

Registration of two and three-dimensional images is regularly applied to find the correspondence between the operative scene and a preoperative image. Viola & Wells (1995), for example, devised a method of using mutual information to register 2D video images to a model of a 3D object (usually based on an MR or a CT image). Bansal et al. (1999) register 2D portal images to a preoperative CT in order to verify the position of the patient with respect to the radiotherapy treatment plan. They propose an iterative approach, which switches between segmenting the images based on the current registration and registering the images based on the current segmentation.

Kim et al. (1999) correct for motion in fMRI acquisitions by registering fMRI slices into a 3D anatomical MR scan. In Kim et al. (1999) out-of-plane deformation of the slices is introduced. Calibration of an ultrasound probe using 2D/3D registration is described by Blackall et al. (2000), registering 2D B-mode ultrasound images to an MR volume to allow reconstruction of a 3D ultrasound image.

A comparison of six intensity-based registration measures, for registration of a 2D fluoroscopy image to a CT volume, has been made by Penney et al. (1998).

3.5.6 Number of Images

Commonly, two images are involved in the registration process. However, in certain situations several images of a scene are to be registered or a series of images taken at different times needs to be compared.

When more than two images are employed, two types of registration problems can be distinguished: with known and with unknown inter-image geometry. In the first case, the transformations between several images are known and only a single transformation has to be determined. In the second case, no knowledge about the

transformation between individual images is available and multiple transformations are to be found to transform the images to a common coordinate system.

More than 2 images, with known inter-image geometry: An example of the first type is the problem of determining the position of a 3D object among a number of 2D images of the object, taking from different, known, angles. Several different mutual information based solutions have been proposed for this problem. One could simply sum the measures of each 2D image and the 3D image or combine the intensity correspondences for each 2D image and the 3D image in a single joint histogram (Bansal et al., 1999; Kim et al., 2001). Clarkson, Rueckert, Hill, & Hawkes (1999) have compared three methods of combining measures (adding the measures of individual 2D images and the 3D image, alternating between separate optimizations or creating a single two-dimensional histogram of intensity correspondences in each of the 2D images and the 3D image). This was applied to registration of 2D video images and a CT volume. A similar problem, that of registering an MR volume to a set of 2D ultrasound images, is tackled by Blackall et al. (2000) by gathering the corresponding intensity pairs for each slice and the volume into a single joint histogram.

More than 2 images, with unknown inter-image geometry. One instance of a registration problem from the second class (requiring more than one transformation) is described by Lynch et al. (1999), who register three images. They circumvent the problem of having to optimize several transformations simultaneously by first registering two images. The third image is registered to the previous two using a two-dimensional intensity distribution for the registered images, which results in a higher-dimensional mutual information measure.

Krücker et al. (2000) register several 3D ultrasound scans, acquired under different angles, to form a compounded 3D image with a better signal-to-noise ratio. The first scan is used as the reference scan to register all subsequent scans to.

Images of a patient that have been taken over a period of time need registration to study changes through time. Usually, the first recorded image acts as a reference to which all subsequent images are registered (Baker et al., 2000; Denton et al., 2000; Holden et al., 1999). Kim et al. (1999) correct for patient motion during the acquisition of fMRI time series by registering each slice into an anatomical volume.

3.5.7 Modalities

Mutual information has been applied to a wide variety of image modalities. These can be subdivided into applications of monomodality images, of multimodality images, of an image and a model and of an image to physical space (e.g. using intraoperative images of a patient).

3.5.7.1 Monomodality

Even though, when first introduced, one of the main advantages of mutual information was its capability to register multimodality images, the measure has also been shown to be well suited to registration of images of the same modality. The following is a brief overview of modalities found in the literature. Although MR images can have very different characteristics for different scanning protocols, we have classified all MR registration problems as monomodality.

MR: Registration of MR images has been described in many publications (Fei et al., 2002; Rodriguez et al., 1998), often for curved transformations (Denton et al., 1999; Gaens et al., 1998; Hata et al., 1998; Hayton et al., 1999; Holden, Denton, et al., 1999; Holden, Hill, et al., 2000; Kim et al., 1999; Kjems et al., 1999; Nikou et al., 1999; Rohde et al., 2001; Rohlfing et al., 2001; Rueckert, Clarkson, et al., 2000; Rueckert, Frangi, et al., 2001; Studholme et al., 2001; Tanner et al., 2000). Some results of registering interventional MR images can be found in Carrillo et al. (2000). Time series of fMRI images require registration to detect changes in brain function. Changes in brain anatomy are studied in Denton et al. (2000) and Holden et al. (2001). Furthermore, registration is needed to map the functional information onto an

anatomical MR scan (Hellier & Barillot, 2000; Rohlfing et al., 2000; Studholme et al., 2001). Registration to correct for eddy current distortions in diffusion-weighted images is described in Horsfield (1999).

CT: In the mutual information literature, CT is usually combined with other modalities and few monomodality cases have been reported. Extraction of the midsagittal plane is described by Liu, Collins, & Rothfus (2001). Martens et al. (1999) use registration of pre- and postoperative CT images to validate pedicle screw placement, whereas Bergmans et al. (2001) validate root canal treatment.

SPECT: Holmes, Camp, & Robb (1996) compare mutual information with a measure similar to Woods' measure (Woods et al., 1992). Radau et al. (2001) compare normalized mutual information with two other measures for the registration of SPECT images to an atlas, created by averaging of a number of SPECT images. Registration of transmission images to achieve alignment of the corresponding emission images is described by Van Laere et al. (2001).

PET: In Holmes' comparison of mutual information and Woods' measure, PET-PET registration is one of the modality combinations described (Holmes et al., 1996).

US: Meyer et al. (1999) use mutual information to match breast ultrasound images, whereas Zagrodsky et al. (2000, 2001) register two series of cardiac images. Shekhar et al. (2002) study the effect of median filtering, number of histogram bins and interpolation method on the smoothness of the registration function of cardiac ultrasound images.

Microscopy: Registration of histological sections has been reported both using rigid transformations (Capek & Krekule, 1999; Ourselin, Roche, Prima, & Ayache, 2000) and curved ones (Likar & Pernus, 2001).

X-ray: Sanjay-Gopal et al. (1999) compare mutual information and the correlation coefficient for registration of intrasubject mammograms. Plattard et al. (2000)

register both 2D portal images and portal to x-ray images to verify the position of the patient with respect to previous radiotherapy treatment sessions.

3.5.7.2 Multimodality

Mutual information has been studied for many combinations of modalities.

MR-CT: A popular combination, and one of the earliest described, is registration of MR and CT images (Capek et al., 1999; Collignon, 1998; Ding et al., 2001; Holmes et al., 1996; Maes et al., 1997; Maes, 1998; Maintz et al., 1998; Pluim et al., 2000; Rodriguez & Loew, 1999; Studholme et al., 1996, 1999; Studholme, 1997; Thévenaz et al., 2000). A category are the papers that report on registering what are commonly known as the "RREP" or "Vanderbilt" images (Capek et al., 1999; Ding et al., 2001; Maes et al., 1997; Studholme et al., 1999; Thévenaz et al., 2000). These images are publicly available [Vanderbilt] .

MR-PET: A variety of applications of MR-PET registration has been recounted (Collignon, 1998; Ding et al., 2001; Holmes et al., 1996; Kjems et al., 1999; Maes et al., 1997; Maes, 1998; Meyer et al., 1997; Pluim et al., 2000; Studholme et al., 1996, 1999; Studholme, 1997; Thévenaz et al., 2000). The RREP images mentioned above also include MR and PET image pairs.

MR-SPECT: Comparisons between mutual information and other measures for registration of MR and SPECT images are made in (Holmes et al., 1996; Nikou et al., 1999). Other publications on the subject of MR-SPECT matching are (Capek et al., 1999; Lynch et al., 1999; Zhu & Cochoff, 2001).

CT-PET: Both Erdi et al. (2002) and Mattes et al. (2001) register CT and PET *transmission* images of the thorax to achieve fusion of CT and PET emission images. CT-PET registration of the thorax is furthermore described by Meyer et al. (1997).

CT-SPECT: Meyer et al. (1997) also registered CT and SPECT images, now focusing on the abdomen. Kagadis et al. (2002) compare a surface-based and a mutual information based registration routine.

3.5.7.3 Modality to Model

By a model denoted any kind of simplified or processed image. A model can be a simulated image [Castellano et al., 2001; Holden et al., 2001; Jenkinson & Smith, 2001] or a segmentation (Gaens et al., 1998). Another possibility is an average image or a statistical model composed of several images (Kjems et al., 1999; Radau et al., 2001; Van Laere et al., 2001).

3.5.7.4 Modality to Physical Space

A previously acquired image of a person can be registered to the actual person, via an intraoperative image. This is what we term 'registration to physical space'. A common application in radiotherapy is the verification of patient position with respect to a treatment plan based on a previously acquired image. Usually, this involves registration of a pretreatment CT to portal (Bansal et al., 1999; Plattard et al., 2000) or fluoroscopy images (Kim et al., 2001).

Registration to physical space is also required in image-guided treatment, for transferring the information of a pretreatment image (and any treatment plans based upon the image) to a patient on the operating table. Preoperative images are registered to intraoperatively acquired images, such as video images from the operation microscope (Clarkson et al., 1999), ultrasound images, fluoroscopy images (Penny et al., 1998), PET transmission images (Erdi et al., 2002) or interventional MR images (Carrillo et al., 2000).

3.5.8 Subject

The subject in the images to be registered can be the same (intrasubject registration), can differ (intersubject registration) or one of the images can be a model. Intersubject registration based on mutual information is a highly relevant topic, because it can form the basis for methods such as tissue segmentation, bias field correction in MR images and analysis of images of groups of subjects. Only a small percentage of the references deal with intersubject registration.

3.5.9 Object

The object in medical image registration is the part of the anatomy involved. There is a varied list of objects, which is summarized in this section.

Brain: A large part of the literature of mutual information based registration concerns head or brain images. (Collignon et al., 1995; Collignon, 1998; Denton et al., 2000; Ding et al., 2001; Freire et al., 1999; Holden, Denton, et al., 1999; Holden, Hill, et al., 2000; Holden, Schnabel, et al., 2001; Jenkinson & Smith, 2001; Kagadis et al., 2002; Kim et al., 1999; Kjems et al., 1999; Maes et al., 1997; Maes, 1998; Meyer et al., 1997; Nikou et al., 1999; Plattard et al., 2000; Radau et al., 2001; Rodriguez & Loew, 1999; Rodriguez et al., 1998; Rueckert et al., 2000; Studholme, 1997; Studholme et al., 1996, 1997, 1999; Studholme, Novotny, et al., 2001; Thévenaz et al., 2000; Viola, 1995; Viola & Wells, 1997; Zhu & Cochoff, 2001;).

Thorax/lungs: there is a small amount of paper about thorax and the lungs (Erdi et al., 2002; Kim et al., 2001; Meyer et al., 1997).

Heart: Zagrodsky et al. (2000, 2001) register two series of cardiac ultrasound images, with each series of images a sequence of heart cycles. In Shekhar et al. (2002), several adaptations to the method are made to allow affine registration.

Breast: Registration of breast images has been described for various imaging modalities, in particular, MR (Hayton et al., 1999; Denton et al., 1999; Rueckert et al., 2001; Tanner et al., 2000), x-ray and ultrasound (Krücker et al., 2000; Meyer et al., 1999).

Abdomen/liver: Several papers have been published on registration of abdominal images (Meyer et al., 1997; Rodriguez et al., 1998) and of the liver (Carrillo et al., 2000; Rohlfing et al., 2001).

Pelvis: As registration of pelvic images is quite a challenging task, almost all references given propose some adaptations (whether large or small) to the standard method of mutual information based matching (Bansal et al., 1999; Fei et al., 2002; Studholme et al., 1996).

Tissue: Histological sections of tissue are the object of registration in a number of studies (Kim et al., 1997; Likar & Pernus, 2001).

CHAPTER FOUR
METHODS

Elements of a general registration process are a geometric transformer, a measure and an optimizer as shown in Figure 4.1. Registration as described here is an iterative process where one image is geometrically transformed within a predetermined parameter space and compared against the other. A measure of similarity is computed between the images at each iteration and used to determine if they are "sufficiently" aligned. This process is controlled by the optimizer that starts from an initial guess and determines subsequent parameter values for the next iteration to reach an optimal alignment.

The methods which are used in thesis are categorized in process order and described in detail with four sections: preprocessing, transformer, measure and optimizer.

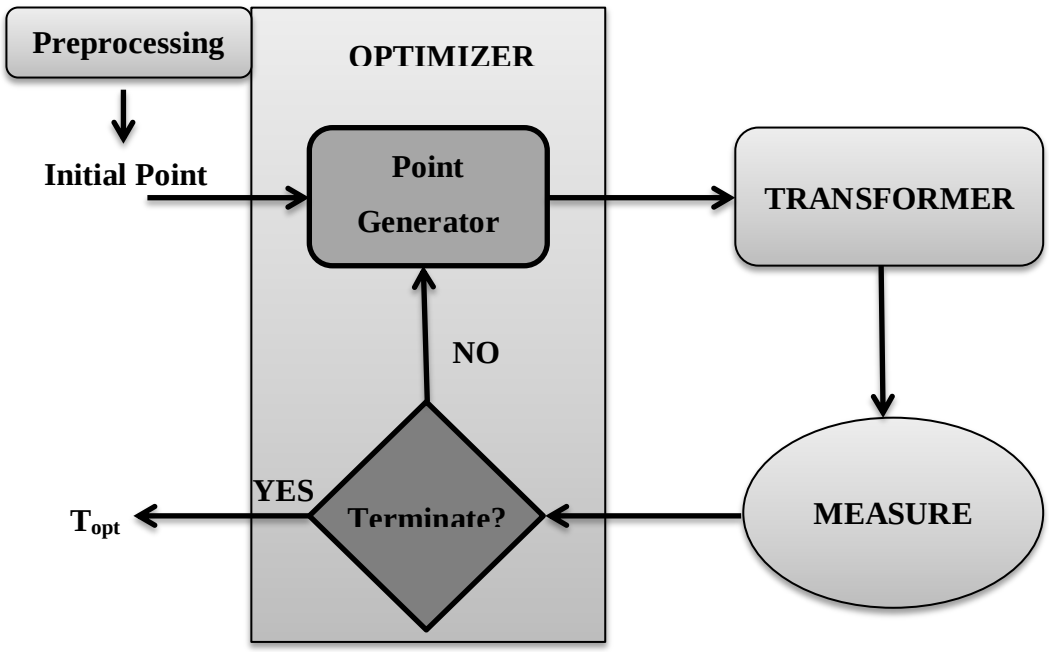

Figure 4.1 Main components of a general registration process.

4.1 Preprocessing

Preprocessing includes a wide range of tasks that may be performed on images outside the optimization loop and at the beginning of the process. Preprocessing may include filtering, resampling, rectification, gradient computation, pyramid construction, feature detection, etc.

As a first step of our preprocessing, volume image pair resolutions are equalized if there is an inequality between corresponding axes'.

Secondly, parameter search space limitations should be defined for SPSA optimization algortihm. There are eight parameters to be searched for in this study. This parameters could vary in a wide range. To overcome this variability, the parameter search space is normalized to [-1, +1].

Thereafter, the centre of massess of image pairs (CM_1, CM_2) are calculated in order to use for initial guess. As mentioned before, initial guess is very important for optimization algorithm's convergence accuracy, especially for stochastic approaches. At this point the initial guess, $p_0 = [\theta_x, \theta_y, \theta_z, t_x, t_y, t_z, S_x, S_y]^T$, should be chosen as close as possible to the real registration parameters. The difference between the center of massess is given as initial points for translation parameters, while the others are chosen as center point of parameter search space. Translation parameters are, for x axes $t_{xm} = CM_{1x} - CM_{2x}$, for y axes $t_{ym} = CM_{1y} - CM_{2y}$ and for z axes $t_{zm} = CM_{1z} - CM_{2z}$.

As a consequence, after all this assumptions the initial guess vector becomes $p_0 = [0, 0, 0, t_{xm}, t_{ym}, t_{zm}, 0,0]^T$.

Finally, for multiresolution image pyramid approach, three levels of decomposition are used. These volumes correspond to a decimation of 4, 2 and 1 of the original volume respectively. At lowest decomposition level, initial guess vector

is given to optimization algorithm. Vector of optimization results from a lower resolution levels is transferred as initial guess vector to the next decomposition level.

Mathematical background of multiresolution image pyramid approach is described in the next section.

4.1.1 Multiresolution Image Pyramid

Multiscale data representation is a powerful idea which captures data in a hierarchical manner where each level corresponds to a reduced-resolution approximation. The representation should allow images to be successively approximated starting from a coarse version and going to fine-resolution version. There are two main category of multiresolution approach; signal based which contains Gaussian and Laplacian filters, and graph based which contains stochatic and adaptive pyramid structure. In this thesis Gaussian pyramid approach is used.

A Gaussian pyramid is a technique involves creating a series of images which are weighted down using a Gaussian average. When this technique is used multiple times, it creates a stack of successively smaller images, with each pixel containing a local average that corresponds to a pixel neighborhood on a lower level of the pyramid.

The Gaussian pyramid consists of low-pass filtered, downsampled images of the preceding level of the pyramid, where the base level is defined as the original image. More formally, let the two-dimensional original image be denoted by *I(x, y)*. The Gaussian pyramid is defined recursively as follows,

$$G_0(x,y) = I(x,y),\ for\ level, l = 0 \tag{4.1}$$

$$G_l(x,y) = \sum_{m=-2}^{2} \sum_{n=-2}^{2} w(m,n) G_{l-1}(2x+m, 2y+n) \tag{4.2}$$

where *w(m,n)* is a weighting function (identical at all levels) termed the generating kernel which adheres to the following properties: separable, symmetric and each

node at level *n* contributes the same total weight to nodes at level *l*+1. An example of *w(m,n)* is the 5-tap filter: 1/16 [1 4 6 4 1]. The weighting function closely approximates the Gaussian function, hence the origins of the pyramids name. Alternatively, the same result can be realized by applying an equivalent weighting function denoted $w_1(m, n)$ directly to the original image, followed by *l* downsampling operations, where *l* denotes the level number. The equivalent weighting function approximates a Gaussian function that doubles in scale with each level. An alternative view of the Gaussian pyramid is that each element of the pyramid represents a local average obtained with the equivalent weighting function applied to the original image. Thus the Gaussian pyramid contains local averages at various scales.

4.2 Transformer

The transformer uses a geometric transformation function to map points in one image (float image) to new coordinates in the transformed image. The type of function used in image registration depends on the type and severity of the geometric difference between the images. Transformation functions are classified into two categories, collinear and deformable. A collinear transformation maps straight lines onto straight lines and preserves the parallelism between lines. Rigid, affine, and projective transformations are collinear and can be stated by a 4 x 4 matrix operating on homogeneous vectors representing 3D points. While rigid transformations consist of only translations and rotations, affine transformations are allowed to include scaling as well. An affine transformation can furthermore include shearing. As to deformable transformations they are either parametric or nonparametric. In parametric methods, a set of control points selected in the image domain are used to perform some piecewise polynomial interpolation of a displacement field. Nonparametric methods define the registration problem in a variational setting such that a transformation is modeled by an arbitrary displacement field regularized by some smoothing criteria (Modersitzki, 2004).

The transformer is also responsible for determining the intensity of the points in the transformed image by interpolating intensity values of corresponding points in the float image. This process is also called resampling. The nearest neighbor interpolation is the simplest and fastest interpolation method, but it is seldom used in practice since it usually results in inaccurate evaluation of similarity functions. While the most commonly used interpolation method is trilinear interpolation that provides a good balance between speed and accuracy, other more complex and accurate methods, namely quadratic, cubic, cubic B-spline, and Gaussian interpolation (Lehmann et al., 1999) can also be preferred depending on the application.

In this study, all of our transformations are chosen to be affine. Although all interpolations are linear in the optimization algorithm, there are both linear and cubic interpolations in image pyramid construction stage. In order to optimize speed/accuracy stability, cubic interpolation is used at lower resolutions, while linear interpolation is used at the highest resolution level of image pyramid.

4.3 Measure

All automatic registration methods need a measure of similarity or distance between images. Where the images are in perfect alignment, a similarity measure reaches its maximum, and declines as the images move farther away. In contrast, a distance measure reaches its minimum where the images are aligned. Various similarity and distance measures exist for unimodality such as sum of squared differences, gradient correlation, ratio-image uniformity, normalized cross correlation, correlation coefficient. For multimodality, mutual information (MI), normalized mutual information (NMI) and correlation ratio are known to be suitable similarity measures. It should be remarked that, different intensity distortions between the images may require different similarity measures depending on their effectiveness. The unimodality measures can be computed independently at each spatial location followed by integrating them over the overlapped image region. Multimodality measures are computed by estimating statistical (mutual information) or functional (correlation ratio) dependence between images that are assumed to be

instances of two discrete random variables. The methods computing multimodality measures have to estimate joint and marginal probability functions of the underlying discrete random variables from image data

4.3.1 Proposed 3-Slice NMI

In image registration methods, defining the objective function to be optimized plays critical role. Especially when working with volume images, objective functions need to be well constructed in a manner of less time consuming as well as computing similarity measure effectively. In this sense, we propose a new fast approach for computing NMI as a similarity measure between two volume images.

According to this approach, the algorithm uses three orthogonal central slices of the volumes to build joint histogram instead of entire volumes. Number of voxels along the three orthogonal axes of a volume is usually not the same, making the orthogonal slices different in size. In order to equalize the contributions of orthogonal slices to joint histogram, incrementing by rational coefficients is used in building histogram instead of slice-independent value of 1.

Rational coefficients are calculated directly from the sizes of slices in voxel. Let us denote volumes to be registered by A and B and central orthogonal slices of A by A_1, A_2, and A_3, and of B by B_1, B_2, and B_3. The total numbers of pixels which will individually be used for histogram are N_1 for A_1 and B_1, N_2 for A_2 and B_2, and N_3 for A_3 and B_3. One of the N_1, N_2 or N_3 can be selected as numerator of the rational coefficients. Our algorithm selects the biggest one as the numerator (say N_3) and builds the partial joint histograms using the following voting law:

$$h_{A_1B_1}(a,b) = h_{A_1B_1}(a,b) + (N_3/N_1) \tag{4.3}$$

$$h_{A_2B_2}(a,b) = h_{A_2B_2}(a,b) + (N_3/N_2) \tag{4.4}$$

$$h_{A_3B_3}(a,b) = h_{A_3B_3}(a,b) + 1 \tag{4.5}$$

where $(a, b) = \left(A_k(x, y), B_k(x, y)\right) for\ 0 \leq a, b \leq 255.$

Finally the joint histogram h_{AB} becomes:

$$h_{AB} = h_{A_1B_1} + h_{A_2B_2} + h_{A_3B_3} \tag{4.6}$$

4.4 Optimizer

The overall performance of a registration algorithm is dependent on the effectiveness of the optimizer. Finding the minimum of dissimilarity measure or the maximum of similarity measure is a multidimensional optimization problem, where the number of dimensions corresponds to the degrees of freedom of the expected geometrical transformation. The only method yielding global extreme solution is an exhaustive search over the entire parameter space of used geometrical transformation. Although it is computationally demanding, it is often used if only translations are to be estimated.

In case of transformations with more degrees of freedom or in case of more complex similarity measures, sophisticated optimization algorithms are required, which help to localize the maxima or minima, respectively. Local methods find a local optimum in the vicinity of an initial point and within their capture range. They may converge to an incorrect alignment if not properly initialized. Global methods, however, find the global optimum within a given range of parameters. They are robust with respect to selection of the initial point but at the cost of slower convergence. Global and local methods may be combined to improve robustness while maintaining a reasonable convergence rate. Frequent additions are multi-resolution (e.g., pyramid) and multi-scale approaches to speed up convergence, to reduce the number of transformations to be examined and to avoid local minima. Many applications use more than one optimization technique, frequently a fast but coarse technique followed by an accurate yet slow one.

Approaches based on the use of gradient estimations tend to be fast, but are sensitive to the presence of local optima. Discussion of these methods can be found in Spall (2003).

In the next section, a stochastic gradient technique which is used in this thesis is explained. It is a gradient-free approach and does not require an explicit derivation of the required gradient vector, but it uses instead an approximation to the gradient.

4.4.1 Simultaneous Perturbation Stochastic Approximation (SPSA)

SPSA was first introduced by Spall in (Spall, 1998), where a detailed description can be found. It has recently attracted attention for solving challenging optimization problems where it is difficult or impossible to obtain an analytic expression for the gradient of the objective function. This is especially true of the MI function, since the probabilities required in the computation of MI are estimated using the joint image histogram. The dependence of the MI function on this discrete histogram makes the computation of its derivative complex. SPSA is based on an easily implemented and highly efficient gradient approximation that relies only on measurements of the objective function to be optimized. It does not rely on explicit knowledge of the gradient of the objective function, or on measurements of this gradient.

Let us denote the objective function to be optimized by L. Three-dimensional affine transformation suggests a parameter search space consisting of 3 rotations, 3 translations and 3 scales. We assume that the scale factor for the z-axis is already known, that is true for most medical volume image datasets taken by 3D biomedical imaging devices. At each iteration, the gradient approximation is based on only two function measurements (regardless of the dimension of the parameter space). An additional function measurement is made at each newly computed point, in order to decide whether to block or to update the parameters. At iteration k, the update law for the parameters is steepest ascent

$$p_{k+1} = p_k + a_k g_k \tag{4.7}$$

where the gradient vector $g_k = [g_k{}^1 g_k{}^2 \dots g_k{}^m]^T$ for the m-dimensional parameter space is determined by

$$(g_k)^i = \frac{L(p_k + c_k\Delta_k) - L(p_k - c_k\Delta_k)}{2c_k(\Delta_k)^i} \tag{4.8}$$

Each element $(\Delta_k)^i$ of the vector Δ_k, takes on a value of +1 or -1, as generated by a Bernoulli distribution, and a_k and c_k are positive sequences of the form

$$a_k = \frac{a}{(k+A+1)^\alpha} \tag{4.9}$$

$$c_k = \frac{c}{(k+1)^\gamma} \tag{4.10}$$

where (a, $c > 0$), ($A \geq 0$), ($0 < \gamma < \alpha < 1$) and (α-2$\gamma > 0$).

The SPSA algorithm is a very powerful technique, which can get through some local maxima of the objective function to find the global maximum because of the stochastic nature of the gradient approximation. All the elements of the parameter vector are randomly perturbed to obtain two measurements of L. The algorithm works by iterating from an initial guess of the optimal parameters, p_0 by using this calculated gradient. Spall (1998) presents sufficient conditions for convergence of the SPSA iterative process in the stochastic almost sure.

4.4.2 Proposed SPSA Gain Parameters Calculation

The choice of the gain sequences (a_k and c_k) is critical to the performance of SPSA (as with all stochastic optimization algorithms and the choice of their respective algorithm coefficients). Spall (1998) gives some guidelines and recommendations to choose and refine the gain sequences. There are almost certain explanations for A, α and γ constants but for constants a and c, the approximations are mostly experimental. Hence, a serious number of experiments is needed for every different dataset. Consequently, calculating these gain sequences automatically

before starting the optimization algortihm will clearly provide substantial convenience.

We propose an algorithm which uses gradients approximations of objective function to be optimized to calculate the gain paramaters *a* and *c*. The algorithm needs the number of gradient approximations (*N*), parameter search space limitations ($\boldsymbol{P}_{max}$, $\boldsymbol{P}_{min}$) and the number of SPSA iteration (N_{SPSA}) wich will be used in SPSA.

At first step, algorithm defines minimum and maximum step sizes, $\boldsymbol{s}_{min}$ and $\boldsymbol{s}_{max}$, for parameters which will be searched in optimization algorithm. To do that, it uses 0.5% of parameter search space range for minimum step size and 5% of parameter search space range for maximum step size:

$$\boldsymbol{s}_{min} = (\boldsymbol{P}_{max} - \boldsymbol{P}_{min})/200 \tag{4.11}$$

$$\boldsymbol{s}_{max} = (\boldsymbol{P}_{max} - \boldsymbol{P}_{min})/20 \tag{4.12}$$

As a second step, the algorithm calculates sum of gradient approximations (g_{sum}) in *N* loops using Eq. (4.8) with two differences: p_k replaced with p_n which is randomly picked within the ranges of parameter space, and c_k replaced with s_n which is randomly picked in the same manner. Hence, Eq. (4.8) becomes:

$$g_n = \frac{L(p_n + s_n\Delta_n) - L(p_n - s_n\Delta_n)}{2s_n\Delta_n} \tag{4.13}$$

where n=1,..,N and Δ_n = +1 or -1, as generated by a Bernoulli distribution.

After *N* iterations, the mean gradient approximation can be found by:

$$g_{mean} = \frac{1}{N}\sum_{n=1}^{N} g_n \tag{4.14}$$

From the SPSA's update rule, Eq. (4.11), the term $a_k g_k$ expected to be a minimum at the last iteration of the optimization process. Based on this assumption

$a_k g_k$ suppose to be equal to the minimum step size, and with replacing g_k with g_{mean}, finally gain sequence a_k can be found by:

$$a_k = s_{min}/g_{mean} \tag{4.15}$$

and combining Eq. (4.15) with Eq. (4.9), the equation for constant *a* becomes:

$$a = (k + A + 1)^{\alpha} s_{min}/g_{mean} \tag{4.16}$$

where k= N_{SPSA}, A = N_{SPSA}/10 and α=0.602. Finally, for gain sequence constant *c*, from SPSA's gradient aprroximation rule $p_k + c_k\Delta_k$, the term $c_k\Delta_k$ expected to be maximum at the first iteration of the optimization process. Therefore, $c_k\Delta_k$ should be equal to maximum step size, $c_k\Delta_k = s_{max}$ and combining this equality with Eq. (4.10) constant *c* becomes:

$$c = (k + 1)^{\gamma} s_{max} \tag{4.17}$$

where k=1, γ = 0.101.

Figure 4.2, the result of an experimental validation, shows that the proposed algorithm for calculation of the gain parameters leads the SPSA algorithm to the registration parameters giving minimum error.

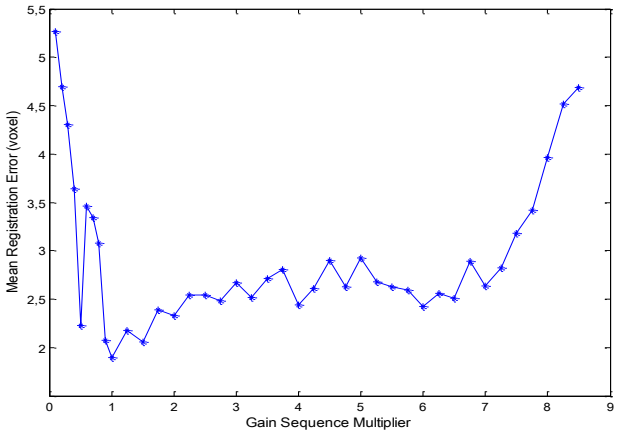

Figure 4.2 Example of mean registration error variation versus multiplier for the calculated gain sequences for a phantom image with 217x181x181 resolutions.

CHAPTER FIVE
RESULTS AND DISCUSSION

In this chapter, our registration method combining multiresolution structure, Studholme's NMI and the SPSA optimization scheme are thoroughly tested and compared using synthetic test data as well as real MR-CT and MR-PET data. SPSA's optimum constants are estimated for every dataset as explained in section 4.4.2. Then based on these results, an automated SPSA optimization scheme is designed and applied to the datasets using NMI. Results are also provided to compare SPSA with Nelder-Mead (NM), Simulated Annealing (SA), Hooke-Jeeves (HJ) and Powell's (PW) optimization algorithms. These algorithms were developed with MathWorks MATLAB® and conducted on a computer which has 8 GB RAM and AMD Phenom™ II X4 955, 3.2 GHz processor.

5.1 Description of Datasets and Parameters

In this work, 3 phantom MR(T1)-MR(T2), 7 MR(PD)-PET, 7 MR(T1)-PET, 6 MR(T2)-PET, 3 MR(PD)-CT, 5 MR(T1)-CT and 4 MR(T2)-CT intrasubject datasets were used (Table 5.1). MR, CT and PET volumes were obtained from web based Retrospective Image Registration Evaluation database (RIRE).

Phantom datasets contain 3 grayscale simulated MR-T1 and MR-T2 brain volumes which were obtained from Brain Web Simulated Brain Database (BrainWeb). Volumes have voxel resolution of 217 x 181 x 181 with 1mm slice thickness 217 x 181 x 60 with 3mm slice thickness and 217 x 181 x 36 with 5 mm slice thickness and all with 3% noise levels and 20% intensity non-uniformities.

MR volumes come from three groups of resolutions which two of them have 1.25 mm pixel spacing and 4mm slice thickness with 256x256x24 and 256 x 256 x 26

voxel resolutions while the last group has 0.82 mm pixel spacing, 3 mm slice thickness and 256 x 256 x 52 voxel resolution.

CT volumes come from four groups of resolutions. Three of them have 0.653595 mm. pixel spacing, 4 mm slice thickness with 512 x 512 x 28, 512 x 512 x 29, 512 x 512 x 230 voxel resolutions, and the last group of CT has 0.42 mm pixel spacing, 3 mm slice thickness with 512 x 512 x 49 resolution.

PET volumes have same resolutions in all datasets which have 2.590723 mm pixel spacing, 8 mm slice thickness and 128 x 128 x 15 voxel resolutions.

For multiresolution approach 3 levels of decomposition are used. These volumes correspond to a decimation of 4, 2 and 1 of the original volume respectively. At the lowest decomposition level, initial guess is given as central point of search space if optimization algorithm used needs an initial guess. Vector of optimization results from a lower level is transferred as initial guess vector to the next decomposition level.

The gain sequences and the constants of the SPSA algorithm are estimated as explained in Section 4.4.2. Furthermore, the upper and lower limits of parameter search space are defined as: (-45^o, 45^o) for rotation parameters in all axes, (-50, +50) for translation parameters in all axes and (0.75, 1.25) for scale parameters in *x* and *y* axes. These parameter search space limits are then normalized to [-1, +1], and initial guess vector is given to SPSA algorithm for lowest decomposition level as central points of search space except the translation parameters which are given as difference between center of masses of volume images to be registered. The numbers of iterations during optimizations are set to 300, 250 and 150 for 3 resolution levels from low to high, respectively. These numbers are 300, 200, and 100 for NM and HJ algorithms, and 500, 300, 200 for SA and PW algorithms. Termination tolerances of objective functions are set to 1e-5 for SA, NM, HJ and PW algorithms. These iteration numbers and termination tolerances are chosen based on a series of experiments which have been conducted using the simulated brain images.

Table 5.1 Descriptions of datasets

DATASET NAME	MODALITIES INVOLVED	RESOLUTIONS
DS-(1-3)	*Phantom MR-T2* *Phantom MR-T1*	*217*181*181 (1mm) - MR T2 - T1 [DS-1]* *217*181*60 (3mm) - MR T2 - T1 [DS-2]* *217*181*36 (5mm) - MR T2 - T1 [DS-3]*
DS-(4-10)	*PET – MR-PD*	*128*128*15 (8mm) - PET* *256*256*26 (4mm) - MR [DS - (4-8)]* *256*256*24 (4mm) - MR [DS - (9-10)]*
DS-(11-17)	*PET – MR-T1*	*128*128*15 (8mm) - PET* *256*256*26 (4mm) - MR [DS-(11-15)]* *256*256*24 (4mm) - MR [DS-(16-17)]*
DS-(18-23)	*PET – MR-T2*	*128*128*15 (8mm) - PET* *256*256*26 (4mm) - MR [DS-(18-21)]* *256*256*24 (4mm) - MR [DS-(22-23)]*
DS-(24-26)	*CT – MR-PD*	*512*512*28 (4mm) - CT [DS-24]* *512*512*29 (4mm) - CT [DS-25]* *512*512*33 (4mm) - CT [DS-26]* *256*256*26 (4mm) - MR [DS-(24-26)]*
DS-(27-31)	*CT – MR-T1*	*512*512*28 (4mm) - CT [DS-27]* *512*512*29 (4mm) - CT [DS-28]* *512*512*33 (4mm) - CT [DS-(29-30)]* *512*512*49 (3mm) - CT [DS-31]* *256*256*26 (4mm) - MR [DS-(27-30)]* *256*256*52 (3mm) - MR [DS-31]*
DS-(32-35)	*CT – MR-T2*	*512*512*28 (4mm) - CT [DS-32]* *512*512*29 (4mm) - CT [DS-33]* *512*512*33 (4mm) - CT [DS-34]* *512*512*49 (3mm) - CT [DS-35]* *256*256*26 (4mm) - MR [DS-(32-34)]* *256*256*52 (3mm) - MR [DS-35]*

5.2 Experimental Results

Since simulated T1 and T2 brain volumes come in exact registration, exact measurement of the registration errors given by the proposed method (or any other registration method) could be determined. However, because the rest of the datasets did not come with a reference registration, the alignments judged by the visual inspection on fused volumes seem to be only validation of registration results with those datasets. Nevertheless, estimated registration parameters are used to calculate relative errors between the registration methods with different optimization algorithms for every dataset. These relative registration errors are tabulated and given just after the related registration results shown figure by figure. All the registration errors are calculated with root-mean-square error (RMSE). The RMSE between the coordinates of the transformed voxels (B'_i) by the estimated transformation matrix and the coordinates of transformed voxels T_i by the actually known transformation matrix is defined as:

$$RMSE = \sqrt{\frac{\sum_{i=1}^{N}(B'_{ix}-T_{ix})^2+(B'_{iy}-T_{iy})^2+(B'_{iz}-T_{iz})^2}{N}} \tag{5.1}$$

where the symbols with sub-indices *x*, *y* and *z* denote the components of the cartesian coordinates of voxels and *N* is the total number of voxels in each volume image.

5.2.1 Phantom Image Registration Results

Phantom image registration results play critical role as it helps to identify algorithm's and/or method's accuracy and robustness.

Phantom image registration results are given below seperately for 1 mm, 3 mm and 5 mm slice thickness' with a figure and a related table. The tables show the average process times and the registration errors with their standart deviations (StD). These averages and standart deviations were calculated from 20 consecutive runs for 3-Slice registration and 10 consecutive runs for 3D registration (which uses entire volume intensites). The results given in the figures are chosen from these runs as close to the averages as possible.

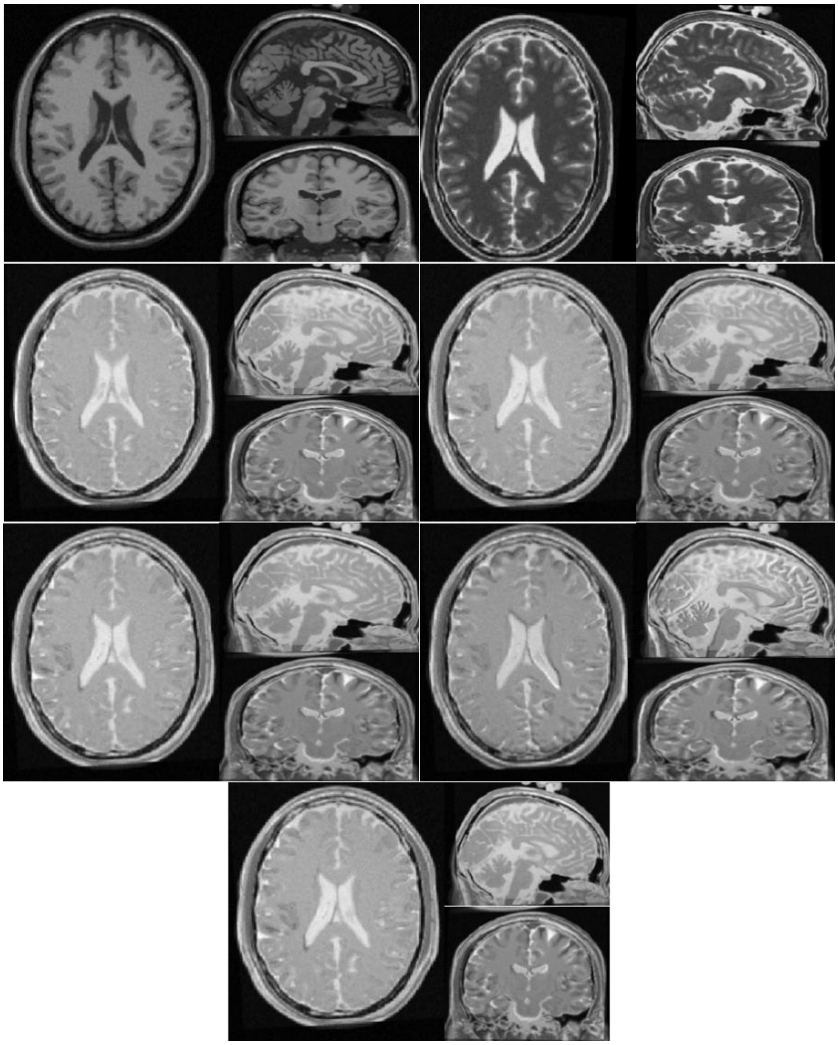

Figure 5.1 DS-1 registration results: First row, reference orthogonal T1 and float orthogonal T2 slices; second row superimposed registered slices from SPSA and NM algorithms; third row superimposed registered slices from HJ and SA algorithms; last row superimposed registered slices from PW

Table 5.2 DS-1 registration results.

Optimization Algorithms	Phantom T2 to T1 Registration (1mm)			
	3-Slice Registration		3D Registration	
	Registration Error (vox) ± StD	Process Time (sec) ± StD	Registration Error (vox) ± StD	Process Time (sec) ± StD
SPSA	*1.42 ± 0.54*	*66.5 ± 0.13*	*1.11 ± 0.08*	*897.6 ± 1.74*
NM	*1,59 ± 0.71*	*242.6 ± 1.4*	*0.63 ± 0.21*	*3602 ± 3.89*
HJ	*1.65 ± 0.13*	*98.5 ± 0.36*	*0.59 ± 0.12*	*2972 ± 2.21*
SA	*2.79 ± 0.44*	*194.3 ± 1.14*	*4.2 ± 3.1*	*8180 ± 58.7*
PW	*1.75 ± 1.02*	*133.5 ± 0.73*	*0.6 ± 0.13*	*2665.1 ± 1.46*

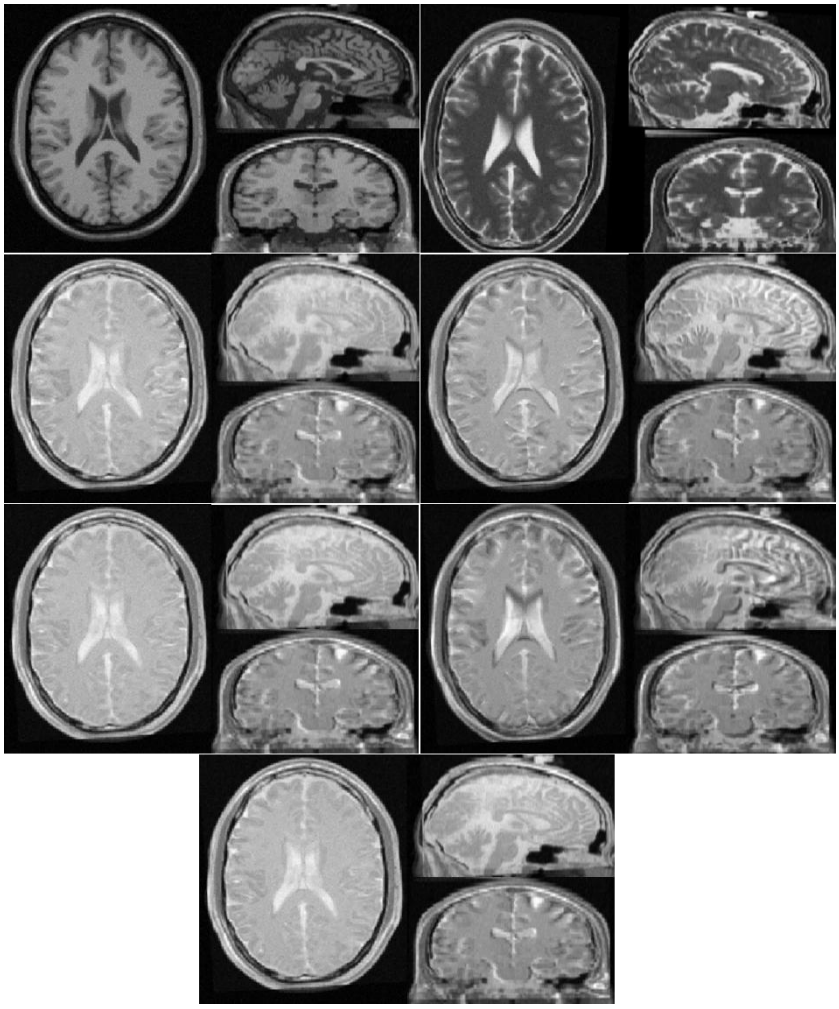

Figure 5.2 DS-2 registration results: First row, reference orthogonal T1 and float orthogonal T2 slices; second row superimposed registered slices from SPSA and NM algorithms; third row superimposed registered slices from HJ and SA algorithms; last row superimposed registered slices from PW.

Table 5.3 DS-2 registration results.

Optimization Algorithms	Phantom T2 to T1 Registration (3mm)			
	3-Slice Registration		3D Registration	
	Registration Error (vox) ± StD	Process Time (sec) ± StD	Registration Error (vox) ± StD	Process Time (sec) ± StD
SPSA	*1.61 ± 0.21*	*38.9 ± 0.07*	*1.92 ± 0.32*	*454.4 ± 2.33*
NM	*2,101 ± 0.32*	*140.8 ± 1.82*	*0.601 ± 0.09*	*1400.1 ± 11.2*
HJ	*0.89 ± 0.04*	*50.7 ± 0.1*	*0.59 ± 0.1*	*1439.5 ± 12.1*
SA	*5.11 ± 2.45*	*85.3 ± 1.27*	*9.96 ± 7.07*	*2587.6 ± 54.04*
PW	*0.97 ± 0.11*	*101 ± 0.55*	*0.61 ± 0.09*	*1378.2 ± 4.73*

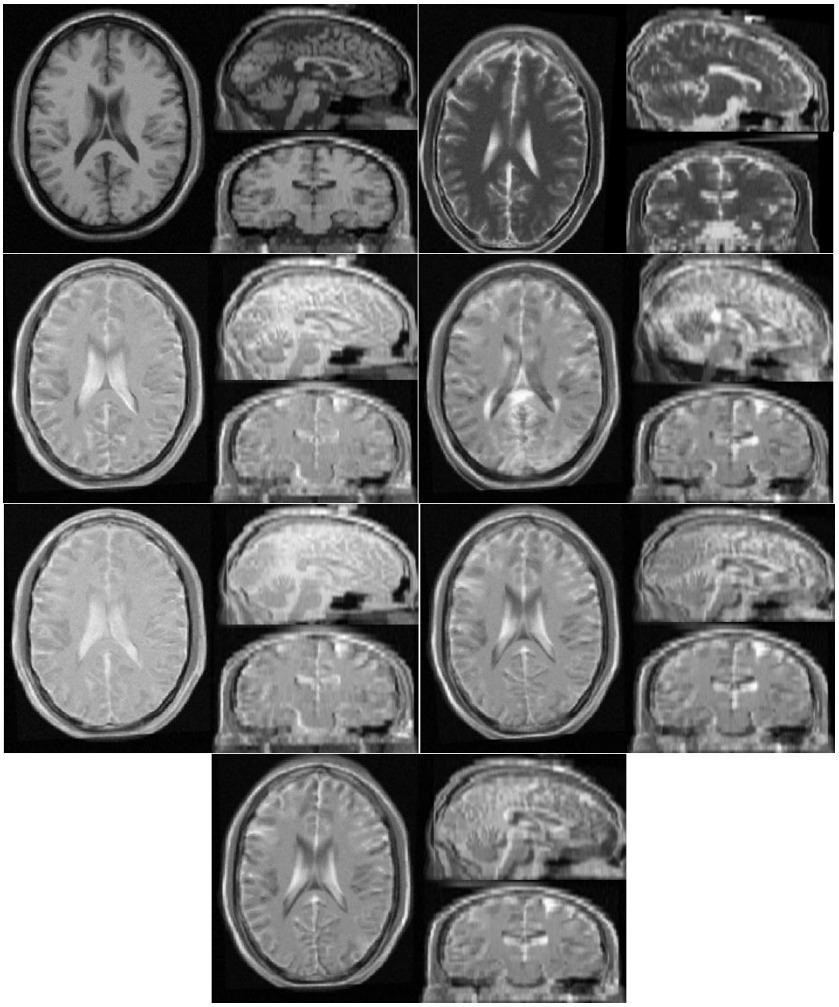

Figure 5.3 DS-3 registration results: First row, reference orthogonal T1 and float orthogonal T2 slices; second row superimposed registered slices from SPSA and NM algorithms; third row superimposed registered slices from HJ and SA algorithms; last row superimposed registered slices from PW.

Table 5.4 DS-3 registration results.

Optimization Algorithms	Phantom T2 to T1 Registration (5mm)			
	3-Slice Registration		3D Registration	
	Registration Error (vox) ± StD	Process Time (sec) ± StD	Registration Error (vox) ± StD	Process Time (sec) ± StD
SPSA	*1.26 ± 0.21*	*30.68 ± 0.21*	*2.01 ± 0.31*	*280.3 ± 1.29*
NM	*10.15 ± 2.11*	*120.25 ± 1.25*	*0.57 ± 0.12*	*868.3 ± 2.68*
HJ	*1.15 ± 0.13*	*41.27 ± 0.08*	*0.77 ± 0.06*	*920.3 ± 0.77*
SA	*6.29 ± 1.97*	*63.3 ± 0.66*	*6.24 ± 6.66*	*1435 ± 22.7*
PW	*7.41 ± 1.32*	*175 ± 1.12*	*0.63 ± 0.09*	*942.9 ± 0.45*

Widely used procedure for validation is to check against the algorithm with a gold standard, which its accuracy is known to be high. Gold standards mostly based on computer simulations (phantoms), cadavers or patient images. Advantage of phantoms is to provide true geometrical transformations between images. Although nothing can replace patient images for absolute validation as most acceptable standard, the difficulty of finding the real transformation between images makes phantom image based validations very important. Hence, to be able to make reliable quantitative assessments of the tested schemes, phantom image registration results are only given in detail in this paper, with which the accuracy and efficiency of the proposed method can be easily revealed by comparisons to those of the other schemes.

As seen in Table 5.1, Table 5.2 and Table 5.3, registration of 3D images with using entire volume intensities is very accurate for all optimization methods, but they are all undesirably time consuming. For a volume image which has 217x181x181 resolutions, there are 7.1 millions of voxels to be transformed in each transformation. While the optimization is in process, this volume image is transformed and similarity measure is calculated at least 2 times (may be more due to the optimization method) in every single iteration. Since the optimization processes spend more than %99 of the entire registration time based on these calculations, reducing the number of voxels to be transformed, clearly, will decrease the entire registration time.

The motivation of this study is to reduce the number of voxels involved in the iterative optimization process as much as possible without significant loss in registration accuracy. From this point of view, the proposed method provides for a reduction in the number of voxels to 0.11 million (111315 voxels) from 7.1 million with using only 3 central slices, of which two have 217×181 and one 181×181 resolutions, from a volume of $217 \times 181 \times 181$ resolution. 3-Slice approach does not transform entire volume to obtain these transformed three orthogonal slices. Since the central slices are only required to compute MI, it obtains them by calculating only corresponding voxels in these slices from volume image using the 4×4

transformation matrix regenerated from the transform parameter vector returned by the optimization algorithm.

As can be seen from Table 5.1, Table 5.2 and Table 5.3, the SPSA optimization algorithm provides the fastest results with both 3D entire volume approach and 3-Slice approach. In spite of being seen worse than NM, HJ and PW algorithms with 3D approach, its accuracy is still in acceptable range of clinical use which is identified by the most researchers as around or below 2 mm (Hajnal et al, 2001, Studholme et al., 1996, 1997, 1999; Thévenaz et al., 2000). However, with 3-Slice approach, SPSA becomes one of the most accurate and stable algorithm besides its high speed, and independence of varying resolutions while the NM and PW algorithms do not, which are very accurate and stable with only 3D entire volume aproach. Any increase in slice thickness causes lack of information in volume image, and accordingly introduces need for better interpolation techniques partly to compensate for the information loss. However, this is usually not preferred because of speed considerations, and it turns out that optimization algorithms fail in finding global maximum giving true transformation parameters. These interactions can be seen from the results in Table 5.3.

The SA optimization method gives the worst and unacceptable results in the way of clinical use for both approaches. The main reason is its sensitivity to the local minimas. In most cases, it stucks in local minimas/maximas and can not be able to exit. This the why being most time consuming method especially in 3D entire approach.

As a final discussion about the registration results obtained on the phantom images, 3-Slice approach provides substantial reductions in process time (10 to 40 times). In terms of accuracy and stability when the optimization algorithms are concerned, 3-Slice approach is considerably viable with SPSA and HJ algorithms but not with the rest of them.

5.2.2 PET - MR Registration Results

Positron emission tomography (PET) and magnetic resonance imaging (MR) are prominent examples of functional and morphological imaging modalities, respectively. The imaging concepts of PET and MR are quite different, as are the images. MR and PET are useful in conjuction precisely because they are complementary. While PET is able to provide information about a specific function such as cerebral blood flow or the density of a receptor in a certain area, MR defines different structures or tissue types, and thus provides information about morphology and the topology of structures. Both imaging modalities are useful for clinical studies as well as for basic scientific investigations. Once the respective images from PET and MR have been realigned and registered, they can be fused for an integrated display, offering the opportunity to describe with precision function and morphology at the same time. Because of this opportunity, PET – MR registrations are extensively used in clinical applications.

In this section, three groups of PET – MR registration results are given. The first group contains 7 different datasets of PET – MR-PD weighting volume pairs, the second group contains 7 datasets of PET – MR-T1 weighting volume pairs and the last group contains 6 datasets of PET – MR-T2 weighting volume pairs.

All the PET volume images have the same resolution, 128x128x15 voxels with 2.590723 mm pixel spacing and 8 mm slice thickness. There are two groups of resolution of MR volumes, both with 4 mm slice thickness and 1.25 mm pixel spacing. These are 256x256x26 and 256x256x24 voxels resolutions.

All the PET and MR volumes, which are given in the following 20 results, were resliced and rescaled to 256x256x52 or 256x256x48 voxels resolutions. The sample results are given in figures together with their relative registration error tables.

The results will be discussed after these figures and tables.

5.2.2.1 PET – MR-PD Registrations

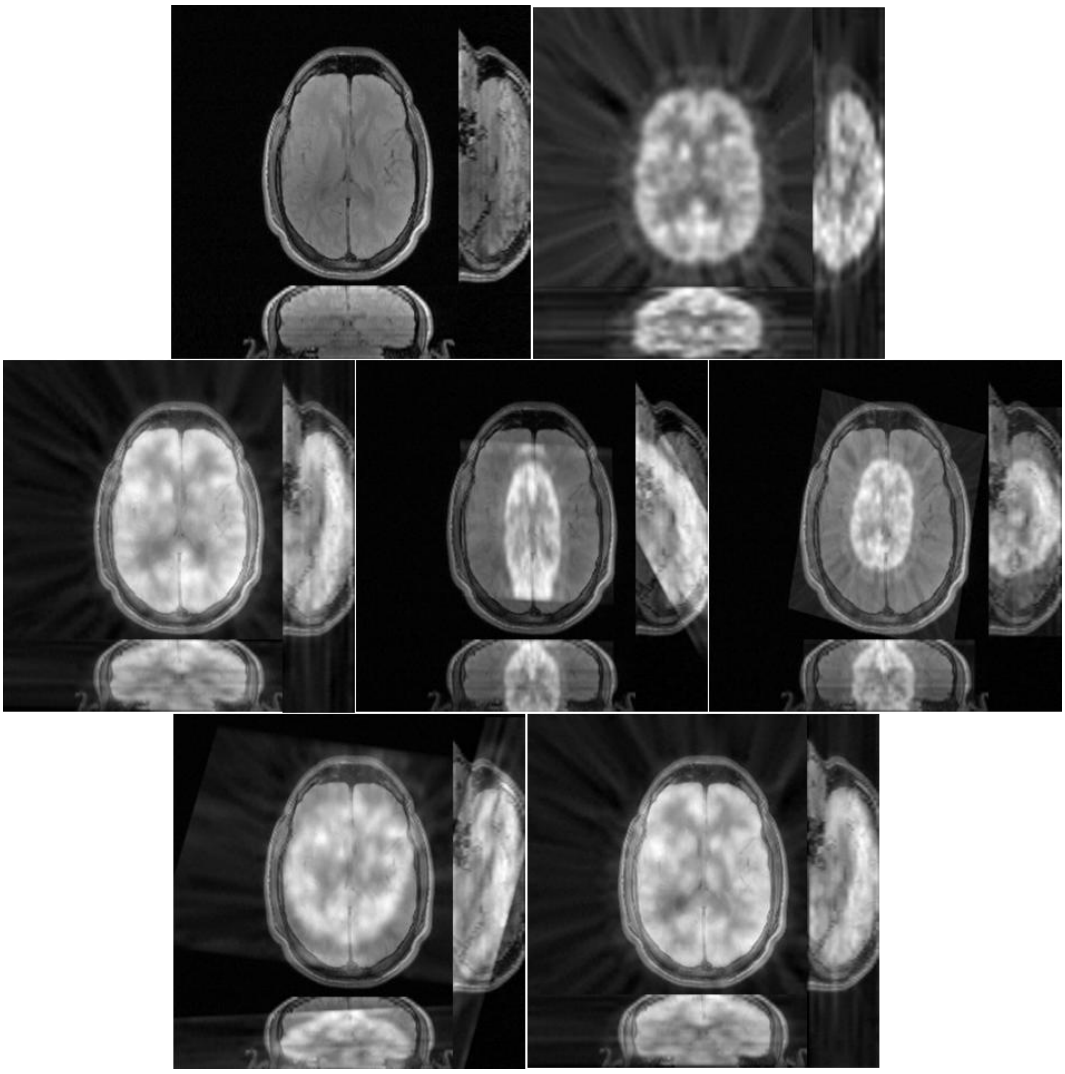

Figure 5.4 DS-4 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row superimposed registered slices from SPSA, NM and HJ algorithms; Last row superimposed registered slices from SA and PW algorithms.

Table 5.5 Relative registration errors for DS-4.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*66,98*	*63,12*	*60,67*	*5,73*
NM	*66,98*	-	*49,51*	*94,35*	*67,15*
HJ	*63,12*	*49,51*	-	*69,22*	*61,83*
SA	*60,67*	*94,35*	*69,22*	-	*60,70*
PW	*5,731*	*67,15*	*61,83*	*60,70*	-

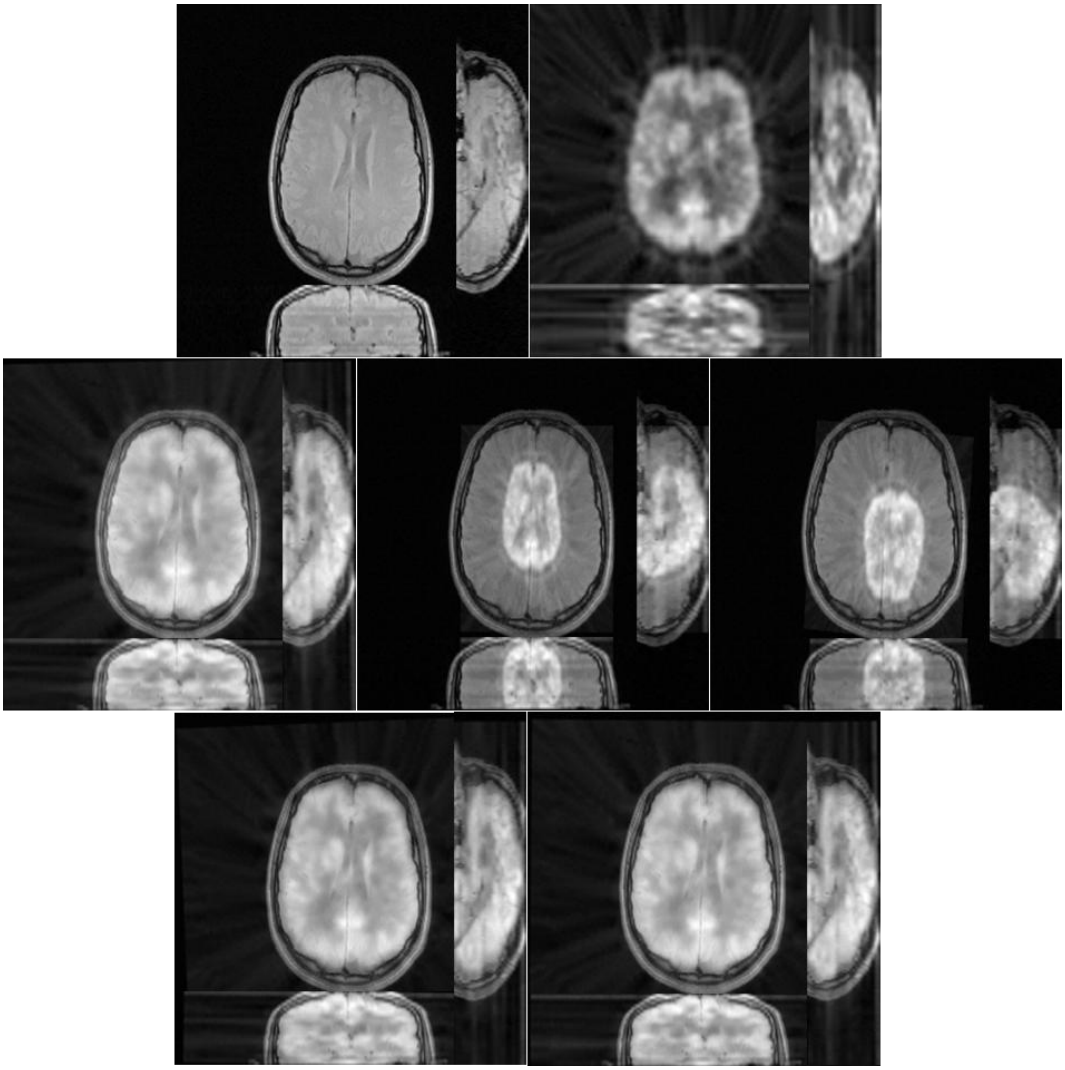

Figure 5.5 DS-5 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row superimposed registered slices from SPSA, NM and HJ algorithms; Last row superimposed registered slices from SA and PW algorithms

Table 5.6 Relative registration errors for DS-5.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*62,49*	*167,56*	*9,37*	*7,48*
NM	*62,49*	-	*108,63*	*53,76*	*56,29*
HJ	*167,56*	*108,63*	-	*159,27*	*161,73*
SA	*9,37*	*53,76*	*159,27*	-	*6,39*
PW	*7,48*	*56,29*	*161,73*	*6,39*	-

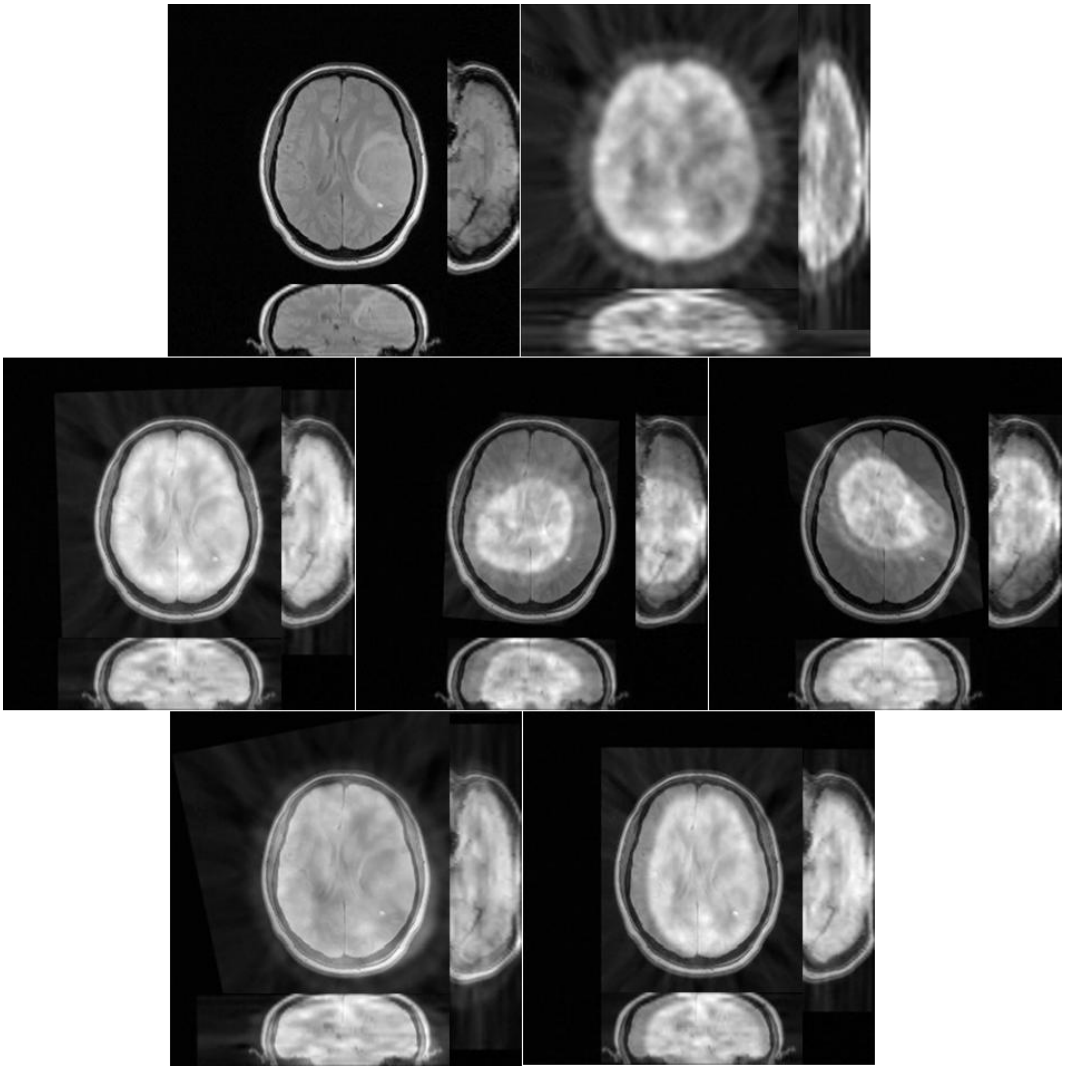

Figure 5.6 DS-6 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row superimposed registered slices from SPSA, NM and HJ algorithms; Last row superimposed registered slices from SA and PW algorithms

Table 5.7 Relative registration errors for DS-6.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	95,45	96,96	24,42	10,13
NM	95,45	-	114,15	102,66	91,77
HJ	96,96	114,15	-	120,09	90,58
SA	24,42	102,66	120,09	-	31,69
PW	10,13	91,77	90,58	31,69	-

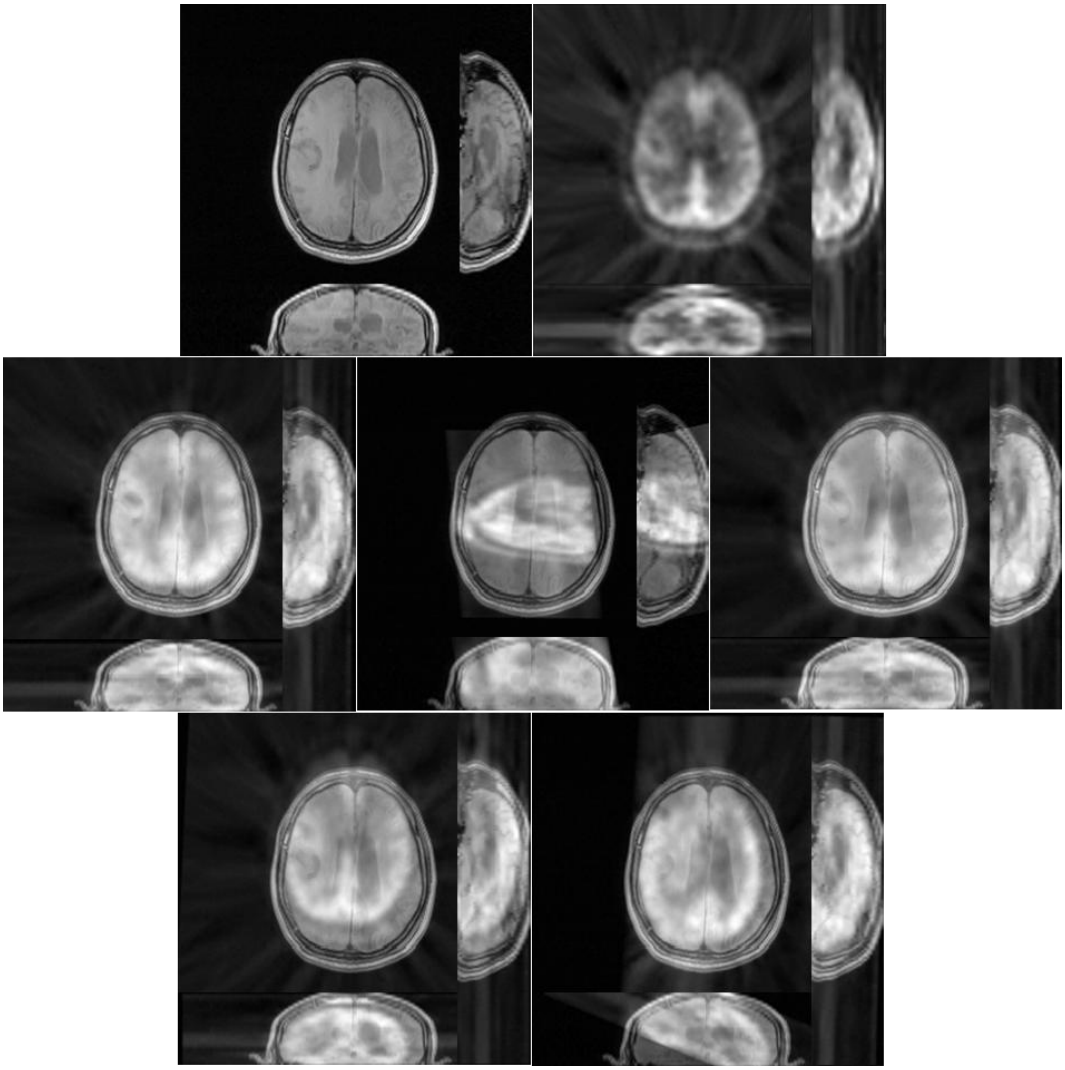

Figure 5.7 DS-7 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row superimposed registered slices from SPSA, NM and HJ algorithms; Last row superimposed registered slices from SA and PW algorithms

Table 5.8 Relative registration errors for DS-7.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*135,19*	*13,13*	*17,99*	*32,10*
NM	*135,19*	-	*145,11*	*130,47*	*143,47*
HJ	*13,13*	*145,11*	-	*24,00*	*39,04*
SA	*17,99*	*130,47*	*24,00*	-	*34,67*
PW	*32,10*	*143,47*	*39,04*	*34,67*	-

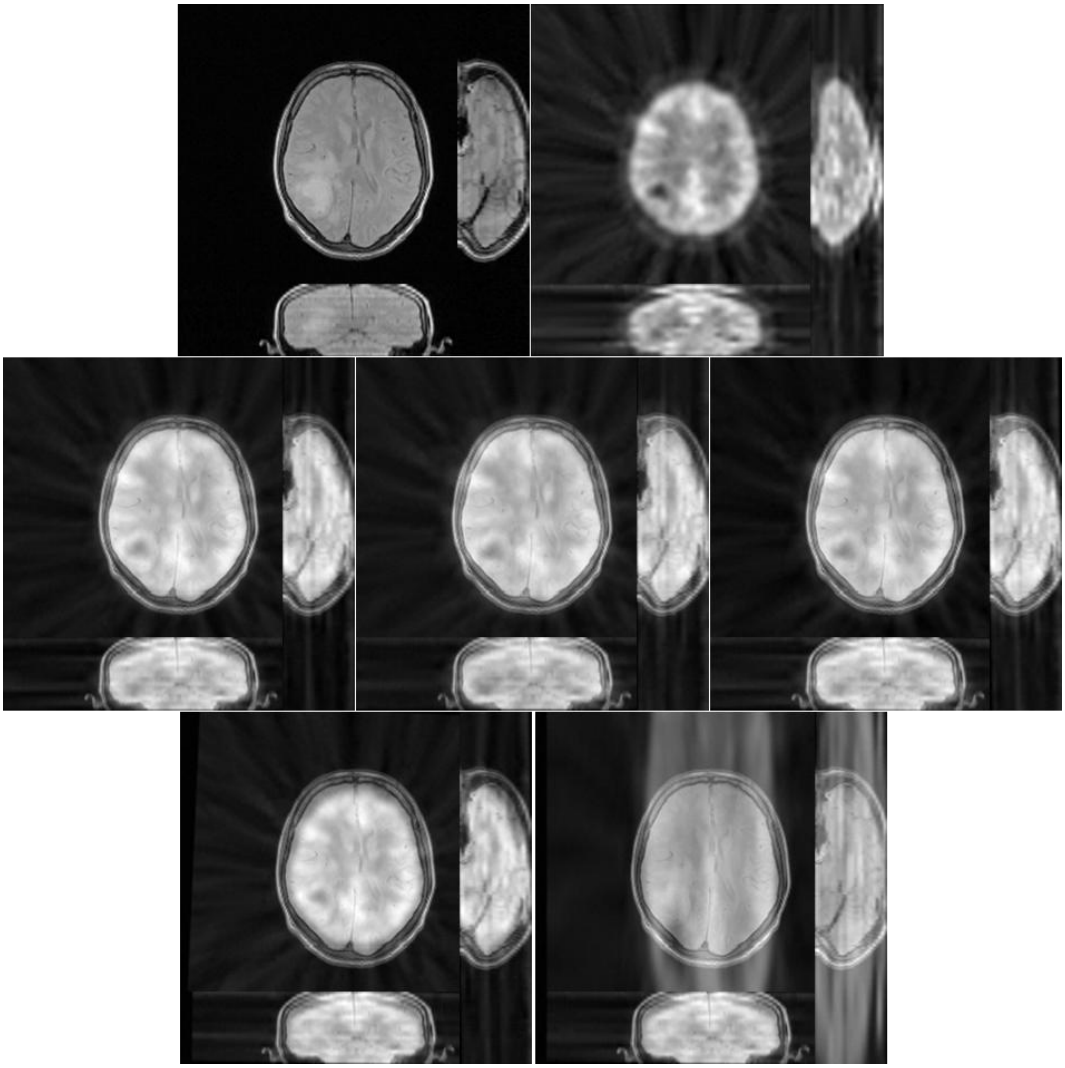

Figure 5.8 DS-8 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row superimposed registered slices from SPSA, NM and HJ algorithms; Last row superimposed registered slices from SA and PW algorithms

Table 5.9 Relative registration errors for DS-8.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*8,08*	*6,50*	*13,79*	*220,25*
NM	*8,08*	-	*6,83*	*21,28*	*218,48*
HJ	*6,50*	*6,83*	-	*19,27*	*215,04*
SA	*13,79*	*21,28*	*19,27*	-	*227,86*
PW	*220,25*	*218,48*	*215,04*	*227,86*	-

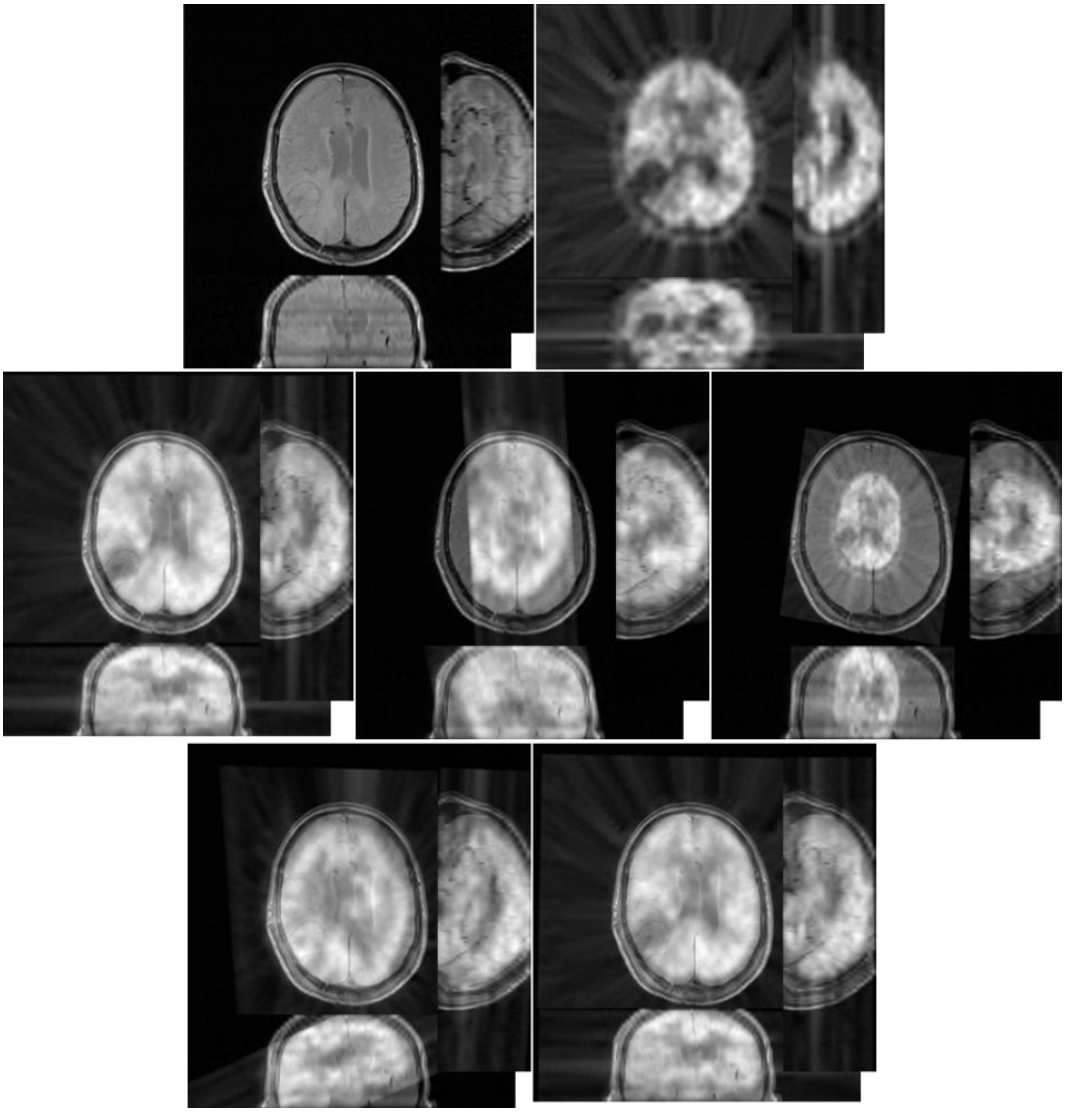

Figure 5.9 DS-9 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row superimposed registered slices from SPSA, NM and HJ algorithms; Last row superimposed registered slices from SA and PW algorithms

Table 5.10 Relative registration errors for DS-9.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*145,22*	*59,25*	*36,34*	*7,28*
NM	*145,22*	-	*115,83*	*151,50*	*140,64*
HJ	*59,25*	*115,83*	-	*58,86*	*52,18*
SA	*36,34*	*151,50*	*58,86*	-	*35,30*
PW	*7,28*	*140,64*	*52,18*	*35,30*	-

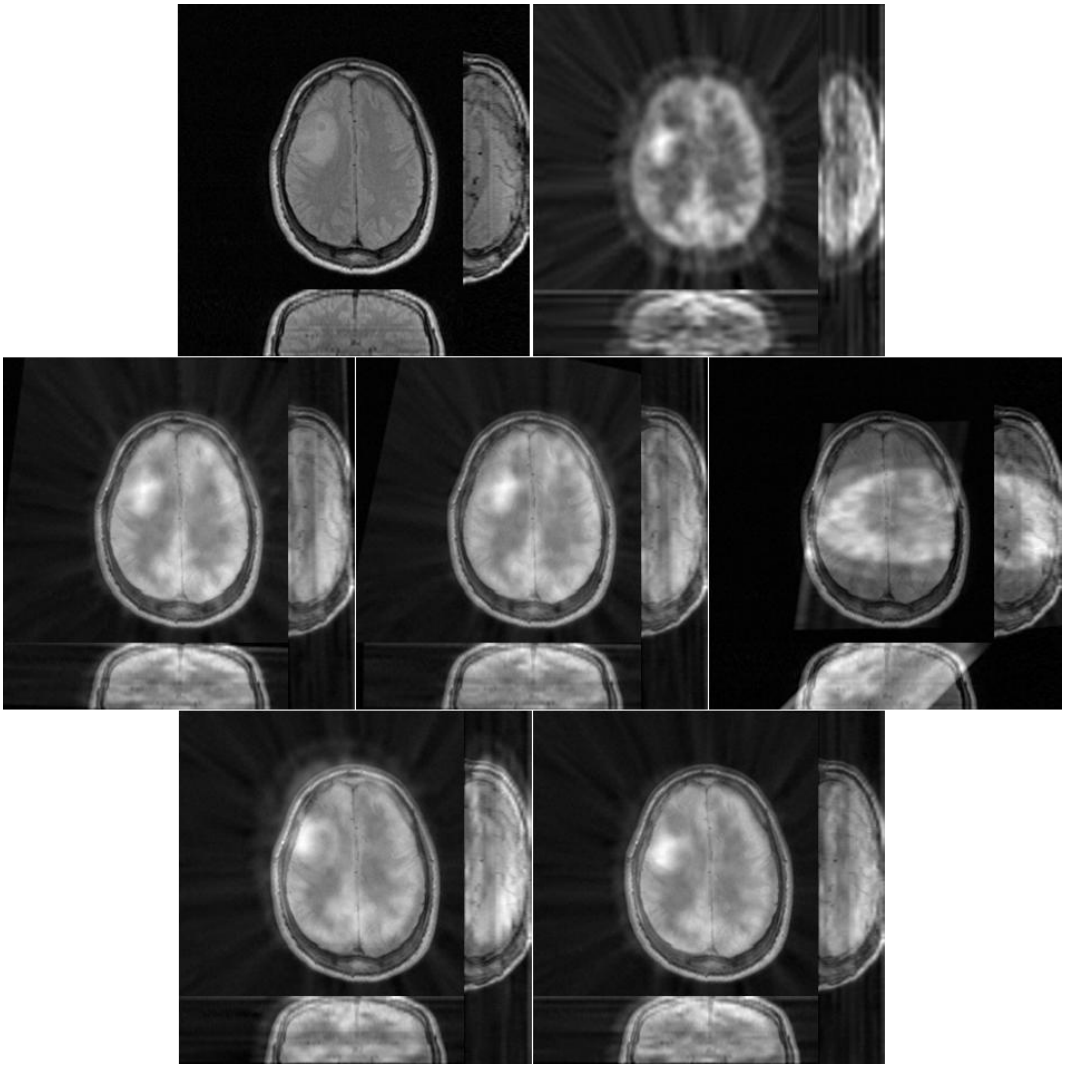

Figure 5.10 DS-10 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row superimposed registered slices from SPSA, NM and HJ algorithms; Last row superimposed registered slices from SA and PW algorithms

Table 5.11 Relative registration errors for DS-10.

	SPSA	**NM**	**HJ**	**SA**	**PW**
SPSA	-	10,33	85,66	27,11	13,95
NM	10,33	-	85,57	35,52	21,87
HJ	85,66	85,57	-	81,99	79,91
SA	27,11	35,52	81,99	-	17,26
PW	13,95	21,87	79,91	17,26	-

5.2.2.2 PET – MR-T1 Registrations

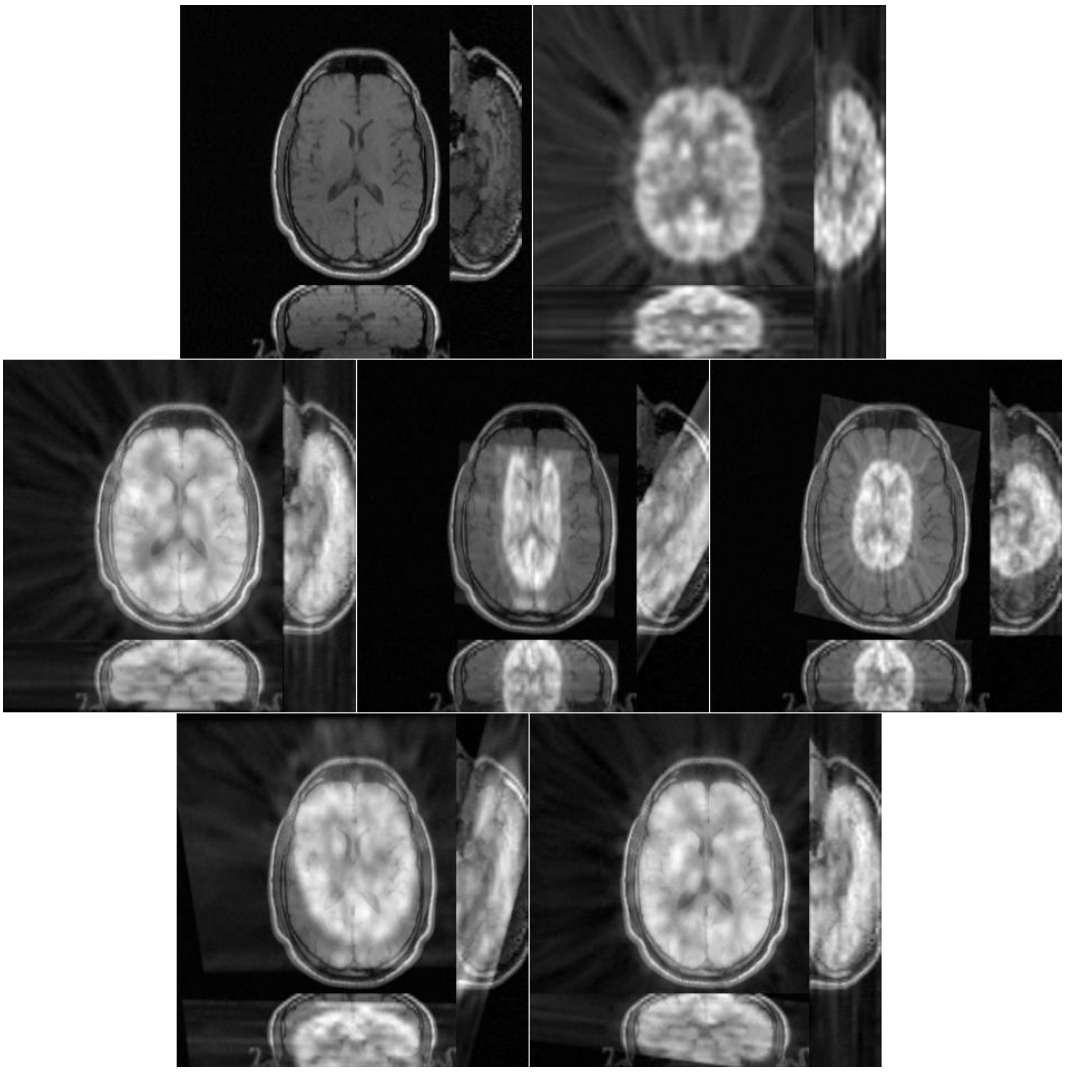

Figure 5.11 DS-11 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.12 Relative registration errors for DS-11

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*69,99*	*62,44*	*53,70*	*11,51*
NM	*69,99*	-	*59,15*	*63,13*	*68,29*
HJ	*62,44*	*59,15*	-	*84,33*	*60,70*
SA	*53,70*	*63,13*	*84,33*	-	*54,38*
PW	*11,51*	*68,29*	*60,70*	*54,38*	-

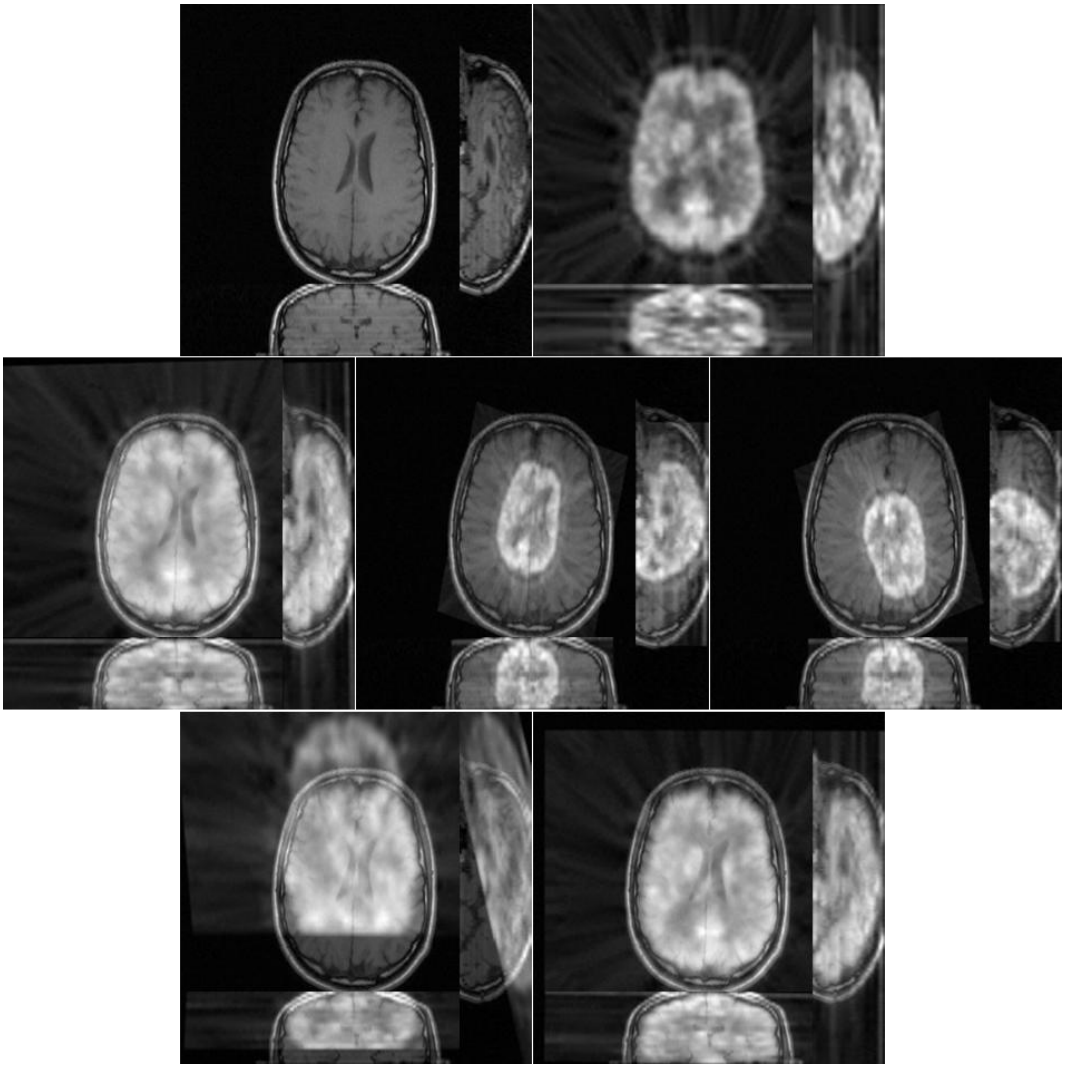

Figure 5.12 DS-12 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.13 Relative registration errors for DS-12.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*64,09*	*164,18*	*48,70*	*10,40*
NM	*64,09*	-	*104,66*	*85,69*	*55,64*
HJ	*164,18*	*104,66*	-	*183,16*	*155,62*
SA	*48,70*	*85,69*	*183,16*	-	*52,98*
PW	*10,40*	*55,64*	*155,62*	*52,98*	-

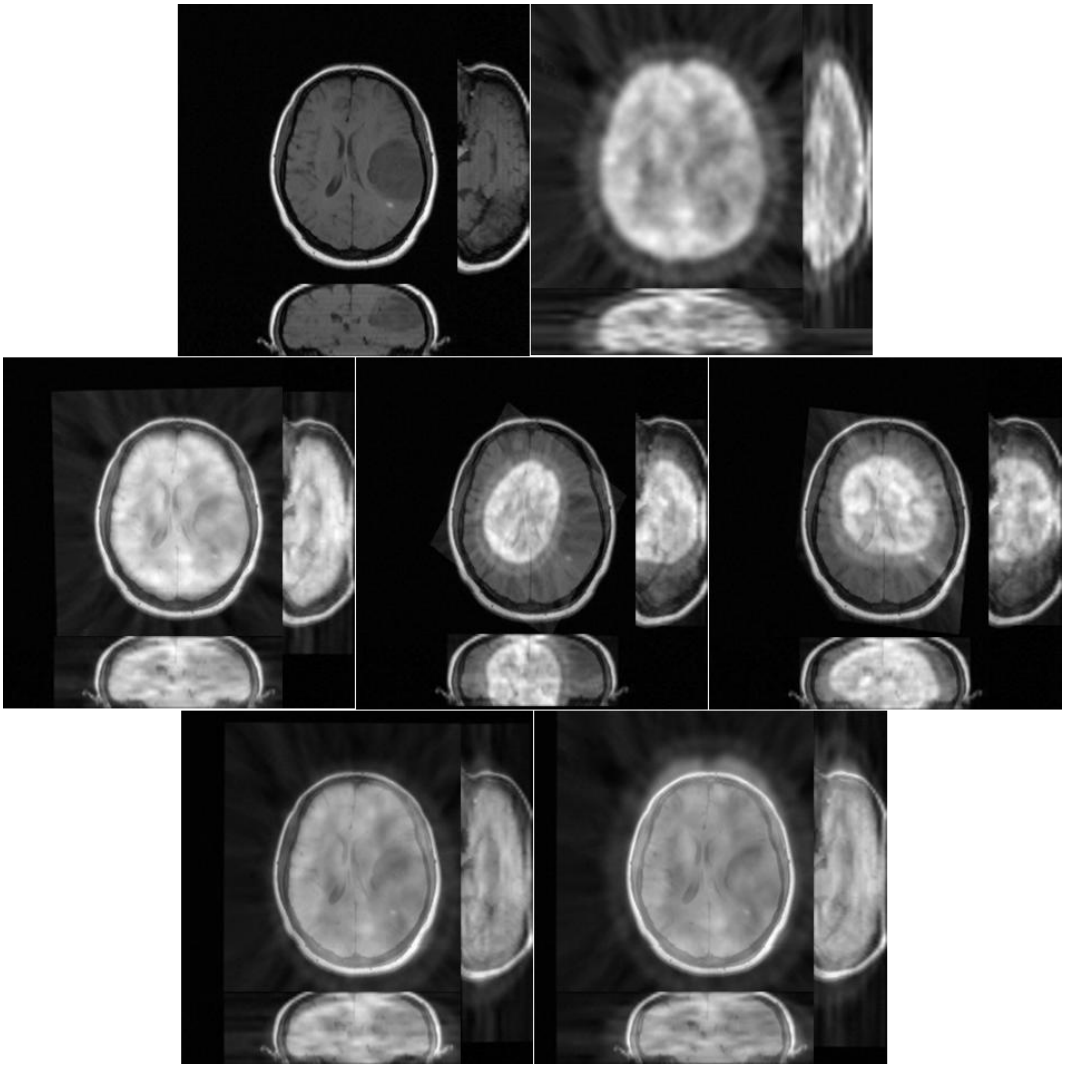

Figure 5.13 DS-13 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.14 Relative registration errors for DS-13.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*53,33*	*106,11*	*15,45*	*23,73*
NM	*53,33*	-	*55,59*	*64,64*	*72,90*
HJ	*106,11*	*55,59*	-	*116,85*	*125,28*
SA	*15,45*	*64,64*	*116,85*	-	*13,58*
PW	*23,73*	*72,90*	*125,28*	*13,58*	-

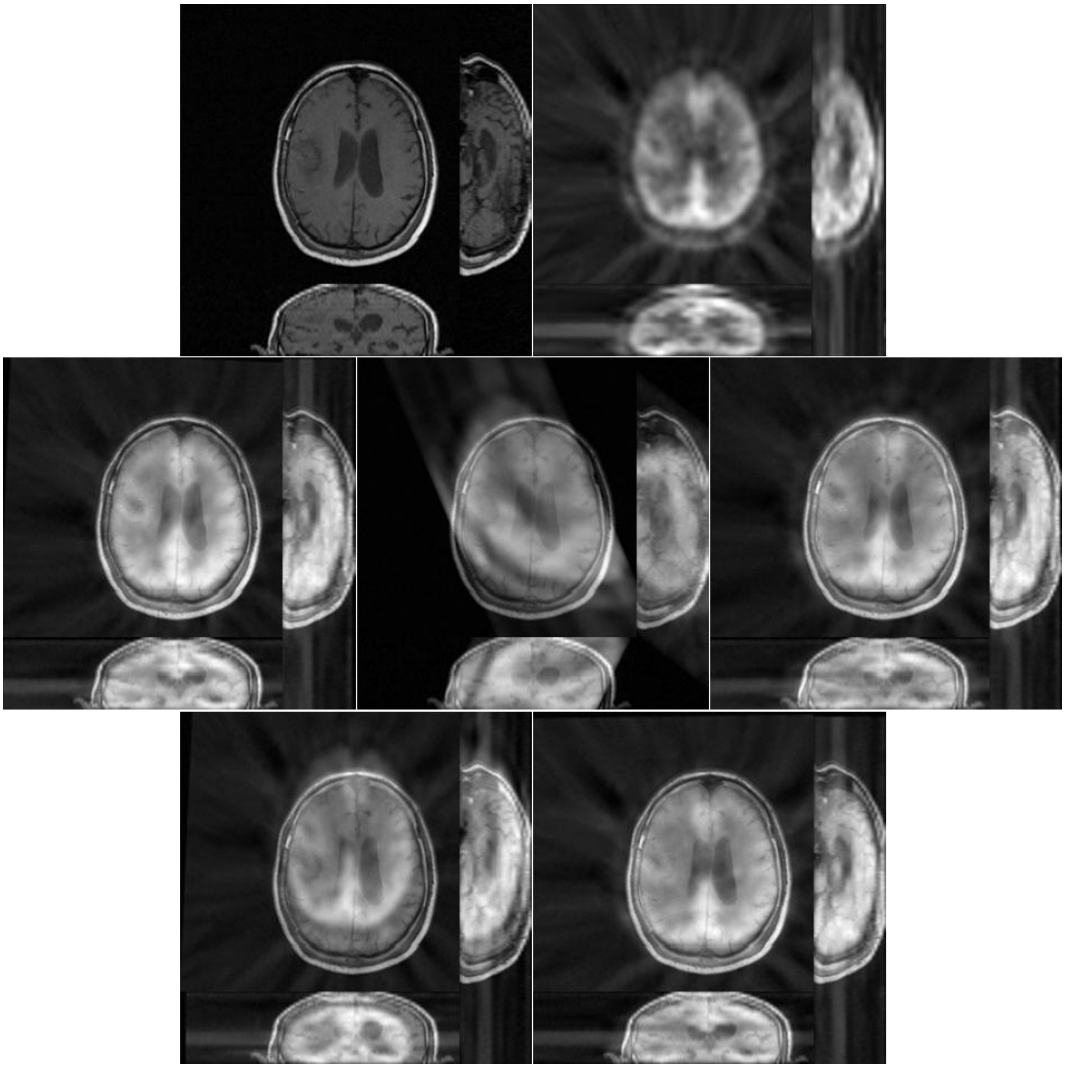

Figure 5.14 DS-14 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms.

Table 5.15 Relative registration errors for DS-14.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*106,61*	*16,32*	*19,98*	*10,76*
NM	*106,61*	-	*113,65*	*104,03*	*107,38*
HJ	*16,32*	*113,65*	-	*21,55*	*16,24*
SA	*19,98*	*104,03*	*21,55*	-	*23,79*
PW	*10,76*	*107,38*	*16,24*	*23,79*	-

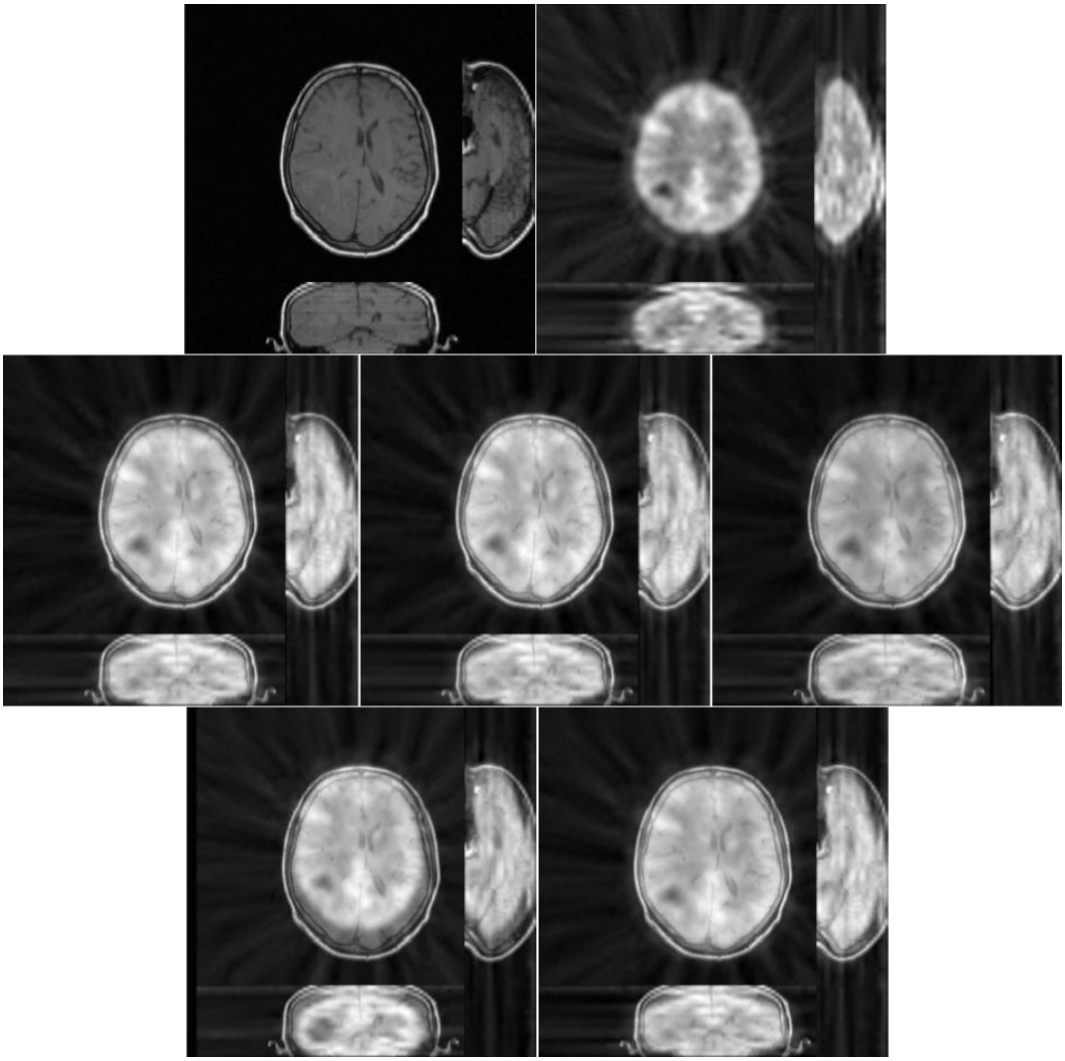

Figure 5.15 DS-15 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.16 Relative registration errors for DS-15.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*5,36*	*8,53*	*13,40*	*5,17*
NM	*5,36*	-	*9,73*	*13,88*	*4,51*
HJ	*8,53*	*9,73*	-	*19,32*	*8,61*
SA	*13,40*	*13,88*	*19,32*	-	*15,59*
PW	*5,17*	*4,51*	*8,61*	*15,59*	-

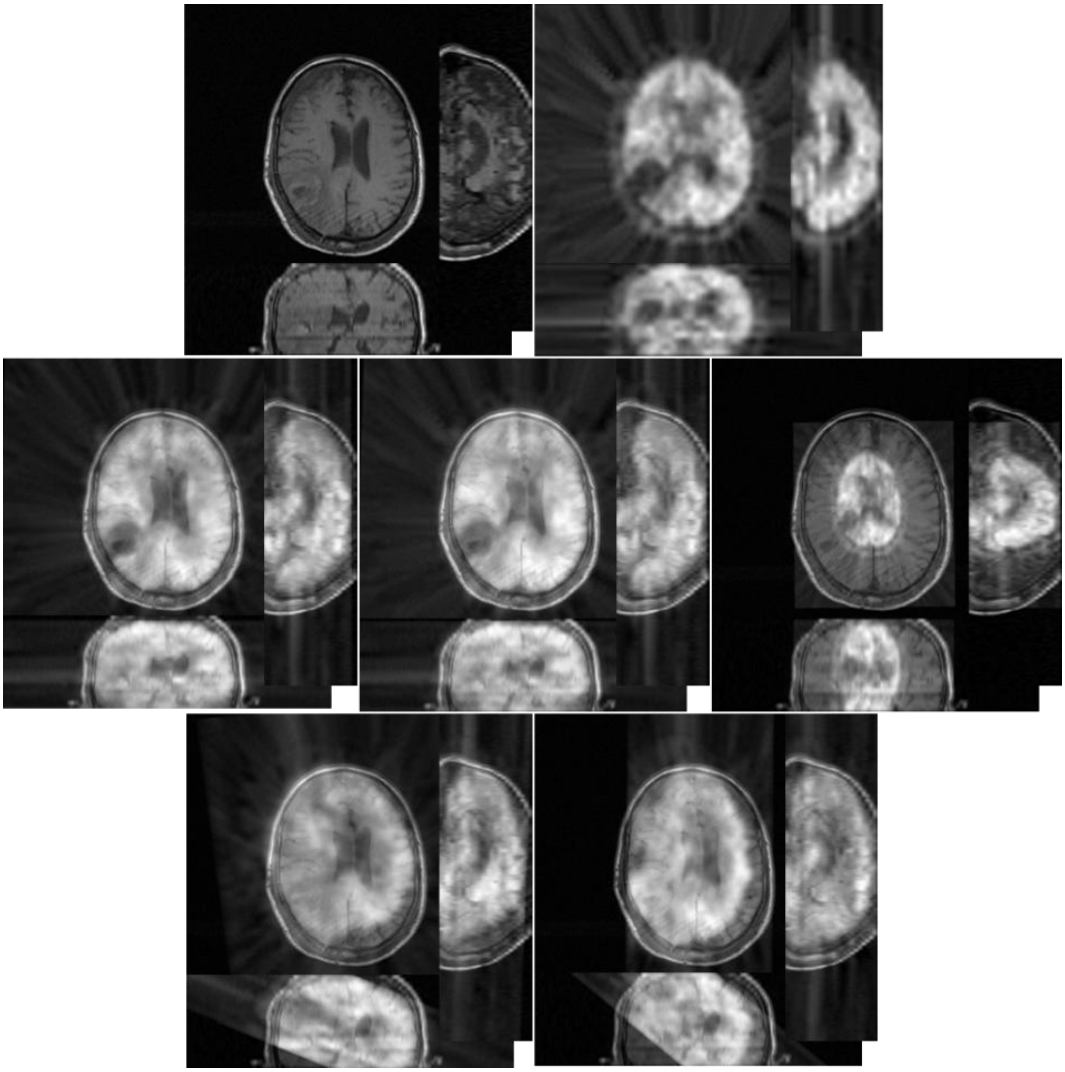

Figure 5.16 DS-16 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.17 Relative registration errors for DS-16.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*3,92*	*58,65*	*36,68*	*53,93*
NM	*3,92*	-	*58,74*	*37,45*	*53,41*
HJ	*58,65*	*58,74*	-	*63,53*	*65,47*
SA	*36,68*	*37,45*	*63,53*	-	*28,89*
PW	*53,93*	*53,41*	*65,47*	*28,89*	-

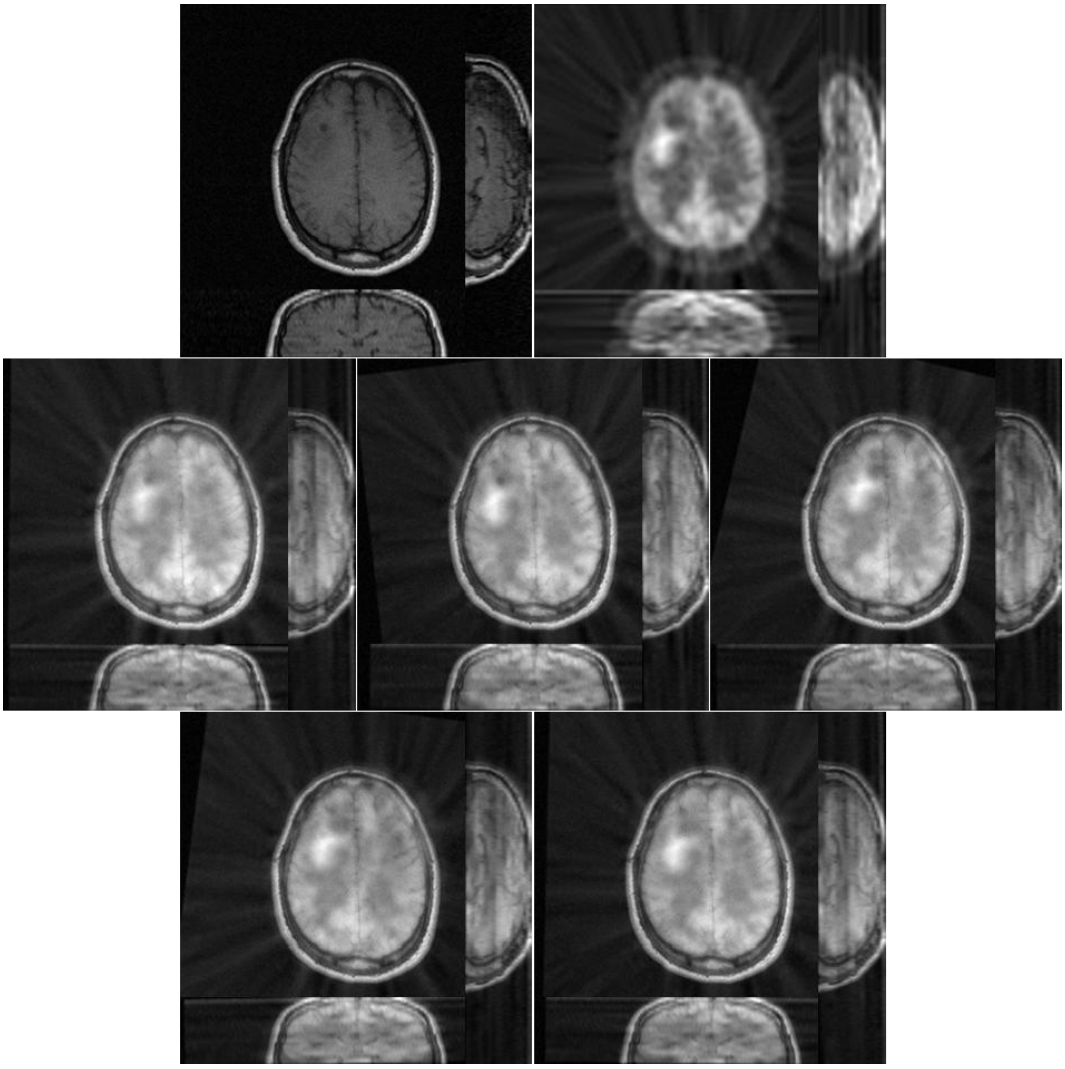

Figure 5.17 DS-17 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.18 Relative registration errors for DS-17.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*14,17*	*23,67*	*12,56*	*3,44*
NM	*14,17*	-	*36,52*	*25,89*	*14,82*
HJ	*23,67*	*36,52*	-	*11,80*	*22,04*
SA	*12,56*	*25,89*	*11,80*	-	*11,27*
PW	*3,44*	*14,82*	*22,04*	*11,27*	-

5.2.2.3 PET – MR-T2 Registrations

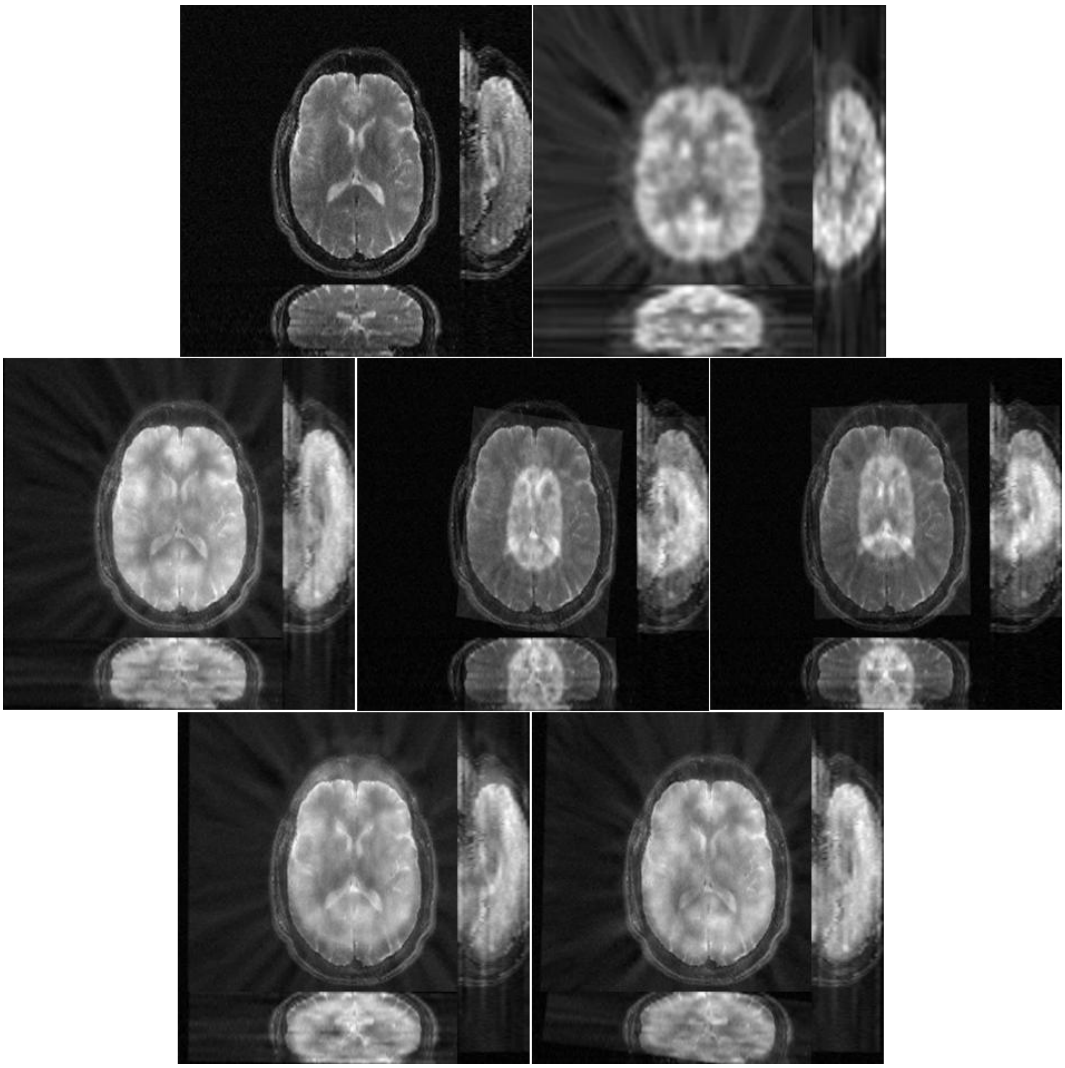

Figure 5.18 DS-18 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.19 Relative registration errors for DS-18.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	*-*	*58,42*	*57,59*	*12,90*	*10,65*
NM	*58,42*	*-*	*12,32*	*56,80*	*53,66*
HJ	*57,59*	*12,32*	*-*	*54,16*	*52,39*
SA	*12,90*	*56,80*	*54,16*	*-*	*13,99*
PW	*10,65*	*53,66*	*52,39*	*13,99*	*-*

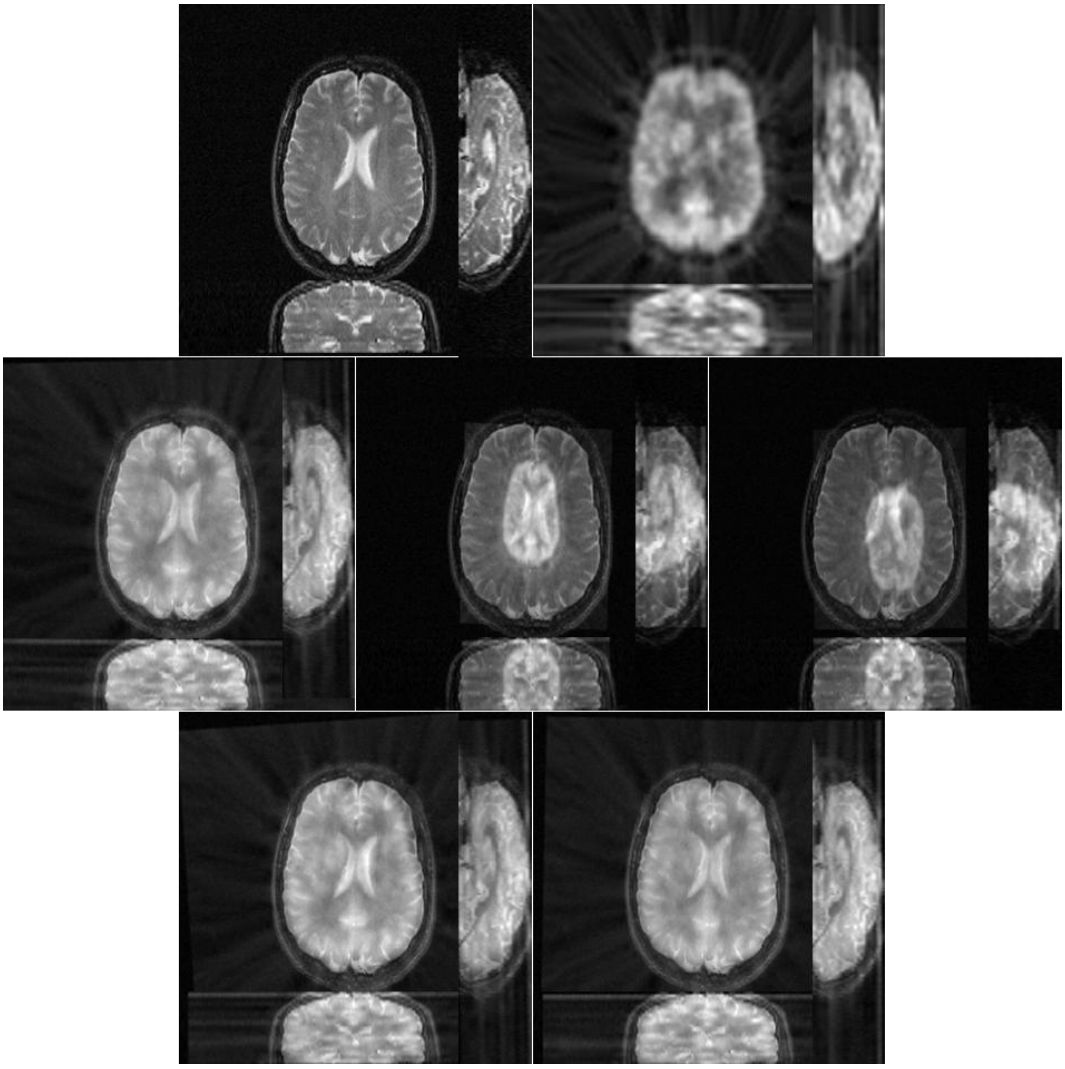

Figure 5.19 DS-19 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.20 Relative registration errors for DS-19.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*60,69*	*162,34*	*7,87*	*7,43*
NM	*60,69*	-	*104,56*	*54,77*	*56,11*
HJ	*162,34*	*104,56*	-	*156,39*	*158,00*
SA	*7,87*	*54,77*	*156,39*	-	*10,00*
PW	*7,43*	*56,11*	*158,00*	*10,00*	-

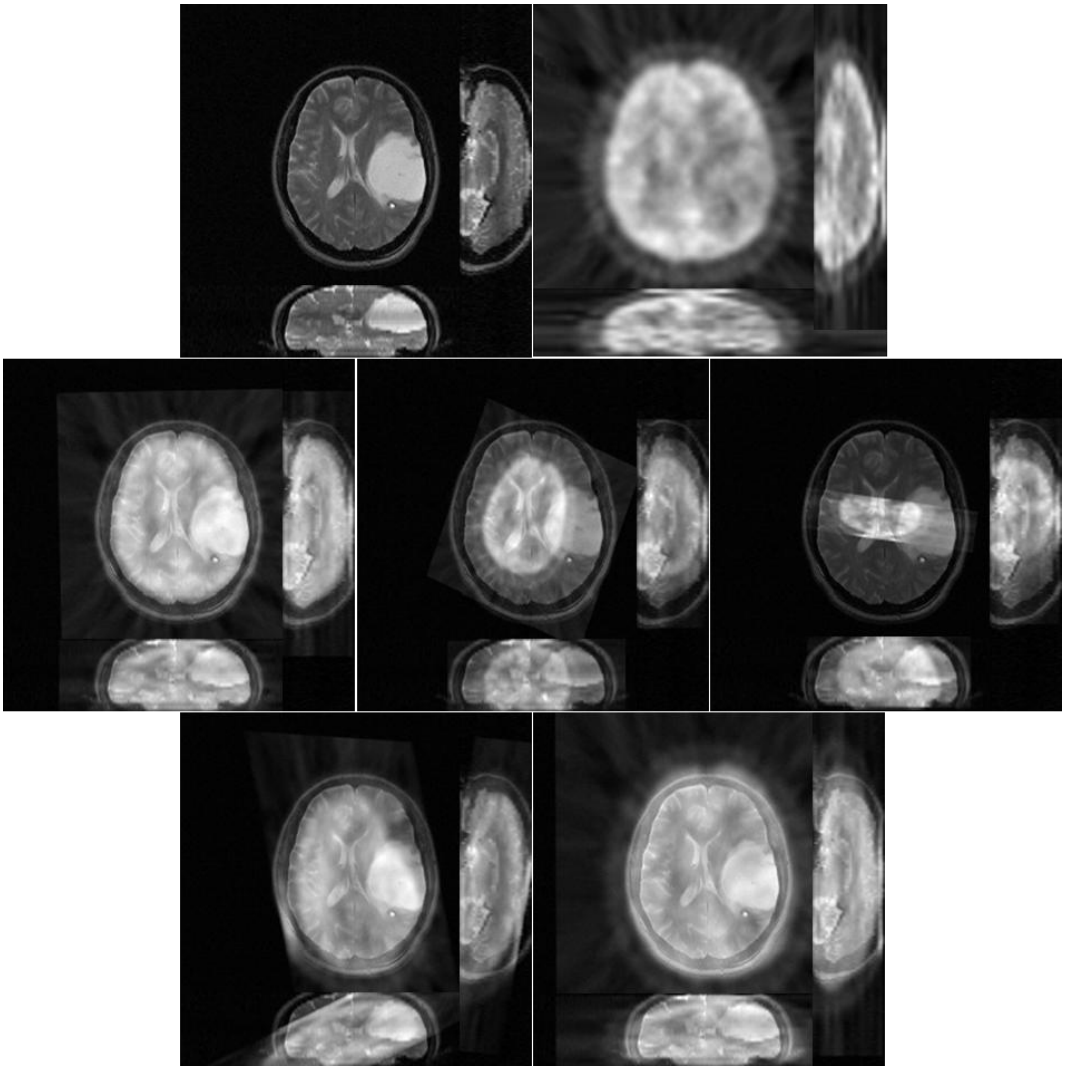

Figure 5.20 DS-20 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.21 Relative registration errors for DS-20.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*39,26*	*96,15*	*42,55*	*25,63*
NM	*39,26*	-	*82,74*	*56,78*	*61,48*
HJ	*96,15*	*82,74*	-	*110,45*	*112,98*
SA	*42,55*	*56,78*	*110,45*	-	*43,78*
PW	*25,63*	*61,48*	*112,98*	*43,78*	-

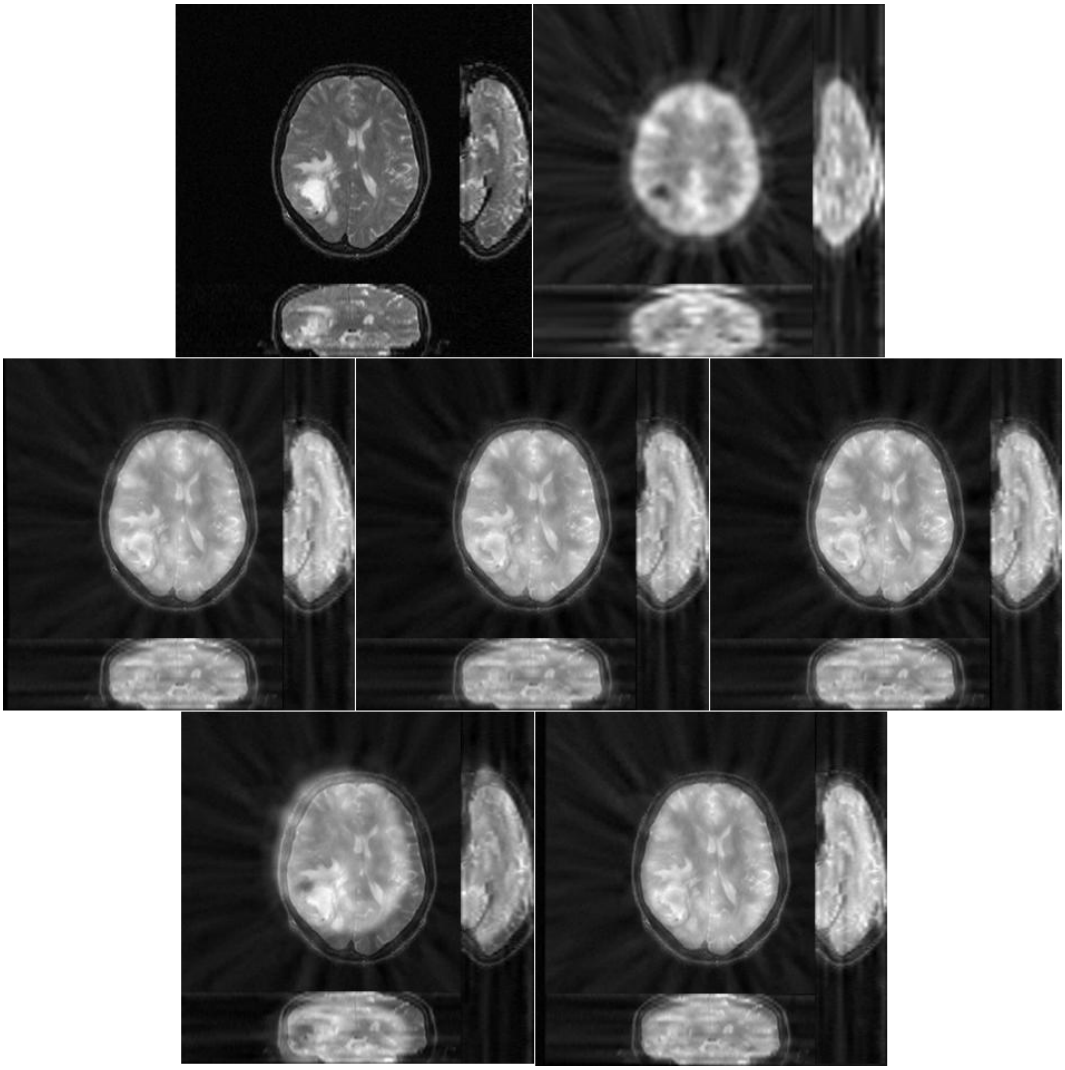

Figure 5.21 DS-21 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.22 Relative registration errors for DS-21.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	*-*	*5,36*	*4,48*	*14,01*	*3,99*
NM	*5,36*	*-*	*2,19*	*15,77*	*6,46*
HJ	*4,48*	*2,19*	*-*	*14,16*	*5,63*
SA	*14,01*	*15,77*	*14,16*	*-*	*14,63*
PW	*3,99*	*6,46*	*5,63*	*14,63*	*-*

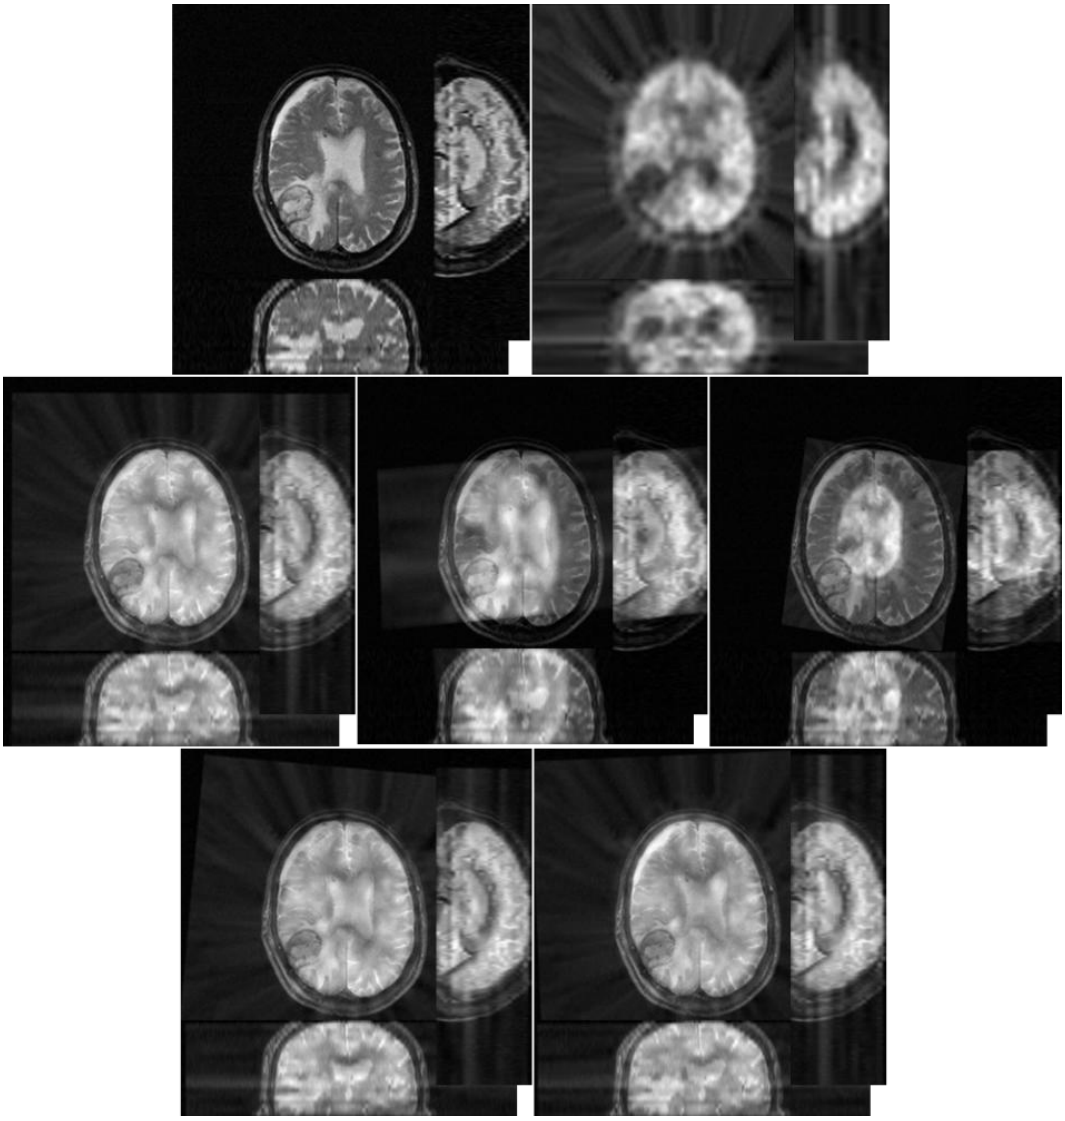

Figure 5.22 DS-22 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.23 Relative registration errors for DS-22.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*132,26*	*53,71*	*9,87*	*6,69*
NM	*132,26*	-	*109,83*	*133,17*	*134,57*
HJ	*53,71*	*109,83*	-	*53,21*	*56,58*
SA	*9,87*	*133,17*	*53,21*	-	*15,59*
PW	*6,69*	*134,57*	*56,58*	*15,59*	-

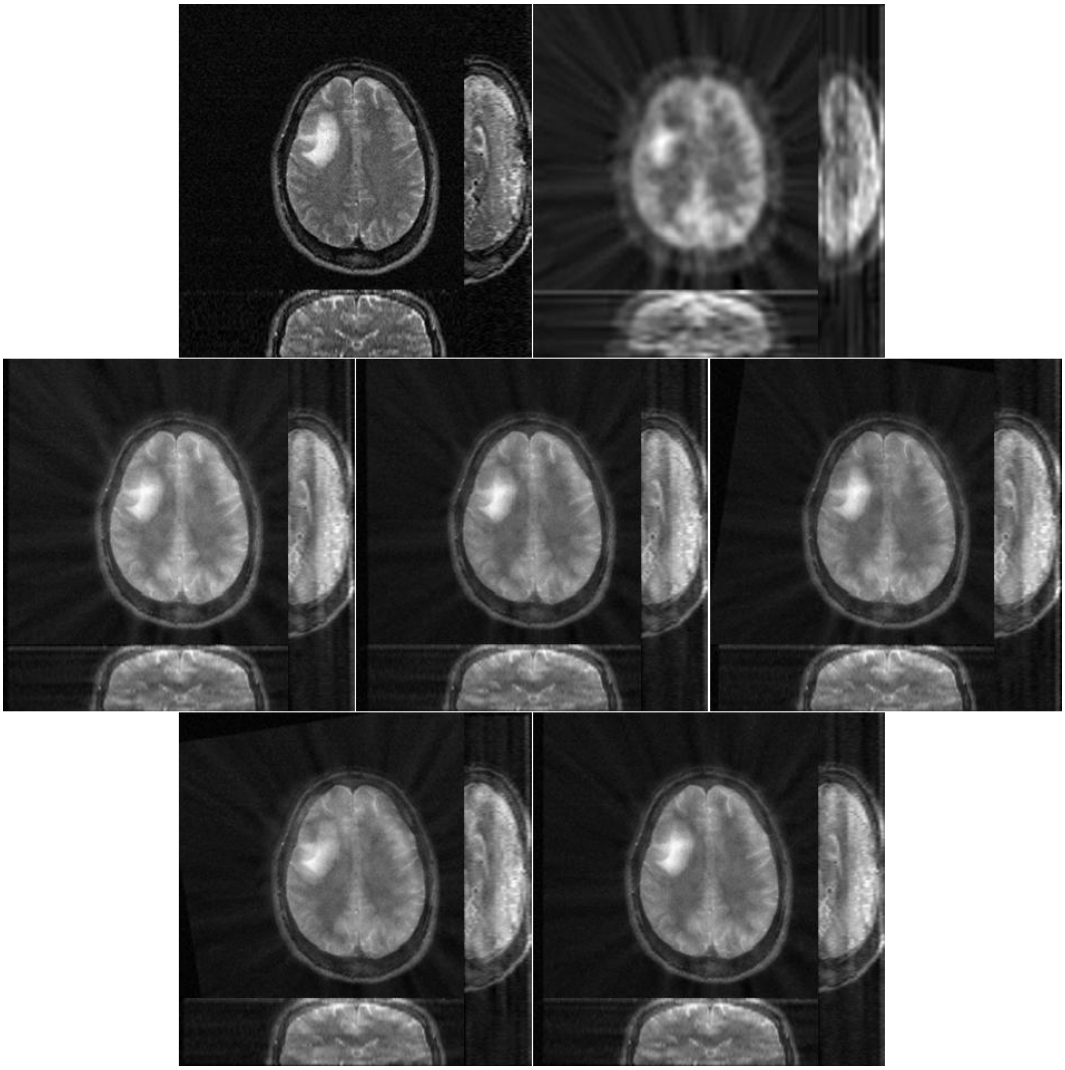

Figure 5.23 DS-23 registration results: First row, reference orthogonal MR and float orthogonal PET slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.24 Relative registration errors for DS-23.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*3,48*	*15,57*	*18,04*	*3,05*
NM	*3,48*	-	*16,62*	*16,08*	*1,45*
HJ	*15,57*	*16,62*	-	*32,21*	*15,58*
SA	*18,04*	*16,08*	*32,21*	-	*17,06*
PW	*3,05*	*1,45*	*15,58*	*17,06*	-

As given by the PET – MR results, from Figure 5.4 to Figure 5.23, the first noticeable difference from phantom image results is the failure of deterministic optimization algorithms which are NM, HJ and PW.

Since PET volume images have very low resolutions (low information), the optimization algorithms mostly have major difficulties in convergence to the true registration parameters, except the SPSA. Because of its advantage of avoiding the local minimas, unlike the deterministic gradient based methods which have tendency of getting trapped in local minimas, the SPSA seems to converge without any difficulty to the optimal registration parameters.

The term of "optimal registration parameter" is identified by our visual judgement of these PET – MR registration results because of absence of true registration parameters. Hence, to define registration error in deterministic way, the relative registration error tables are given which show root mean square errors between any optimization method and the others. This approximation provides to compare any optimization method results with any others numerically.

As can be seen from the results, the 3-Slice method with the SPSA algorithm gives the most accurate (mostly only the one) results. Beside its accuracy and speed (Table 5.25), there was no significant failure of the method in the given 20 results obtained with different datasets. Therefore the SPSA deserves the description of the "stable and robust algorithm".

Table 5.25 PET – MR Registrations' average process time with 3-Slice method (Sec ± StD).

	PET - MR-PD	***PET - MR-T1***	***PET - MR-T2***
SPSA	*52.4 ± 1.18*	*52.6 ± 1.34*	*52.1 ± 1.1*
NM	*101.2 ± 16.2*	*107.3 ± 15.7*	*96.2 ± 13.1*
HJ	*63.61 ± 11.5*	*64.7 ± 9.1*	*64.1 ± 9.5*
SA	*51.3 ± 6.1*	*51.6 ± 4.2*	*49.6 ± 5.75*
PW	*198.5 ± 84.1*	*259.3 ± 126*	*345 ± 106.1*

In Table 5.26, the relative errors can be seen for all 3-Slice PET – MR registrations. These values were calculated by avareging out the results of 20 consecutive runs of the registration method on each dataset type of PET –MR.

Table 5.26 PET – MR Registrations' average relative errors

		SPSA	***NM***	***HJ***	***SA***	***PW***
PET MR-PD	***SPSA***	-	*74.81 ± 50.5*	*70.3 ± 50.6*	*27.1 ± 16.01*	*42.4 ± 73.07*
	NM	*74.81±50.5*	-	*89.4 ± 43.4*	*84.2 ± 45.3*	*105.6 ± 61.5*
	HJ	*70. 3 ± 50.5*	*89.4 ± 43.4*	-	*76.1 ± 46.5*	*100.04 ± 59.7*
	SA	*27.1±16.01*	*84.2±45.3*	*76.1±46.5*	-	*59.1±70.6*
	PW	*42.4±73.07*	*105.6±61.5*	*100.04±59.7*	*59.1±70.6*	-
		SPSA	***NM***	***HJ***	***SA***	***PW***
PET MR-T1	***SPSA***	-	*45.3±35.9*	*62.8±51.8*	*28.6±16.1*	*16.9±16.2*
	NM	*45.3±35.9*	-	*62.5±33.6*	*56.3±30.1*	*53.8±32.5*
	HJ	*62.8±51.8*	*62.5±*	-	*71.5±57.9*	*64.5±52.4*
	SA	*28.6±16.1*	*56.3±30.1*	*71.5±57.9*	-	*28.6±16.8*
	PW	*16.9±16.2*	*53.8±32.5*	*64.5±52.4*	*28.6±16.8*	-
		SPSA	***NM***	***HJ***	***SA***	***PW***
PET MR-T2	***SPSA***	-	*49.9±43.2*	*64.9±52.8*	*17.5±11.6*	*9.5±7.5*
	NM	*49.9±43.2*	-	*54.7±45.3*	*55.5±39*	*52.2±43.8*
	HJ	*64.9±52.8*	*54.7±45.3*	-	*70.1±52.6*	*66.8±58.1*
	SA	*17.5±11.6*	*55.5±39*	*70.1±52.6*	-	*19.1±36.8*
	PW	*9.5±7.5*	*52.2±43.8*	*66.8±58.1*	*19.1±36.8*	-

5.2.3 CT – MR Registration Results

MR and CT scanners both produce images of anatomical structure, but the images they generate look quite different. One of the most striking differences is that cortical bone has high x-ray attenuation, so is bright in CT, whereas cortical bone contains virtually no MR visible protons, so is black in MR. This means that the information they provide is complemantary. Because of this complemantary nature of MR and CT, patients are frequently imaged with both and CT-MR registration is widely used for the purpose of diagnosis, or during the workup for treatment.

In this section, three groups of CT – MR registration results are given. The first group contains 3 CT – MR-PD volume pairs, the second group contains 5 CT - MR-T1 weighting volume pairs and the last group contains 4 PET - MR-T2 weighting volume pairs. All the CT and MR volumes, which are given in the following 12 results, are resliced and rescaled to 256x256x52 voxels resolution. The sample results are given in figures together with their relative registration error tables.

The results will be discussed after these figures and tables.

5.2.3.1 CT – MR-PD Registrations

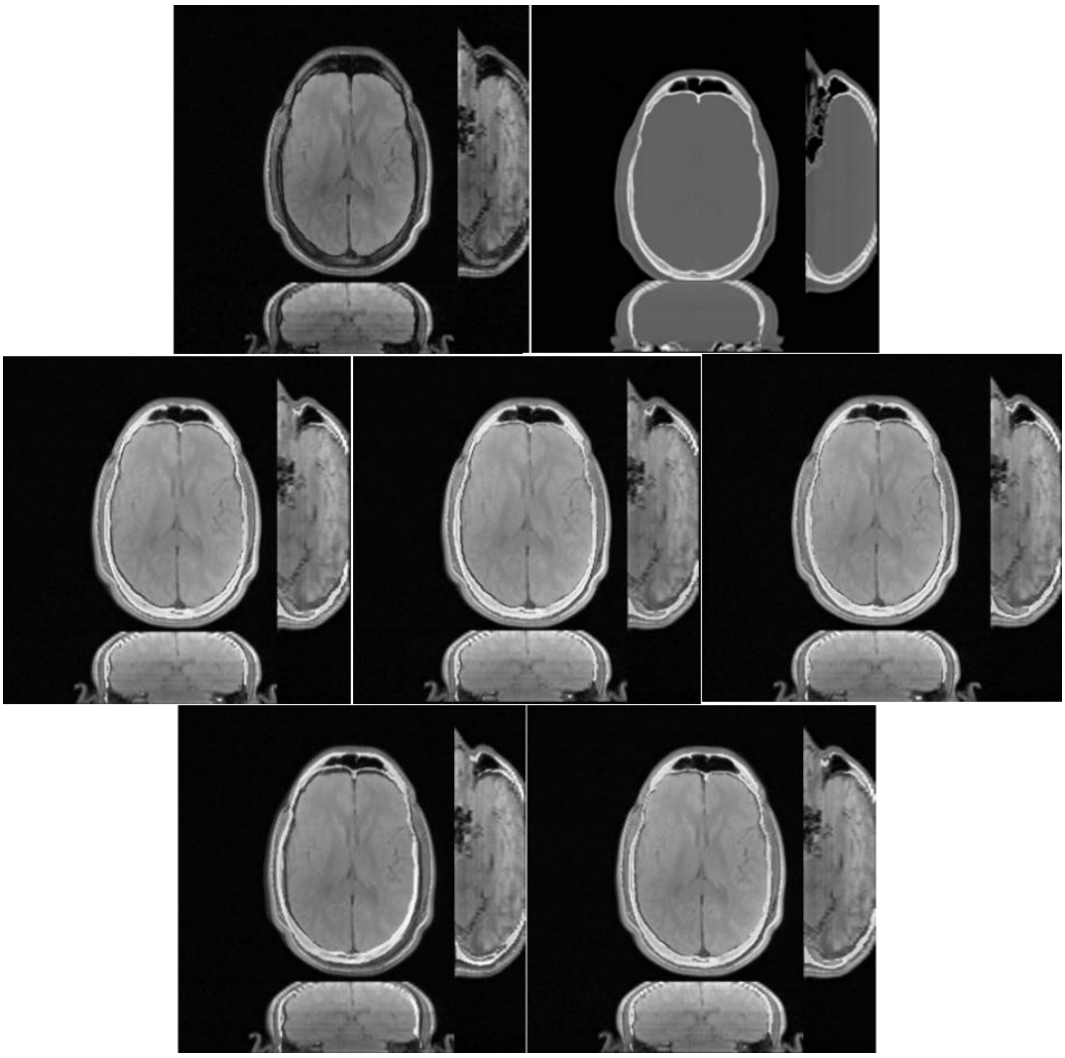

Figure 5.24 DS-24 registration results: First row, reference orthogonal MR and float orthogonal CT slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.27 Relative registration errors for DS-24.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*3,48*	*15,57*	*18,05*	*3,05*
NM	*3,48*	-	*16,62*	*16,08*	*1,45*
HJ	*15,57*	*16,62*	-	*32,21*	*15,58*
SA	*18,05*	*16,08*	*32,21*	-	*17,06*
PW	*3,05*	*1,45*	*15,58*	*17,06*	-

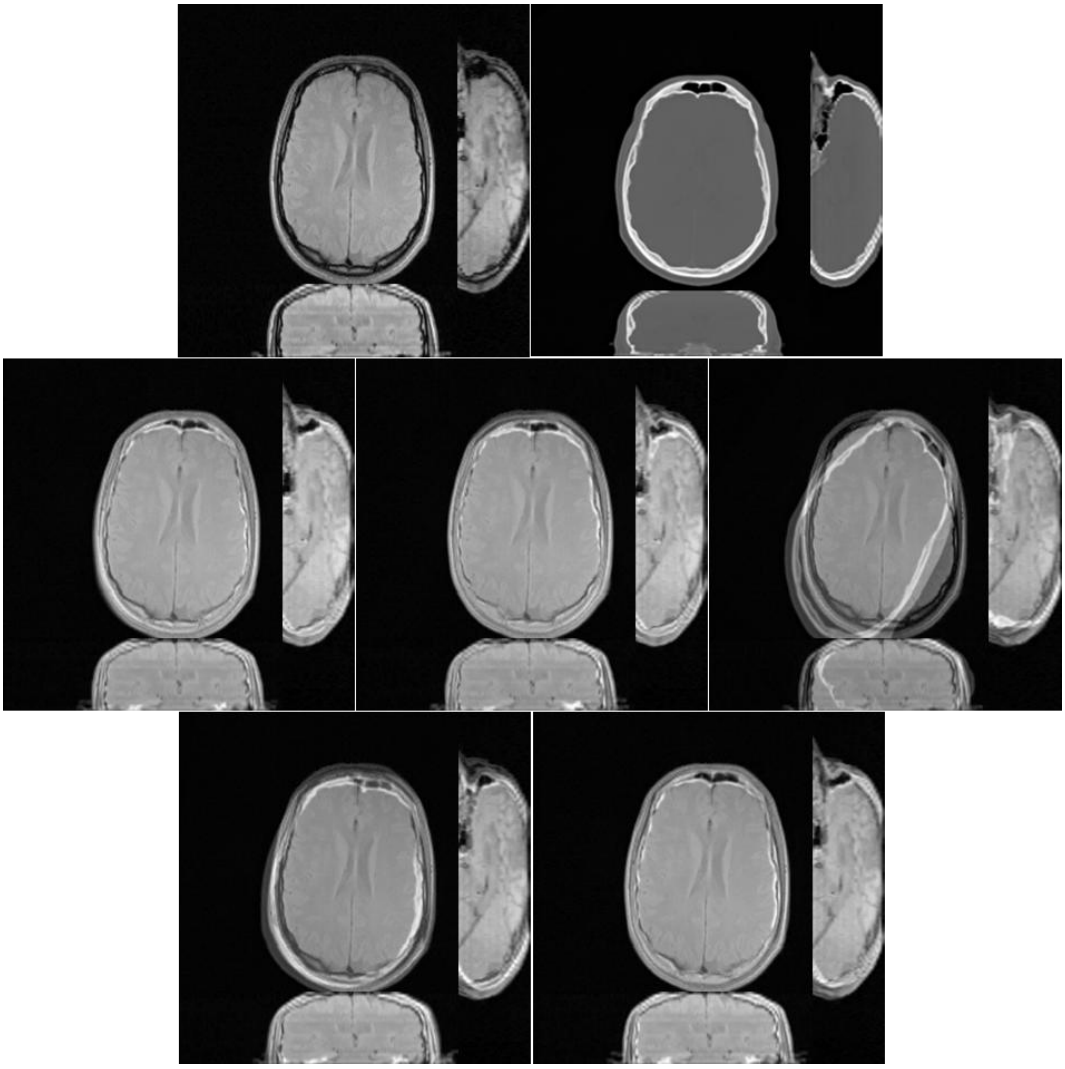

Figure 5.25 DS-25 registration results: First row, reference orthogonal MR and float orthogonal CT slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.28 Relative registration errors for DS-25.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*5,06*	*77,23*	*12,85*	*2,32*
NM	*5,06*	-	*80,00*	*16,12*	*5,42*
HJ	*77,23*	*80,00*	-	*66,29*	*76,45*
SA	*12,85*	*16,12*	*66,29*	-	*12,84*
PW	*2,32*	*5,42*	*76,45*	*12,84*	-

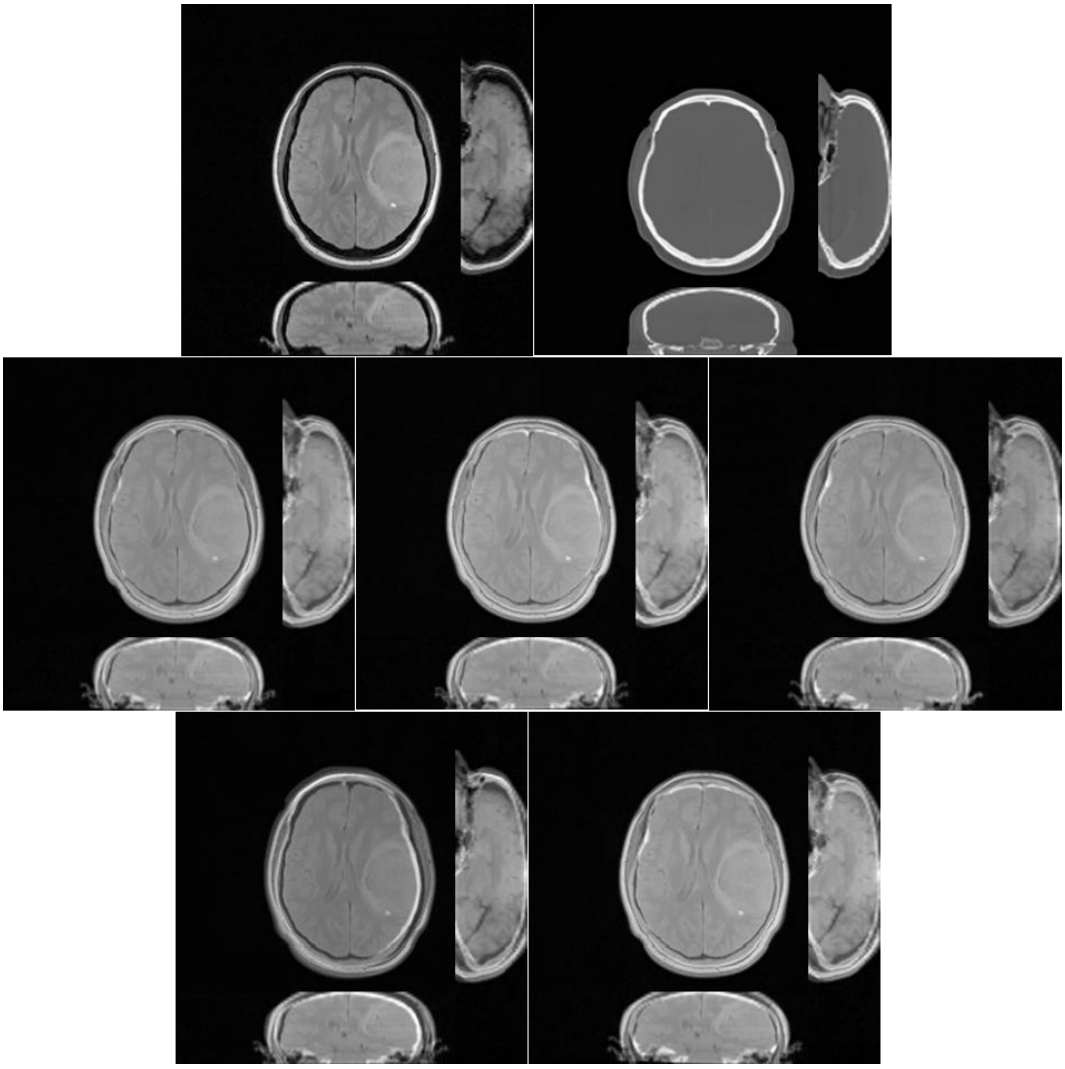

Figure 5.26 DS-26 registration results: First row, reference orthogonal MR and float orthogonal CT slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.29 Relative registration errors for DS-26.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*5,52*	*5,71*	*5,71*	*4,31*
NM	*5,52*	-	*7,82*	*7,82*	*4,68*
HJ	*5,71*	*7,82*	-	-	*3,51*
SA	*5,71*	*7,82*	-	-	*3,51*
PW	*4,31*	*4,68*	*3,51*	*3,51*	-

5.2.3.2 CT – MR-T1 Registrations

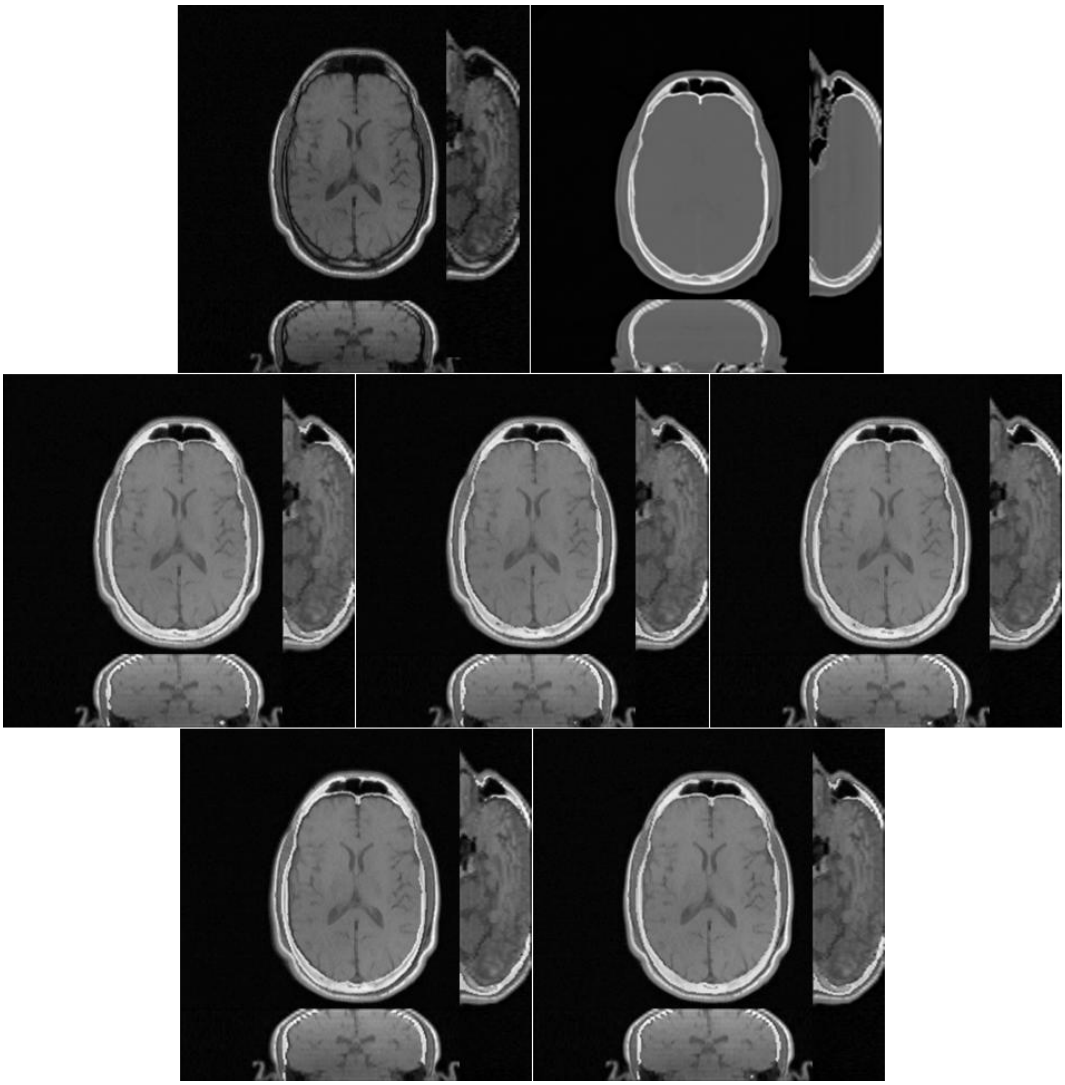

Figure 5.27 DS-27 registration results: First row, reference orthogonal MR and float orthogonal CT slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.30 Relative registration errors for DS-27.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*2,40*	*1,11*	*2,10*	*0,88*
NM	*2,40*	-	*2,48*	*1,28*	*2,29*
HJ	*1,11*	*2,48*	-	*2,41*	*0,85*
SA	*2,10*	*1,28*	*2,41*	-	*2,13*
PW	*0,88*	*2,29*	*0,85*	*2,13*	-

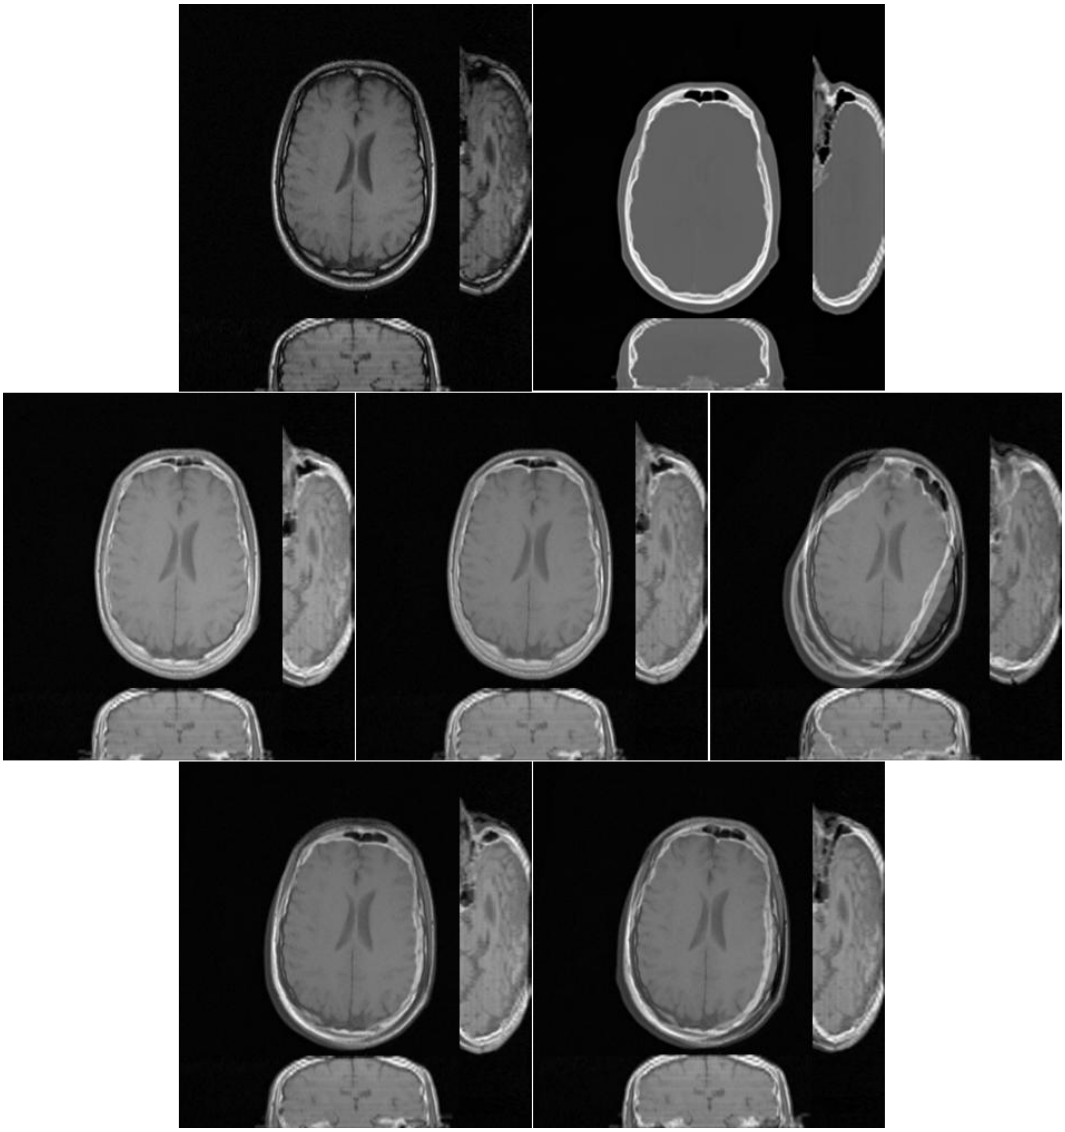

Figure 5.28 DS-28 registration results: First row, reference orthogonal MR and float orthogonal CT slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.31 Relative registration errors for DS-28.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*3,30*	*77,40*	*13,06*	*20,20*
NM	*3,30*	-	*75,22*	*10,66*	*18,14*
HJ	*77,40*	*75,22*	-	*65,82*	*58,27*
SA	*13,06*	*10,66*	*65,82*	-	*8,77*
PW	*20,20*	*18,14*	*58,27*	*8,77*	-

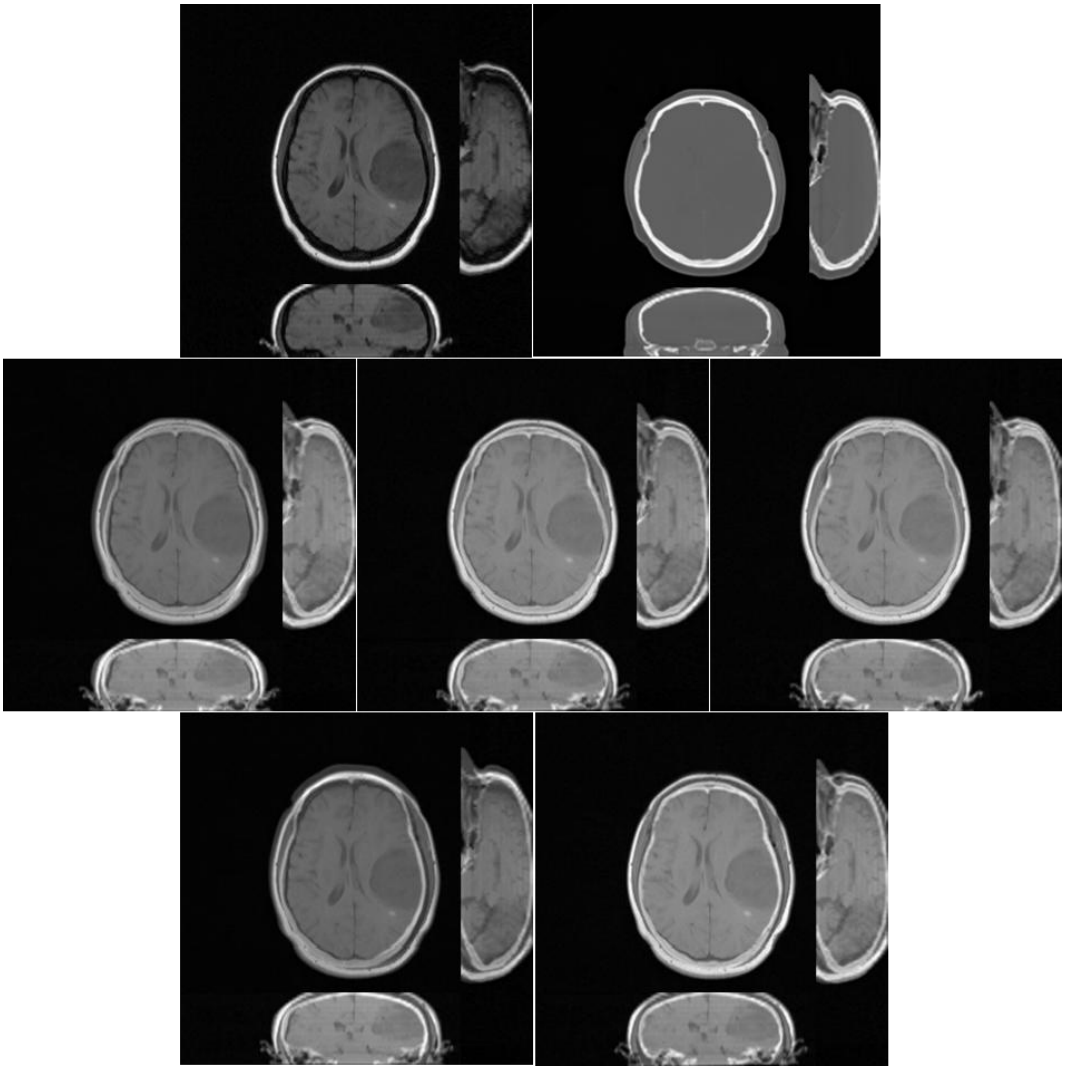

Figure 5.29 DS-29 registration results: First row, reference orthogonal MR and float orthogonal CT slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.32 Relative registration errors for DS-29.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*4,36*	*4,63*	*6,92*	*6,62*
NM	*4,36*	-	*6,19*	*6,38*	*3,50*
HJ	*4,63*	*6,19*	-	*5,86*	*6,16*
SA	*6,92*	*6,38*	*5,86*	-	*6,38*
PW	*6,62*	*3,50*	*6,16*	*6,38*	-

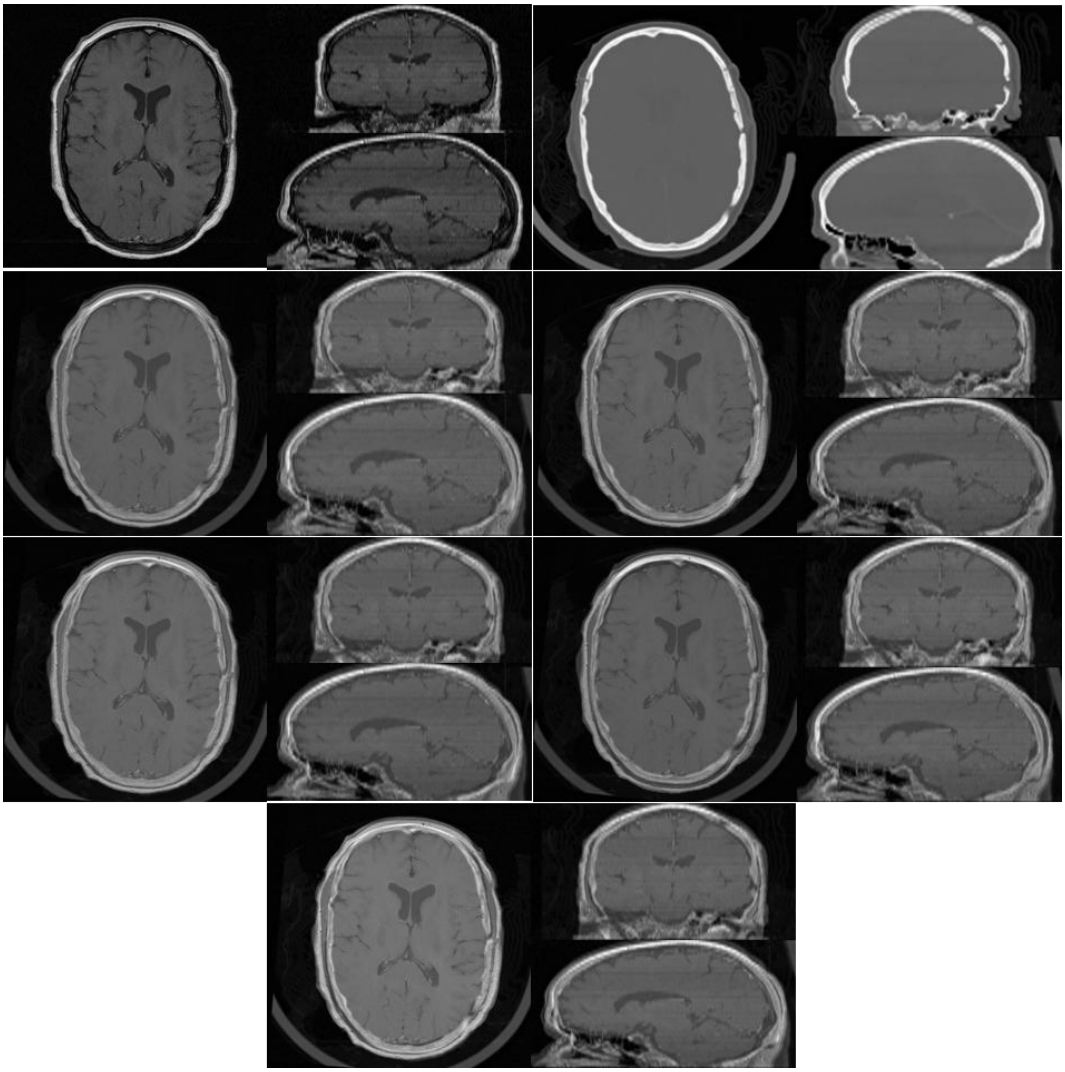

Figure 5.30 DS-30 registration results: First row, reference orthogonal MR and float orthogonal CT slices; second row, superimposed registered slices from SPSA and NM; third row HJ and SA algorithms; Last row, superimposed registered slices from PW algorithm.

Table 5.33 Relative registration errors for DS-30.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*6,44*	*2,24*	*2,58*	*3,15*
NM	*6,44*	-	*6,90*	*7,98*	*3,97*
HJ	*2,24*	*6,90*	-	*3,90*	*3,88*
SA	*2,58*	*7,98*	*3,90*	-	*5,37*
PW	*3,15*	*3,97*	*3,88*	*5,37*	-

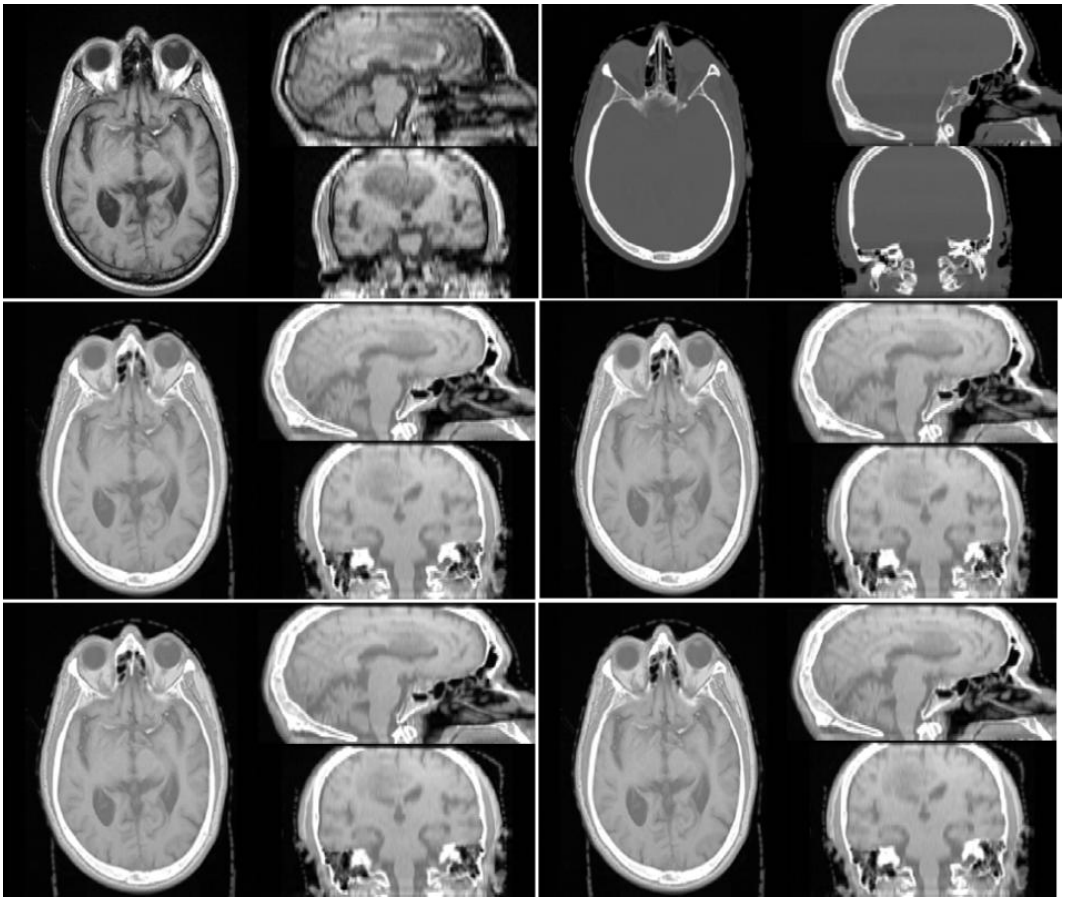

Figure 5.31 DS-31 registration results: First row, reference orthogonal MR and float orthogonal CT slices; middle row, superimposed registered slices from SPSA and NM algorithms; Last row, superimposed registered slices from HJ and SA algorithms

Table 5.34 Relative registration errors for DS-31.

	SPSA	***NM***	***HJ***	***SA***
SPSA	-	*1,71*	*1,00*	*2,93*
NM	*1,71*	-	*2,34*	*2,35*
HJ	*1,00*	*2,34*	-	*3,71*
SA	*2,93*	*2,35*	*3,71*	-

5.2.3.3 CT – MR-T2 Registrations

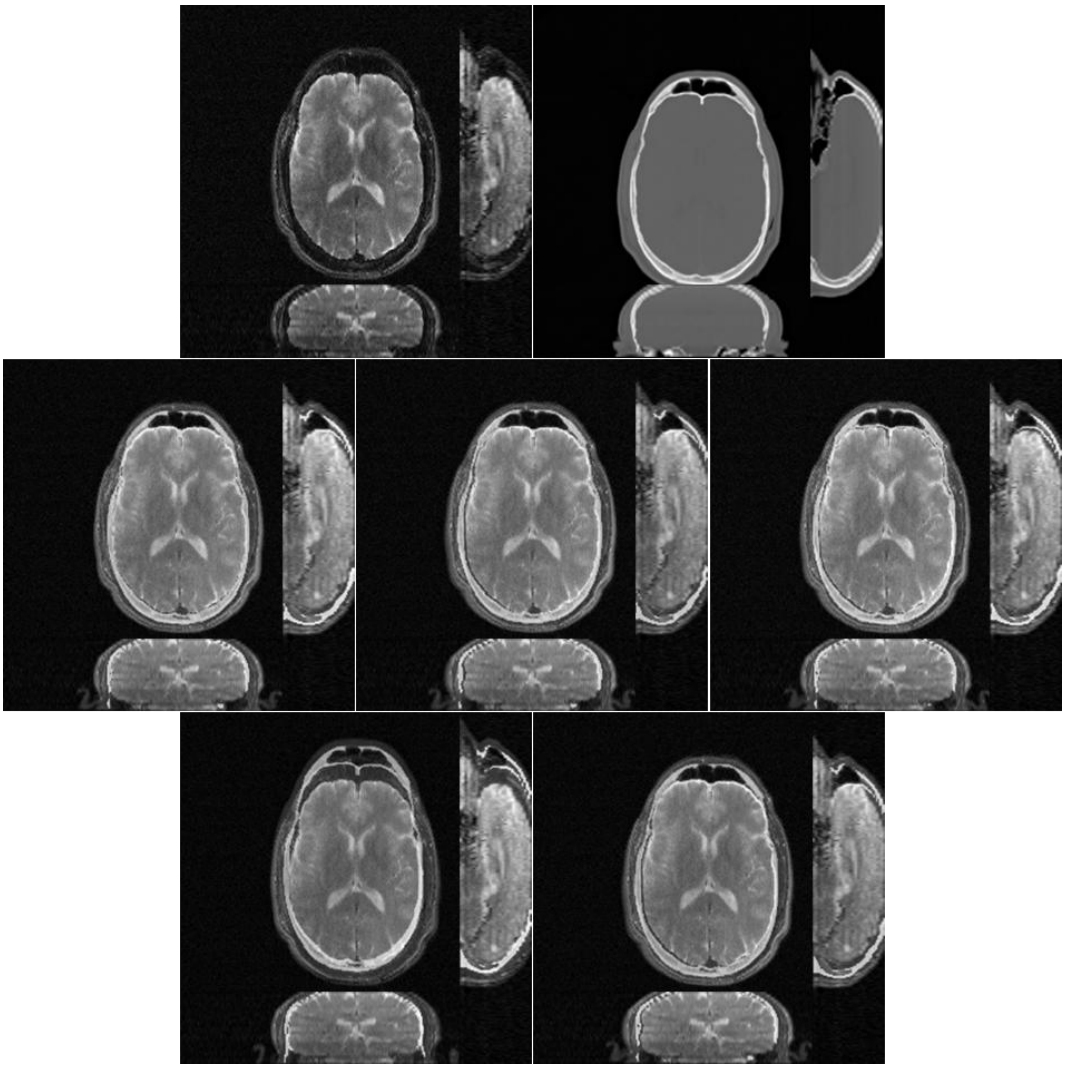

Figure 5.32 DS-32 registration results: First row, reference orthogonal MR and float orthogonal CT slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.35 Relative registration errors for DS-32.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*2,07*	*3,27*	*10,03*	*2,27*
NM	*2,07*	-	*2,60*	*9,64*	*1,11*
HJ	*3,27*	*2,60*	-	*8,90*	*2,20*
SA	*10,03*	*9,64*	*8,90*	-	*10,03*
PW	*2,27*	*1,11*	*2,20*	*10,03*	-

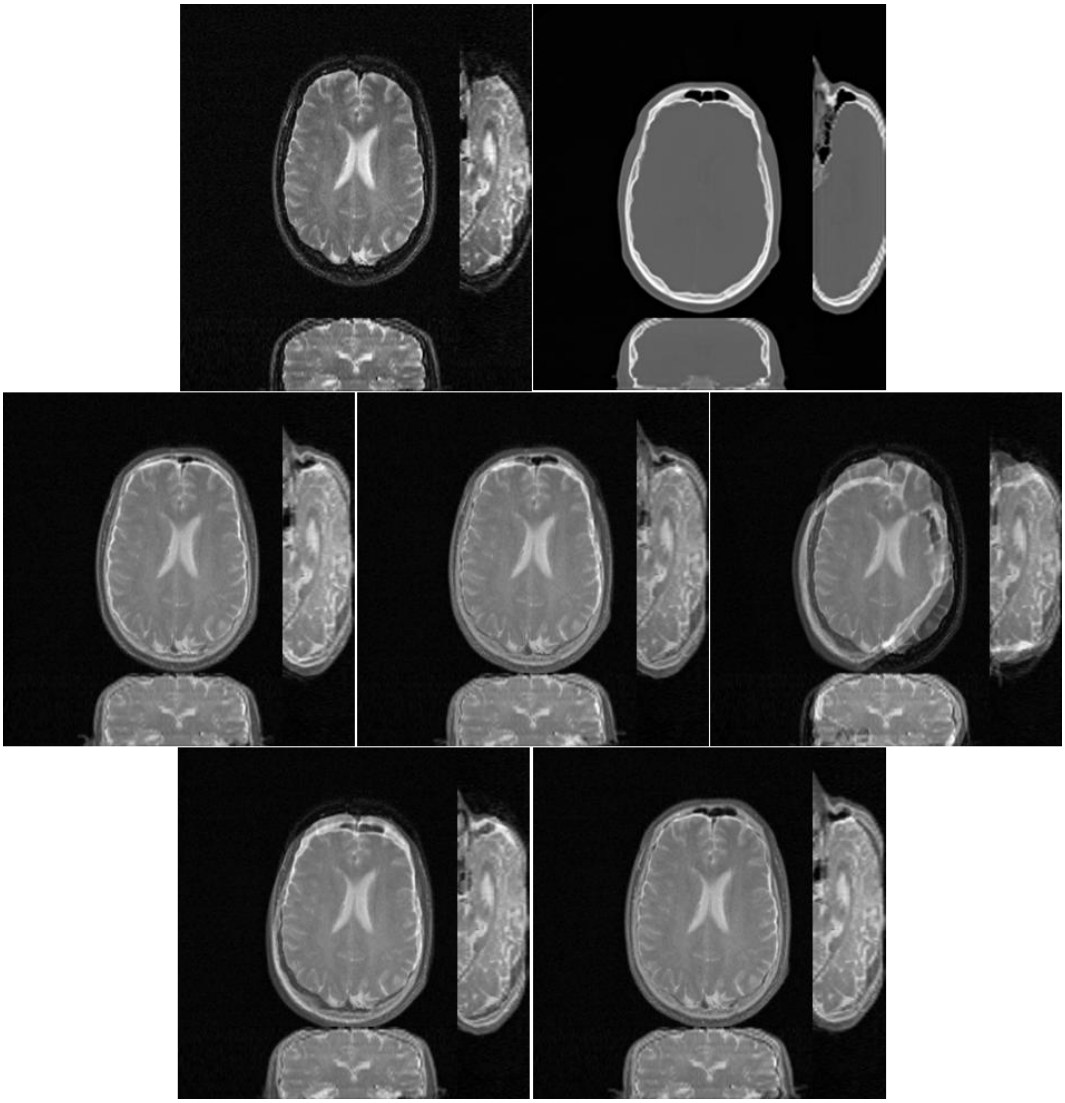

Figure 5.33 DS-33 registration results: First row, reference orthogonal MR and float orthogonal CT slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.36 Relative registration errors for DS-33.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*5,38*	*120,77*	*11,80*	*3,68*
NM	*5,38*	-	*120,27*	*10,05*	*3,76*
HJ	*120,77*	*120,27*	-	*113,65*	*120,68*
SA	*11,80*	*10,05*	*113,65*	-	*11,10*
PW	*3,68*	*3,76*	*120,68*	*11,10*	-

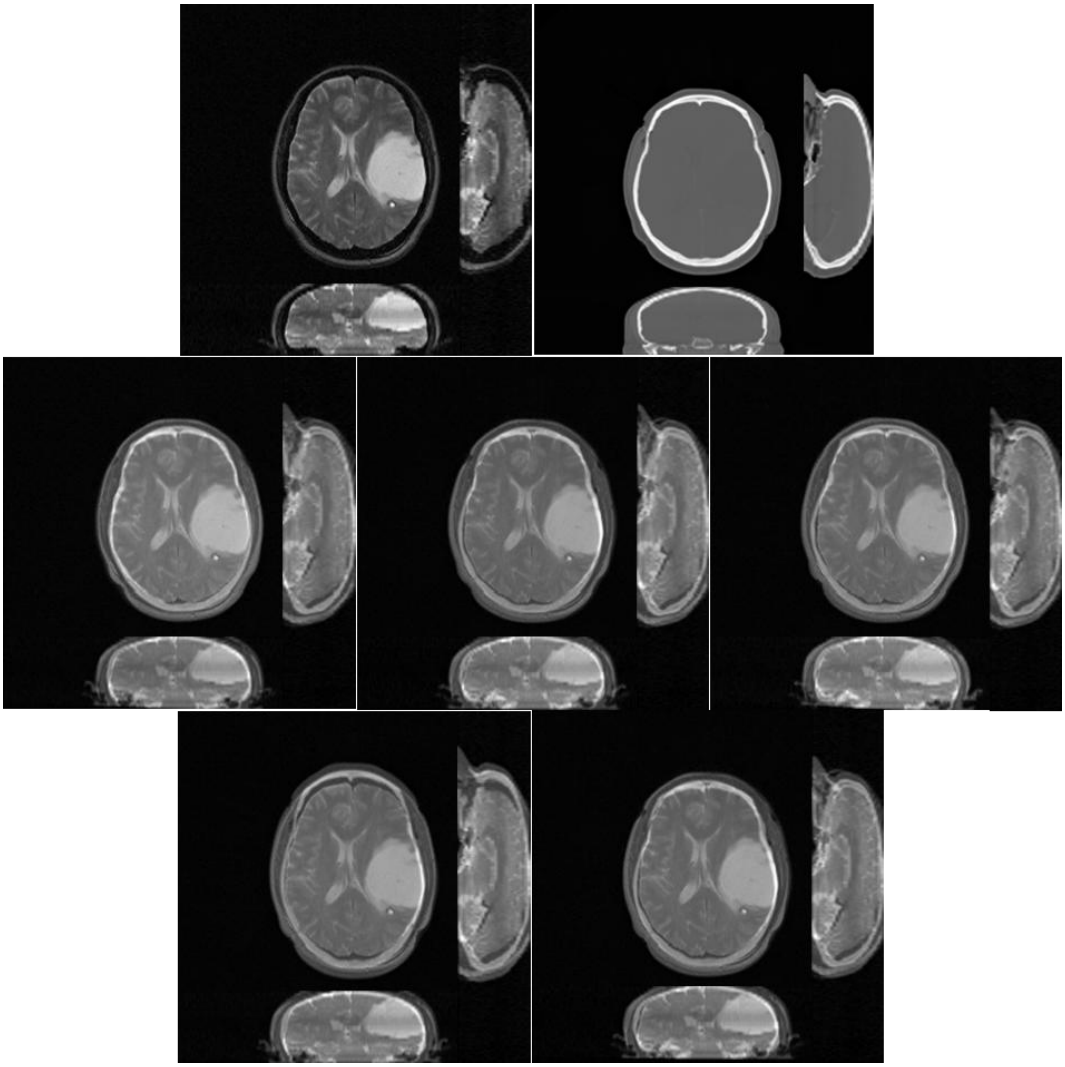

Figure 5.34 DS-34 registration results: First row, reference orthogonal MR and float orthogonal CT slices; middle row, superimposed registered slices from SPSA, NM and HJ algorithms; Last row, superimposed registered slices from SA and PW algorithms

Table 5.37 Relative registration errors for DS-34.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*6,64*	*9,36*	*8,56*	*5,46*
NM	*6,64*	-	*5,05*	*7,26*	*1,78*
HJ	*9,36*	*5,05*	-	*9,06*	*6,50*
SA	*8,56*	*7,26*	*9,06*	-	*7,19*
PW	*5,46*	*1,78*	*6,50*	*7,19*	-

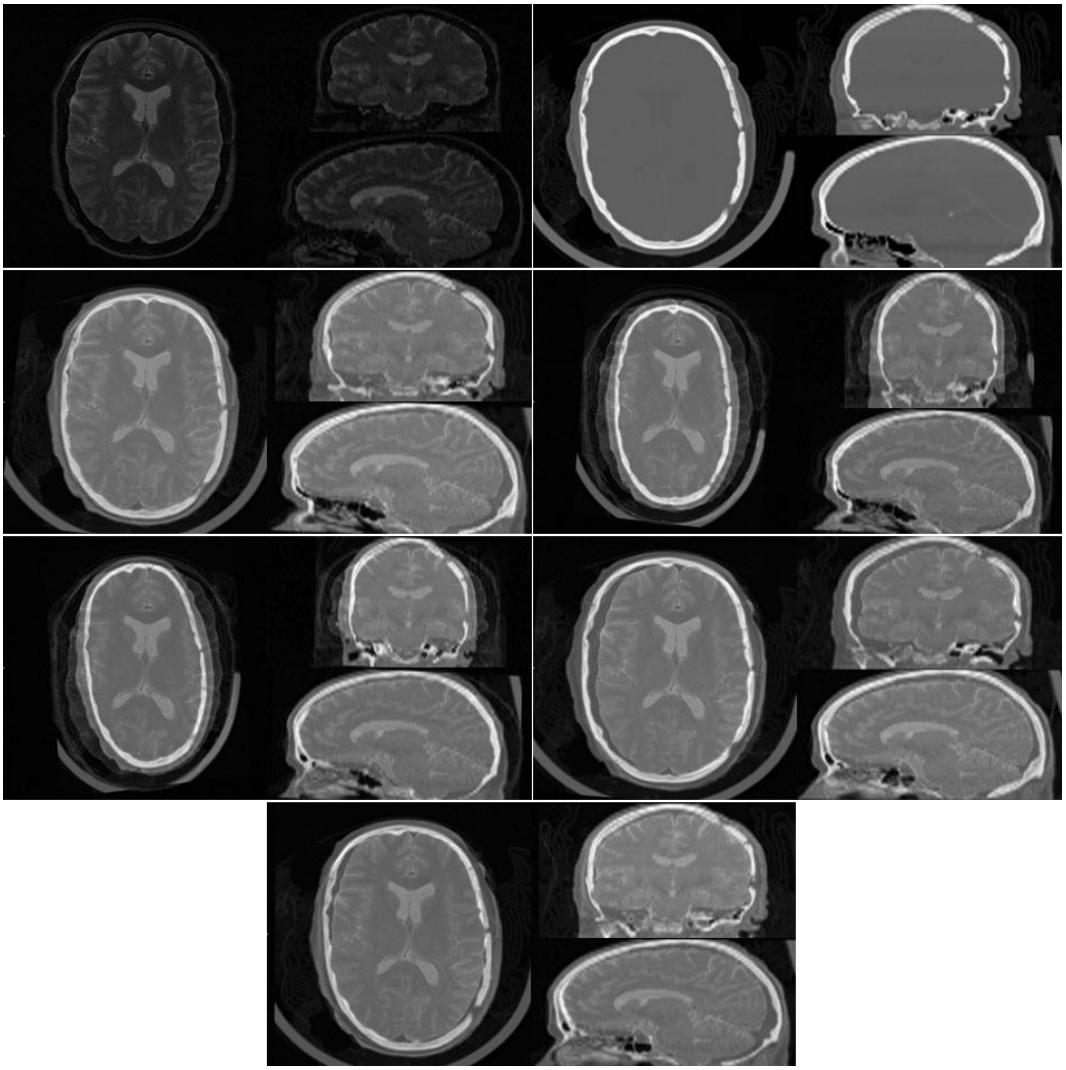

Figure 5.35 DS-35 registration results: first row, reference orthogonal MR and float orthogonal CT slices; second row, superimposed registered slices from SPSA and NM; third row HJ and SA algorithms; last row, superimposed registered slices from PW algorithm.

Table 5.38 Relative registration errors for DS-35.

	SPSA	***NM***	***HJ***	***SA***	***PW***
SPSA	-	*24,24*	*31,75*	*8,22*	*11,56*
NM	*24,24*	-	*18,16*	*27,12*	*25,50*
HJ	*31,75*	*18,16*	-	*32,20*	*26,32*
SA	*8,22*	*27,12*	*32,20*	-	*10,13*
PW	*11,56*	*25,50*	*26,32*	*10,13*	-

As given by the results, from Figure 5.24 to Figure 5.35, 3-Slice - SPSA combination method gives the most stable and accurate results, on the bases of our visual evaluation in the absence of true registration parameters. Since the human eye can not detect the registration error below 2 mm (Hajnal et al., 2001), then by visual evaluation they can be considered as in the range of clinical acceptance.

Table 5.39 shows that SPSA, NM and PW algorithms give relatively close results as over-all performance. Although SA and HJ optimizations do not seem to be stable at all, SA gives relatively closer results to the others respectively. With a higher resolution of CT volumes there are much less difficulty of getting trapped in local minimas with respect to the low resolution PET volumes.

Table 5.39 CT – MR Registrations' average relative errors

		SPSA	***NM***	***HJ***	***SA***	***PW***
CT MR-PD	***SPSA***	-	*4.68±0.8*	*32.8±31.6*	*12.2±5.02*	*3.2±0.8*
	NM	*4.68±0.8*	-	*30.03±32.1*	*8.5±3.9*	*3.5±1.7*
	HJ	*32.8±31.6*	*30.03±32.1*	-	*22.8±27.06*	*20.9±31.9*
	SA	*12.2±5.02*	*8.5±3.9*	*22.8±27.06*	-	*11.1±5.6*
	PW	*3.2±0.8*	*3.5±1.7*	*20.9±31.9*	*11.1±5.6*	-
		SPSA	***NM***	***HJ***	***SA***	***PW***
CT MR-T1	***SPSA***	-	*2.5±1.6*	*10.9±30.1*	*5.01±4.1*	*4.3±7.3*
	NM	*2.5±1.6*	-	*18.6±28.3*	*5.7±3.4*	*5.5±6.4*
	HJ	*10.9±30.1*	*18.6±28.3*	-	*16.3±24.7*	*13.8±22.3*
	SA	*5.01±4.1*	*5.7±3.4*	*16.3±24.7*	-	*4.5±3.1*
	PW	*4.3±7.3*	*5.5±6.4*	*13.8±22.3*	*4.5±3.1*	-
		SPSA	***NM***	***HJ***	***SA***	***PW***
CT MR-T2	***SPSA***	-	*9.5±8.6*	*41.2±47.1*	*9.6±1.4*	*5.7±3.5*
	NM	*9.5±8.6*	-	*36.5±44.9*	*13.5±8.9*	*8.03±9.1*
	HJ	*41.2±47.1*	*36.5±44.9*	-	*40.9±43.1*	*38.9±48.1*
	SA	*9.6±1.4*	*13.5±8.9*	*40.9±43.1*	-	*9.6±1.4*
	PW	*5.7±3.5*	*8.03±9.1*	*38.9±48.1*	*9.6±1.4*	-

Table 5.40 gives the average process time of the optimization algorithms with the 3-Slice algorithm. According to these results, SPSA, HJ and SA are close in over-all speed performance. But when considering the accuracy/speed efficiency SPSA outperforms the other two algorithms.

Table 5.40 CT – MR Registrations' average process time with 3-Slice method (Sec ± StD).

	CT MR-PD	***CT MR-T1***	***CT MR-T2***
SPSA	*51.5 ± 0.51*	*51.3 ± 0.34*	*51.1 ± 0.35*
NM	*120.7 ± 4.8*	*115.7 ± 6.3*	*122.1 ± 3.15*
HJ	*66.2 ± 7.1*	*69.1 ± 3.1*	*65.9 ± 7.06*
SA	*54.04 ± 2.61*	*54.01 ± 3.71*	*55.2 ± 5.1*
PW	*225 ± 118.9*	*156.2 ± 66.7*	*213.2 ± 73*

CHAPTER SIX
CONCLUSION

There now exists a large number of technologies to represent the various anatomic and metabolic processes of the human body. The information within these different modalities finds application in many facets of medicine. However, a greater clinical potential exists in combining the information of different modalities into a single representation, such as the task of multimodal image registration.

Image registration involves determining the parameters which bring into alignment the information from different images. There has been considerable researches invested in developing reliable methods of registering medical images from different modalities. Increasingly, the technologies employed in registering image volumes achieve a high level of automation and accuracy.

Registration by maximization of normalized mutual information is an intensity based method of registering images derived from principles of information theory. An attractive feature of the method is that its correlation is statistically based and is not dependent on specific structures in an image. As such, it avoids many of the complications which befall other registration routines.

Based on the theory of mutual information, a fully automated robust and fast registration routine is proposed which has high level accuracy in registering various sets of multimodal volume images.

A novel approach was introduced, called *3-Slice Method*, to reduce the computational burden of the transformation and NMI cost function. Proposed volume image registration method optimizes normalized mutual information measure using SPSA algorithm. The 3-Slice method uses only three orthogonal central slices of the volumes to build joint histogram instead of entire volumes. The major advantages of the proposed method are: 1) fast and efficient computation of NMI for a statistical

similarity between volume images, and 2) faster and accurate convergence with using SPSA algorithm in comparison with other optimization techniques. Beside these major advantages provided by the method, we present an approach to calculate the SPSA gain parameters, a_k and c_k, automatically by developing an algorithm instead of experimental and manual pre-registration calculations which are usually time consuming and tedious.

We also believe that the algorithm proposed to calculate the gain paramaters automatically will increase the popularity of the SPSA algorithm in not only image registration field but also in any other field requiring fast optimization of a multiparameter function. The 3-Slice method provides substantial computational efficiency for both calculating similarity measure and transforming the volume images. For a volume image which has 217x181x81 resolution, it reduces the number of voxels, which will be transformed, by almost a factor of 65, consequently, decreases the overall registration process time by a factor of 10 to 40 depending on the used optimization algorithm.

Experimental results show that 3-Slice - SPSA combination provides sufficiently accurate results in case of PET-MR registrations as well as CT–MR registrations. While 3-Slice approach provides fast calculation of NMI, SPSA optimization improves the overall performance of speed and accuracy with its fast convergence capability. Since PET images have very low resolutions, the optimization algorithms mostly have major difficulties in convergence to the true registration parameters, except the SPSA. The SPSA has the advantage of avoiding the local minimas, unlike the deterministic gradient based methods which have tendency of getting trapped in local minimas. The SPSA promises high computational efficiency and global search capabilities even if the volumes have considerably low resolutions.

Nowadays, it is common to use high resolution images (i.e. their dimensions are in the order of thousand square pixels) in a wide range of applications. Registration of such high resolution images requires more computational resources to obtain satisfactory performances in implementation. As a future work, several opportunities

can be investigated to increase speed and accuracy. In order to speed up the implementation of the proposed algorithm, parallel processing and\or special hardware solutions such as graphics processing units can be employed for computationally intensive tasks, and for more accurate registration, a deterministic optimization algorithm could be combined with SPSA as a hybrid optimization.

REFERENCES

Aczél J. & Daroczy Z. (1975). *On measures of information and their characterizations*. New York: Academic Press.

Bajcsy R. & Kovacic S. (1989). Multiresolution elastic matching. *Comput. Vision Graph. Image Processing*. 46, 1–21.

Baker M., Busse H., & Vogt M. (2000). An automatic registration and segmentation algorithm for multiple electrophoresis images. In K. M. Hanson (Ed.). *Medical Imaging: Image Processing, vol. 3979 of Proc. SPIE* (426–436). Bellingham, WA: SPIE Press.

Bansal R., Staib L., Chen Z., Rangarajan A., Knisely J., Nath R., & Duncan J. (1999). A minimax entropy registration framework for patient setup verification in radiotherapy. *Computer Aided Surgery, 4* (6), 287–304.

Bergmans L., Van Cleynenbreugel J., Wevers M., & Lambrechts P. (2001). A methodology for quantitative evaluation of root canal instrumentation using microcomputed tomography. *Internationl Endodontic Journal, 34* (5), (390–398).

Besl P. J. & McKay N. D. (1992). A method for registration of 3-D shapes. *IEEE Trans. Pattern Anal. Mach. Inteligence. 14*, 239–256.

Blackall J. M., Rueckert D., Maurer C. R., Jr., Penney G. P., Hill D. L. G., & Hawkes D. J. (2000). An image registration approach to automated calibration for freehand 3D ultrasound. In S. L. Delp, A. M. DiGioia, & B. Jaramaz, (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol. 1935 of Lecture Notes in Computer Science* (462–471). Berlin: Springer-Verlag.

BrainWeb, McConnell Brain Image Centre Montreal Neurological Inst, McGill University, Montreal, QC, Canada. Retrieved November 8, 2010, from http://www.bic.mni.mcgill.ca/brainweb

Broit C. (1981). *Optimal registration of deformed images*. Ph.D. thesis, University of

Pennsylvania, Pennsylvania, USA.

Capek M. & Krekule I. (1999). Alignment of adjacent Picture frames captured by a CLSM. *IEEE Transactions on Information Technology in Biomedicine, 3* (2), (119–124).

Carrillo A., Duerk J. L., Lewin J. S., & Wilson D. L. (2000). Semi-automatic 3-D image registration as applied to interventional MRI liver cancer treatment. *IEEE Transactions on Medical Imaging, 19* (3), 175–185.

Castellano-Smith A. D., Hartkens T., Schnabel J., Hose D. R., Liu H., Hall W. A., Truwit C. L., Hawkes D. J., & Hill D. L. G. (2001). Constructing patient specific models for correcting intraoperative brain deformation. In W. J. Niessen & M. A. Viergever (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol. 2208 of Lecture Notes in Computer Science* (1091–1098). Berlin: Springer-Verlag.

Castro-Pareja C. R., Jagadeesh J. M., & Shekhar R. (2004). FPGA-based acceleration of mutual information calculation for real-time 3D image registration. *Proc. SPIE 5297*, 212–219.

Chen H. & Varshney P. K. (2001). A cooperative search algorithm for mutual information based image registration. In B. V. Dasarathy (Ed.). *Sensor Fusion: Architectures, Algorithms, and Applications, vol. 4385 of Proc. SPIE* (117–128). Bellingham, WA: SPIE Press.

Chirstensen G. E., Rabbitt R. D., & Miller M. I. (1996). Deformable templates using large deformation kinematics. *IEEE Trans. Med. Imaging.* 5, 1435–1447.

Clarkson M. J., Rueckert D., Hill D. L. G., & Hawkes D. J. (1999). Registration of multiple video images to pre-operative CT for image guided surgery. In K. M. Hanson (Ed.). *Medical Imaging: Image Processing, vol.3661 of Proc. SPIE* (14–23). Bellingham, WA: SPIE Press.

Collignon A. (1998). *Multi-modality medical image registration by maximization of mutual information.* PhD thesis. Catholic University of Leuven, Leuven, Belgium.

Collignon A., Maes F., Delaere D., Vandermeulen D., Suetens P., & Marchal G. (1995). Automated multi-modality image registration based on information theory. In Y. Bizais, C. Barillot, & R. Di Paola (Eds.). *Information Processing in Medical Imaging* (263–274). Dordrecht: Kluwer Academic Publishers.

Collignon A., Vandermeulen D., Suetens P., & Marchal G. (1995). 3D multi-modality medical image registration using feature space clustering. In N. Ayache (Eds.). *Computer Vision, Virtual Reality, and Robotics in Medicine, vol. 905 of Lecture Notes in Computer Science* (195–204). Berlin: Springer-Verlag.

Collins D. L., Holmes C. J., Peters T. M., & Evans A. C. (1995). Automatic 3D model- based neuroanatomical segmentation. *Human Brain Mapping*. 3, 190–208.

Collins D. L., Zijdenbos A. P., Killokian V., Sled J. G., Kabani N. J., Holmes C. J., & Evans A. C. (1998). Design and construction of a realistic digital brain phantom. *IEEE Trans. Med. Imaging*. 17, 463–468.

Cover T. M. & Thomas J. A. (1991). *Elements of Information Theory*. New York: John Wiley & Sons.

Cox R.W., & Jesmanowicz A. (1999). Real-time 3D image registration for functional MRI. *Magn. Reson. Med*. 42, 1014-1018.

Davis M. H., Khotanzad A., Flamig D. P., & Harms S. E. (1997). A physics-based coordinate transformation for 3-D image matching. *IEEE Trans. Pattern Anal. Mach. Intelligence*. 16, 317–328.

Denton E. R. E., Holden M., Christ E., Jarosz J. M., Russell- Jones D., Goodey J., Cox T. C. S., & Hill D. L. G. (2000). The identification of cerebral volume changes in treated growth hormone-deficient adults using serial 3D MR image processing. *Journal of Computer Assisted Tomography, 24* (1), 139–145.

Denton E. R. E., Sonoda L. I., Rueckert D., Rankin S. C., Hayes C., Leach M. O., Hill D. L. G., & Hawkes D. J. (1999). Comparison and evaluation of rigid, affine,

and nonrigid registration of breast MR images. *Journal of Computer Assisted Tomography, 23* (5), 800–805.

Ding L., Goshtasby A., & Satter M. (2001). Volume image registration by template matching. *Image and Vision Computing, 19* (12), 821–832.

Eddy W. F., Fitzgerald M., & Noll D. C. (1996). Improved image registration by using Fourier interpolation. *Magn. Reson. Med.* 36, 923–931.

Erdi Y. E., Rosenzweig K., Erdi A. K., Macapinlac H. A., Hu Y. C., Braban L. E., Humm J. L., Squire O. D., Chui C. S., Larson S. M., & Yorke E. D. (2002). Radiotherapy treatment planning for patients with non-small cell lung cancer using positron emission tomography (PET). *Radiotherapy and Oncology, 62* (1), 51–60.

Estevez P. A., Tesmer M., Perez C. A., & Zurada J. M. (2009). Normalized mutual information feature selection. *IEEE Transactions on Neural Networks, 20* (2), 189 – 201.

Fei B., Wheaton A., Lee Z., Duerk J. L., & Wilson D. L. (2002). Automatic MR volume registration and its evaluation for the pelvis and prostate. *Physics in Medicine and Biology, 47* (5), 823–838.

Fitzpatrick J. M., Hill D. L. G., & Maurer C. R. Jr. (2000). Image Registration. In Sonka M. & Fitzpatrick J. M (Ed). *Handbook of Medical Imaging: Medical image processing and analysis* (447-513). Washington: SPIE Press.

Freeborough P. A. & Fox N. C. (1997). The boundary shift integral: an accurate and robust measure of cerebral volume changes from registered repeat MRI. *IEEE Trans. Med. Imaging.* 16, 623–629.

Freire L. & Mangin J. F. (2001). Motion correction algorithms may create spurious brain activations in the absence of subject motion. *NeuroImage, 14* (3), 709–722.

Friston K. J., Williams S., Howard R., Frackowiak R. S., & Turner R. (1996). Movement- related effects in fMRI time-series. *Magn. Reson. Med.* 35, 346–355.

Fu K. C., Gonzalez R. C., & Lee C. S. G. (1987). *Robotics: Control, Sensing, Vision, and Intelligence*. New York: McGraw-Hill.

Gaens T., Maes F., Vandermeulen D., & Suetens P. (1998). Non-rigid multimodal image registration using mutual information. In W. M. Wells, A. Colchester, & S. Delp, (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol.1496 of Lecture Notes in Computer Science* (1099–1106). Berlin Springer-Verlag.

Gee J. C., Reivich M., & Bajcsy R. (1993). Elastically deforming a three-dimensional atlas to match anatomical brain images. *J. Comput. Assist. Tomogr.* 17, 225–236.

Goldstein H. (1950). *Classical Mechanics*. Reading, MA: Addison-Wesley.

Goshtasby A. (1986). Piecewise linear mapping functions for image registration. *Pattern Recognition*. 19, 459–466.

Goshtasby A. (1988). Registration of images with geometric distortions. *IEEE Trans. Geosci. Remote Sensing*. 26, 60–64.

Hajnal J. V., Hill D. L. G. & Hawkes D. J. (Eds.). (2001). *Medical Image Registration*. Florida: CRC Press LLC.

Hajnal J. V., Saeed N., Soar E. J., Oatridge A., Young I. R., & Bydder G. M. (1995). A registration and interpolation procedure for subvoxel matching of serially acquired MR images. *J. Comput. Assist. Tomography.* 19, 289–296.

Haralick R. M. & Shapiro L. G. (1993). *Computer and Robot Vision*, Vol. II. Reading, MA: Addison-Wesley.

Hartley R. V. L. (1928). Transmission of information. *Bell System Technical Journal, 7*, 535–563.

Hata N., Dohi T., Warfield S., Wells W., Kikinis R., & Jolesz F. A. (1998). Multimodality deformable registration of pre- and intraoperative images for MRI-guided brain surgery. In W. M. Wells, A. Colchester, & S. Delp (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol.1496 of Lecture Notes in Computer Science* (1067–1074). Berlin: Springer-Verlag.

Hayton P. M., Brady M., Smith S. M., & Moore N. (1999). A non-rigid registration algorithm for dynamic breast MR images. *Artificial Intelligence, 114* (1-2), 125–156.

Hellier P. & Barillot C. (2000). Multimodal non-rigid warping for correction of distortions in functional MRI. In S. L. Delp, A. M. DiGioia, & B. Jaramaz (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol. 1935 of Lecture Notes in Computer Science* (512–520). Berlin: Springer-Verlag.

Hermosillo G. & Faugeras O. (2001). Dense image matching with global and local statistical criteria: a variational approach. In *Computer Vision and Pattern Recognition* (vol.1) (73–78). Los Alamitos, CA: IEEE Computer Society Press.

Hill D. L. G., Maurer C. R., Jr., Maciunas R. J., Barwise J. A., Fitzpatrick J. M., & Wang M. Y. (1998). Measurement of intraoperative brain surface deformation under a craniotomy. *Neurosurgery, 43* (3), 514–526.

Hill D. L. G., Batchelor P. G., Holden M., & Hawkes D. J. (2001). Medical image registration. *Physics in Medicine and Biology, 46* (3) , 1 – 45.

Hill D. L. G., Hawkes D. J., Gleeson M. J., Cox T. C. S., Strong A. J., Wong W.L., Ruff C. F., Kitchen N. D., Thomas D. G. T., Crossman J. E., Studholme C., Gandhe A. J., Green S. E. M., & Robinson G. P. (1995). Accurate frameless registration of MR and CT images of the head: Applications in surgery and radiotherapy planning. *Radiology*, 191, 447–454.

Hill D. L. G., Hawkes D. J., Harrison N. A., & Ruff C. F. (1993) . A strategy for automated multimodality image registration incor porating anatomical knowledge and imager characteristics. In H. H. Barrett & A. F. Gmitro (Eds.). *Information*

Processing in Medical Imaging: vol. 687 of Lecture Notes in Computer Science (182–196). Berlin: Springer-Verlag.

Hill D. L. G., Maurer, Jr. C. R., Studholme C., Fitzpatrick J. M., & Hawkes D. J. (1998). Correcting scaling errors in tomographic images using a nine degree of freedom registration algorithm. *Journal of Computer Assisted Tomography, 22* (2), 317–323.

Hill D. L. G., Studholme C., & Hawkes D. J. (1994). Voxel similarity measures for automated image registration. In R. A. Robb (Ed.). *Visualization in Biomedical Computing, vol. 2359 of Proc. SPIE* (205–216). Bellingham, WA: SPIE Press.

Holden M. (2001). *Registration of 3D serial MR brain images*. Ph.D. thesis. London University, London, UK.

Holden M., Denton E. R. E., Jarosz J. M., Cox T. C. S., Studholme C., Hawkes D. J., & Hill D. L. G. (1999). Detecting small anatomical changes with 3D serial MR subtraction images. In K. M. Hanson (Ed.). *Medical Imaging: Image Processing, vol. 3661 of Proc. SPIE* (44–55). Bellingham, WA: SPIE Press.

Holden M., Hill D. L. G., Denton E. R. E., Jarosz J. M., Cox T. C. S., Rohlfing T., Goodey J., & Hawkes D. J. (2000). Voxel similarity measures for 3-D serial MR brain image registration. *IEEE Transactions on Medical Imaging, 19* (2), 94–102.

Holden M., Schnabel J. A., & Hill D. L. G. (2001). Quantifying small changes in brain ventricular volume using non-rigid registration. In W. J. Niessen & M. A. Viergever (Eds.). *Medical Image Computing and Computer-Assisted Intervention: Lecture Notes in Computer Science vol. 2208* (49–56). Berlin: Springer-Verlag.

Holmes D. R., Camp J. J., & Robb R. A. (1996). Evaluation of search strategies and cost functions in optimizing voxel-based image registration. In Y. Kim (Ed.). Medical Imaging: Image Display, vol. 2707 of Proc. SPIE (554–562). Bellingham, WA: SPIE Press.

Horn B. K. P. (1986). *Robot Vision*. Cambridge, MA: MIT Press.

Horn B. K. P. (1987). Closed-form solution of absolute orientation using unit quaternions. *JOSA A.* 4(4), 629–642.

Horsfield M. A. (1999). Mapping eddy current induced fields for the correction of diffusion-weighted echo planar images. *Magnetic Resonance Imaging, 17* (9), 1335–1345.

Ioannides A. A., Liu L. C., Kwapien J., Drozdz S., & Streit M. (2000). Coupling of regional activations in a human brain during an object and face affect recognition task. *Human Brain Mapping, 11* (2), 77–92.

Jenkinson M. & Smith S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis, 5* (2), 143–156.

Jumarie G. (1997). A new information theoretic approach to the entropy of non-random discrete maps relation to fractional dimension and temperature of curves. *Chaos, Solitions and Fractals, 8* (6), 953–970.

Kagadis G. C., Delibasis K. K., Matsopoulos G. K., Mouravliansky N. A., & Asvestas P. A. (2002). A comparative study of surface and volume-based techniques for the automatic registration between CT and SPECT brain images. *Medical Physics, 29* (2), 201–213.

Kim B., Boes J. L., Bland P. H., & Meyer C. R. (1999). Out-of-plane non-linear warping of a slice into volume. In C. Taylor & A. Colchester (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol. 1679 of Lecture Notes in Computer Science* (638–645), Berlin: Springer-Verlag.

Kim B., Boes J. L., Bland P. H., Chevenert T. L., & Meyer C. R. (1999). Motion correction in fMRI via registration of individual slices into an anatomical volume. *Magnetic Resonance in Medicine, 41* (5), 964–972.

Kim B., Boes J. L., Frey K. A., & Meyer C. R. (1997). Mutual information for automated unwarping of rat brain autoradiographs. *NeuroImage, 5* (1), 31–40.

Kim J., Fessler J. A., Lam K. L., Balter J. M., & Ten Haken R. K. (2001). A feasibility study on mutual information based set-up error estimator for radiotherapy. *Medical Physics, 28* (12), 2507–2517.

Kim J., Fisher J. W., Tsai A., Wible C., Willsky A. S., & Wells W. M. (2000). Incorporating spatial priors into an information theoretic approach for fMRI data analysis. In S. L. Delp, A. M. DiGioia, & B. Jaramaz (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol. 1935 of Lecture Notes in Computer Science* (62–71). Berlin: Springer-Verlag.

Kim J., Li S., Pradhan D., Hammoud R., Chen Q., Yin F. F., Zhao Y., Kim J. H., & Movsas B. (2007). Comparison of similarity measures for rigid-body CT/dual x-ray image registrations. *Technol. Cancer Res. Treat., 6,* 337–346.

Kjems U., Strother S. C., Anderson J., Law I., & Hansen L. K. (1999). Enhancing the multivariate signal of [^{15}O] water PET studies with a new nonlinear neuroanatomical registration algorithm. *IEEE Transactions on Medical Imaging, 18* (4), 306–319.

Knops Z. F., Maintz J. B. A., Viergevera M. A., & Pluim J. P. W. (2006). Normalized mutual information based registration using k-means clustering and shading correction. *Medical Image Analysis, 10* (3), 432-439.

Krücker J. F., Meyer C. R., LeCarpentier G. L., Fowlkes J. B., & Carson P. L. (2000). 3D spatial compounding of ultrasound images using image-based nonrigid registration. *Ultrasound in Medicine and Biology, 26* (9), 1475–1488.

Lavallee S., Troccaz J., Sautot P., Mazier B., Cinquin P., Merloz P., & Chirossel J. P. (1996). Computer-assisted spinal surgery using anatomy-based registration. In R. H. Taylor, S. Lavallee, G. Burdea, & R. Mosges (Ed.). *Computer-Integrated Surgery: Technology and Clinical Applications* (425) , Cambridge, MA: MIT Press,.

Lehmann T.M., Gonner C., & Spitzer K. (1999). Interpolation methods in medical image processing. *IEEE Trans. Med. Imaging*. 18, 1049-1075.

Lemieux L. & Barker G. J. (1998). Measurement of small interscan fluctuations in voxel dimensions in magnetic resonance images using registration. *Medical Physics*. 25, 1049–1054.

Lemieux L., Kitchen N. D., Hughes S. W., & Thomas D. G. T. (1994). Voxel-based localization in frame-based and frameless stereotaxy and its accuracy. *Med. Physics*. 21, 1301–1310.

Likar B. & Pernus F. (2001). A hierarchical approach to elastic registration based on mutual information. *Image and Vision Computing, 19* (1-2), 33–44.

Liu Y., Collins R. T., & Rothfus W. E. (2001). Robust midsagittal plane extraction from normal and pathological 3-D neuroradiology images. *IEEE Transactions on Medical Imaging, 20* (3), (175–192).

Lynch J. A., Peterfy C. G., White D. L., Hawkins R. A., & Genant H. K. (1999). MRI-SPECT image registration using multiple MR pulse sequences to examine osteoarthritis of the knee. In K. M. Hanson (Ed.). *Medical Imaging: Image Processing vol. 3661 of Proc. SPIE* (68–77). Bellingham, WA: SPIE Press.

Maes F. (1998). *Segmentation and registration of multimodal medical images: from theory, implementation and validation to a useful tool in clinical practice*. Ph.D. thesis. Catholic University of Leuven, Leuven, Belgium.

Maes F., Collignon A., Vandermeulen D., Marchal G., & Suetens P. (1997). Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging, 16* (2), 187–198.

Maguire G. Q., Noz Jr., M. E., Rusinek H., Jaeger J., Kramer E. L., Sanger J. J., & Smith G. (1991). Graphics applied to medical image registration. *IEEE Comput. Graph. Applications*. 11, 20–29.

Maintz J. B. A. & Viergever M. A. (1998). A survey of medical image registration. *Medical Image Analysis, 2* (1), 1–36.

Maintz J. B. A., Meijering E. H. W., & Viergever M. A. (1998). General multimodal elastic registration based on mutual information. In K. M. Hanson (Ed.). *Medical Imaging: Image Processing, vol. 3338 of Proc. SPIE* (144–154). Bellingham, WA: SPIE Press.

Martens K., Verstreken K., Van Cleynenbreugel J., Van Brussel K., Goffin J., Marchal G., & Suetens P. (1999). Image-based planning and validation of C1-C2 transarticular screw fixation using personalized drill guides. In C. Taylor & A. Colchester (Eds.). Medical Image Computing and Computer-Assisted Intervention, vol. 1679 of Lecture Notes in Computer Science (860–867). Berlin: Springer-Verlag.

Mattes D., Haynor D. R., Vesselle H., Lewellen T. K., & Eubank,W. (2001). Nonrigid multimodality image registration. In M. Sonka & K. M. Hanson (Eds.). *Medical Imaging: Image Processing, vol. 4322 of Proc. SPIE* (1609–1620). Bellingham, WA: SPIE Press.

Maurer Jr. C. R., Gaston R. P., Hill D. L. G., Gleeson M. J., Taylor M. G., Felon M. R., et al. (1999). AcouStick: A tracked A-mode ultrasonography system for registration in image-guided surgery. In C. J. Taylor & A. C. F. Colchester (Ed.). *Medical Image Compuing and Computer-Assisted Intervention (MICCAI)* (953–962). Berlin: Springer-Verlag.

Maurer, C.R. Jr. & Fitzpatrick J. M. (1993). A review of medical image registration. In R. J. Maciunas (Ed.). *Interactive Image-Guided Neurosurgery* (17–44). Park Ridge, IL: American Association of Neurological Surgeons.

Maurer, Jr. C. R., Aboutanos G. B., Dawant B. M., Maciunas R. J., & Fitzpatrick J. M. (1996). Registration of 3-D images using weighted geometrical features. *IEEE Trans. Med. Imaging.*15, 836–849.

Maurer, Jr. C. R., Maciunas R. J., & Fitzpatrick J. M. (1998). Registration of head CT images to physical space using a weighted combination of points and surfaces. *IEEE Trans. Med. Imaging.* 17, 753–761.

Meyer C. R., Boes J. L., Kim B., Bland P. H., Lecarpentier G. L., Fowlkes J. B., Roubidoux M. A., & Carson P. L. (1999). Semiautomatic registration of volumetric ultrasound scans. *Ultrasound in Medicine and Biology, 25* (3), 339–347.

Meyer C. R., Boes J. L., Kim B., Bland P. H., Zasadny K. R., Kison P. V., Koral K., Frey K. A., & Wahl R. L. (1997). Demonstration of accuracy and clinical versatility of mutual information for automatic multimodality image fusion using affine and thin-plate spline warped geometric deformations. *Medical Image Analysis, 1* (3), 195–206.

Miga M. I., Paulsen K. D., Lemery J. M., Eisner S. D., Hartov A., Kennedy F. E., & Roberts D. W. (1999). Model-updated image guidance: initial clinical experiences with gravity-induced brain deformation. *IEEE Trans. Med. Imaging*. 18, 866–874.

Nikou C., Heitz F., & Armspach J. P. (1999). Robust voxel similarity metrics for the registration of dissimilar single and multimodal images. Pattern Recognition, 32 (8), 1351–1368.

Ourselin S., Roche A., Prima S., & Ayache N. (2000). Block matching: a general framework to improve robustness of rigid registration of medical images. In S. L. Delp, A. M. DiGioia, & B. Jaramaz (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol. 1935 of Lecture Notes in Computer Science* (557–566). Berlin: Springer-Verlag.

Parker J., Kenyon R.V., & Troxel D. (1983). Comparision of interpolating methods for image resampling. *IEEE Trans. Med. Imaging*. 2, 31-39.

Pelizzari C. A., Chen G. T. Y., Spelbring D. R., Weichselbaum R. R., & Chen C. T. (1989). Accurate three-dimensional registration of CT, PET, and/or MR images of the brain. *J. Comput. Assist. Tomography.* 13, 20–26.

Penny G., Weese J., Desmedt P., Hill D. L. G., & Hawkes D. J. (1998). A comparison of similarity measures for use in 2D-3D medical image registration. *IEEE Transactions on Medical Imaging, 17* (4), 586–595.

Plattard D., Soret M., Troccaz J., Vassal P., Giraud J. Y., Champleboux G., Artignan X., & Bolla M. (2000). Patient set-up using portal images: 2D/2D image registration using mutual information. *Computer Aided Surgery, 5* (4), 246–262.

Pluim J. P. W. (2001). *Mutual information based registration of medical images*. Ph.D. thesis. Utrecht University, Utrecht, The Netherlands.

Pluim J. P. W., Maintz J. B. A., & Viergever M. A. (2000). Image registration by maximization of combined mutual information and gradient information. *IEEE Transactions on Medical Imaging, 19* (8), 809–814.

Pluim J. P. W., Maintz J. B. A., & Viergever M. A. (2000). Interpolation artefacts in mutual information-based image registration. *Computer Vision and Image Understanding*, 77 (2), 211-232.

Pluim J. P. W., Maintz J. B. A., & Viergever M. A. (2001). Mutual information matching in multiresolution contexts. *Image and Vision Computing, 19* (1-2), 45–52.

Pluim J. P. W., Maintz J. B. A., & Viergever M. A. (2003). Mutual information-based registration of medical images: A survey. *IEEE Trans. Med. Imaging*. 22(8), 986–1004.

Pompe B., Blidh P., Hoyer D., & Eiselt M. (1998). Using mutual information to measure coupling in the cardiorespiratory system. *IEEE Engineering in Medicine and Biology Magazine, 17* (6), 32–39.

Press W. H., Teukolsky S. A., Vetterling W. T., & Flannery B. P. (1992). *Numerical Recipes in C: The Art of Scientific Computing* (2nd ed). Cambridge, England: Cambridge University Press.

Radau P. E., Slomka P. J., Julin P., Svensson L., & Wahlund L. O. (2001). Evaluation of linear registration algorithms for brain SPECT and the errors due to hypoperfusion lesions. *Medical Physics, 28* (8), 1660–1668.

RIRE, Retrospective Image Registration Evaluation. Retrieved April 20, 2011, from http://www.insight-journal.org/rire/download_data.php

Ritter N., Owens R., Cooper J., Eikelboom R. H., & van Saarloos P. P. (1999). Registration of stereo and temporal images of the retina. *IEEE Transactions on Medical Imaging, 18* (5), 404–418.

Rodriguez-Carranza C. E. & Loew M. H. (1998). A weighted and deterministic entropy measure for image registration using mutual information. In K. M. Hanson (Ed.). *Medical Imaging: Image Processing, vol. 3338 of Proc. SPIE* (155–166). Bellingham, WA: SPIE Press.

Rodriguez-Carranza C. E. & Loew M. H. (1999). Global optimization of weighted mutual information for multimodality image registration. In K. M. Hanson (Ed.). *Medical Imaging: Image Processing, vol. 3338 of Proc. SPIE* (89–96). Bellingham, WA: SPIE Press.

Rohde G. K., Aldroubi A., & Dawant B. M. (2001). Adaptive free-form deformation for inter-patient medical image registration. In M. Sonka & K. M. Hanson, (Eds.). *Medical Imaging: Image Processing, vol. 4322 of Proc. SPIE* (1578–1587). Bellingham, WA: SPIE Press.

Rohlfing T. & Maurer C. R. Jr. (2001). Intensity-based non-rigid registration using adaptive multilevel free-form deformation with an incompressibility constraint. In W. J. Niessen & M. A. Viergever (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol. 2208 of Lecture Notes in Computer Science* (111–119). Berlin: Springer-Verlag.

Rohlfing T., Maurer C. R. Jr., O'Dell W. G., & Zhong J. (2001). Modeling liver motion and deformation during the respiratory cycle using intensity-based free-form registration of gated MR images. In S. K. Mun (Ed.). *Medical Imaging:*

Visualization, Image Display, and Image-Guided Procedures, vol. 4319 of Proc. SPIE (337–348). Bellingham, WA: SPIE Press.

Rohlfing T., West J. B., Beier J., Liebig T., Taschner C. A., & Thomale U. W. (2000). Registration of functional and anatomical MRI: accuracy assessment and application in navigated neurosurgery. *Computer Aided Surgery, 5* (6), 414–425.

Rosenfield G. H. (1959). The problem of exterior orientation in photogrammetry. *Photogrammetric Engineering*. 25, 536–553.

Rueckert D., Clarkson M. J., Hill D. L. G., & Hawkes D. J. (2000). Non-rigid registration using higher-order mutual information. In K. M. Hanson (Ed.). *Medical Imaging: Image Processing, vol. 3979 of Proc. SPIE* (438–447). Bellingham, WA: SPIE Press.

Rueckert D., Frangi A. F., & Schnabel J. A. (2001). Automatic construction of 3D statistical deformation models using non-rigid registration. In W. J. Niessen and M. A. Viergever (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol. 2208 of Lecture Notes in Computer Science*, (77–84). Berlin: Springer-Verlag.

Rueckert D., Hayes C., Studholme C., Summers P., Leach M., & Hawkes D. J. (1998). Non-rigid registration of breast MR images using mutual information. In A. C. F. Colchester, S. L. Delp, & W. M. Wells (Eds.). *Med. Image Comput. and Computer-Assisted Intervention* (3rd ed.). (1144–1152). Berlin: Springer-Verlag.

Rueckert D., Sonoda L. I., Hayes C., Hill D., Leach M. O., & Hawkes D. J. (1999). Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging, 18* (8), 712–721.

Sanjay-Gopal S., Chan H. P., Wilson T., Helvie M., Petrick N., & Sahiner B. (1999). A regional registration technique for automated interval change analysis of breast lesions on mammograms. *Medical Physics, 26* (12), (2669–2679).

Schnabel J. A., Rueckert D., Quist M., Blackall J. M., Castellano-Smith A. D., Hartkens T., Penney G. P., Hall W. A., Liu H., Truwit C. L., Gerritsen F. A., Hill D. L. G., & Hawkes D. J. (2001). A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations. In W. J. Niessen & M. A. Viergever (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol. 2208 of Lecture Notes in Computer Science* (573–581). Berlin: Springer-Verlag.

Schnabel J. A., Tanner C., Castellano-Smith A. D., Leach M. O., Hayes C., Degenhard A., Hose R., Hill D. L. G., & Hawkes D. J. (2001). Validation of non-rigid registration using finite elements methods. In M. F. Insana & R. M. Leahy (Eds.). *Information Processing in Medical Imaging, vol. 2082 of Lecture Notes in Computer Science* (344–357). Berlin: Springer-Verlag.

Shams R., Sadeghi P., Kennedy R. A., & Hartley R. (2010). A Survey of medical image registration on multi-core and GPU. *Signal Processing Magazine*. 27(2), 50–60.

Shannon C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27*, 379–423/623–656.

Shekhar R. & Zagrodsky V. (2002). Mutual information-based rigid and nonrigid registration of ultrasound volumes. *IEEE Transactions on Medical Imaging*, *21* (1), 9 – 22.

Simon D. A. & Kanade T. (1997). Geometric constraint analysis and synthesis: Methods for improving shape-based registration accuracy. In J. Troccaz, E. Grimson, & R. Mosges,(Ed.). *CVRMed-MRCAS '97* (181–190). Berlin: Springer-Verlag.

Skouson M. B., Guo Q., & Liang Z. P. (2001). A bound on mutual information for image registration. *IEEE Transactions on Medical Imaging, 20* (8), 843–846.

Slomka P. J., Mandel J., Downey D., & Fenster A. (2001). Evaluation of voxel-based registration of 3-D power Doppler ultrasound and 3-D magnetic resonance

angiographic images of carotid arteries. *Ultrasound in Medicine and Biology*. 27 (7), 945–955.

Spall J. C. (1998). An overview of the simultaneous perturbation method for efficient optimization. *Johns Hopkins APL Thecnical Digest.* 19(4), 482–492.

Spall J. C. (2003). *Introduction to Stochastic Search and Optimization: Estimation, Simulation and Control*. Hoboken, NJ: John Wiley & Sons Inc.

Studholme C. (1997). *Measures of 3D medical image alignment*. Ph.D. thesis, University of London, London, UK.

Studholme C., Constable R. T., & Duncan J. S. (2000). Accurate alignment of functional EPI data to anatomical MRI using a physics-based distortion model. *IEEE Transactions on Medical Imaging, 19* (11), 1115–1127.

Studholme C., Hill D. L. G. & Hawkes D. J. (1997). Automated 3D registration of MR and PET brain images by multi-resolution optimization of voxel similarity measures. *Medical Physics, 24* (1) , 25-35.

Studholme C., Hill D. L. G., & Hawkes D. J. (1995). Multiresolution voxel similarity measures for MR-PET registration. In Y. Bizais, C. Barillot, & R. Di Paola (Eds.). *Information Processing in Medical Imaging* (287–298). Dordrecht: Kluwer Academic Publishers.

Studholme C., Hill D. L. G., & Hawkes D. J. (1996). Automated 3-D registration of MR and CT images of the head. *Medical Image Analysis, 1* (2), 163–175.

Studholme C., Hill D. L. G., & Hawkes D. J. (1996). Incorporating connected region labelling into automated image registration using mutual information. In A. A. Amini, F. L Bookstein, & D. C. Wilson (Eds.). *Mathematical Methods in Biomedical Image Analysis* (23–31). Los Alamitos, CA: IEEE Computer Society Press.

Studholme C., Hill D. L. G., & Hawkes D. J. (1999). An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition, 32* (1), 71–86.

Studholme C., Novotny E., Zubal I. G., & Duncan J. S. (2001). Estimating tissue deformation between functional images induced by intracranial electrode implantation using anatomical MRI. *NeuroImage, 13* (4), 561–576.

Tanner C., Schnabel J. A., Chung D., Clarkson M. J., Rueckert D., Hill D. L. G., & Hawkes D. J. (2000). Volume and shape preservation of enhancing lesions when applying non-rigid registration to a time series of contrast enhancing MR breast images. In S. L. Delp, A. M. DiGioia, & B. Jaramaz, (Eds.). *Medical Image Computing and Computer-Assisted Intervention, vol. 1935 of Lecture Notes in Computer Science* (327–337). Berlin: Springer-Verlag.

Thévenaz P. & Unser M. (2000). Optimization of mutual information for multiresolution image registration. *IEEE Transactions on Image Processing, 9* (12), 2083–2099.

Thompson P. & Toga A. W. (1996). A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans. Med. Imaging.*15, 402–417.

Unser M. (1999). Splines: a perfect fit for signal and image processing. *IEEE Signal Processing Magazine*. 16, 22-38.

Vajda I. (1989). *Theory of statistical inference and information*. Dordrecht: Kluwer Academic Publishers.

Van den Elsen P. A. (1993). *Multimodality matching of brain images*. Ph.D. thesis. Utrecht University, Utrecht, Netherlands.

Van den Elsen P. A., Pol E. J. D., & Viergever M. A. (1993). Medical image matching – a review with classification. *IEEE Engineering in Medicine and Biology, 12* (1), 26–39.

Van Laere K., Koole M., D'Asseler Y., Versijpt J., Audenaert K., Dumont F., & Dierckx R. (2001). Automated stereotactic standardization of brain SPECT receptor data using single-photon transmission images. *Journal of Nuclear Medicine, 42* (2), 361–375.

Veninga T., Huisman H., van der Maazen R. W. M., Huizenga H. (2004). Clinical validation of the normalized mutual information method for registration of CT and MR images in radiotherapy of brain tumors. *Journal of Applied Clinical Medical Physics, 5* (3), 66-79.

Viola P. & Wells W. M. (1995). Alignment by maximization of mutual information. In E. Grimson, S. Shafer, A. Blake, & K. Sugihara (Eds.). *International Conference on Computer Vision*, (16–23). Los Alamitos, CA: IEEE Computer Society Press.

Viola P. & Wells W. M. (1997). Alignment by maximization of mutual information. *International Journal of Computer Vision, 24* (2), 137–154.

Viola P. A. (1995). *Alignment by maximization of mutual information*. Ph.D. thesis. Massachusetts Institute of Technology, Massachusetts, USA.

Wang M. Y., Fitzpatrick J. M. & Maurer Jr. C. R. (1995). Design of fiducials for accurate registration of CT and MR volume images. *Proc. SPIE, Medical Imaging 1995: Image Processing*, 2434, 96–108.

Wells W. M., Viola P., & Kikinis R. (1995). Multi-modal volume registration by maximization of mutual information. In *Medical Robotics and Computer Assisted Surgery* (55–62). New York: John Wiley & Sons.

West J. B., Fitzpatrick J. M., Wang M. Y., Dawant B. M., Maurer, Jr. C. R., Kessler R. M., Maciunas R. J., Barillot C., Lemoine D., Collignon A., Maes F., Suetens P., Vandermeulen D., Van den Elsen P. A., Napel S., Sumanaweera T. S., Harkness B., Hemler P. F., Hill D. L. G., Hawkes D. J., Studholme C., Maintz J. B. A., Viergever M. A., Malandain G., Pennec X., Noz M. E., Maguire, Jr. G. Q., Pollack M., Pelizzari C. A., Robb R. A., Hanson D., & Woods R. P. (1997).

Comparison and evaluation of retrospective intermodality brain image registration techniques. *J. Comput. Assist. Tomography.* 21, 554–566.

Woods R. P., Cherry S. R., & Mazziotta J. C. (1992). Rapid automated algorithm for aligning and reslicing PET images. *Journal of Computer Assisted Tomography, 16* (4), 620–633.

Woods R. P., Grafton S. T., Holmes C. J., Cherry S. R., & Mazziotta J. C. (1998). Automated image registration: I. General methods and intrasubject, intramodality validation. *Journal of Computer Assisted Tomography, 22* (1), 139–152.

Woods R. P., Mazziotta J. C., & Cherry S. R. (1993). MRI-PET registration with automated algorithm. *J. Comput. Assist. Tomography.* 17, 536–546.

Wu J., Kim M., Peters J., Chung H., & Samant S. S. (2009). Evaluation of similarity measures for use in the intensity-based rigid 2D-3D registration for patient positioning in radiotherapy. *Medical Physics, 36* (12), 5391–5403.

Yokoi T., Soma T., Shinohara H., & Matsuda H. (2004). Accuracy and reproducibility of co-registration techniques based on mutual information and normalized mutual information for MRI and SPECT brain images. *Annals of Nuclear Medicine, 18* (8), 659-667.

Zagrodsky V., Shekhar R., & Cornhill J. F. (2000). Mutual information based registration of cardiac ultrasound volumes. In K. M. Hanson (Ed.). *Medical Imaging: Image Processing, vol. 3979 of Proc. SPIE* (1605–1614). Bellingham, WA: SPIE Press.

Zagrodsky V., Shekhar R., & Cornhill J. F. (2001). Multifunction extension of simplex optimization method for mutual information-based registration of, ultrasound volumes. In M. Sonka & K. M. Hanson (Eds.). *Medical Imaging: Image Processing, vol. 4322 of Proc. SPIE* (508–515). Bellingham, WA: SPIE Press.

Zhu Y. M. & Cochoff S. M. (2002). Influence of implementation parameters on registration of MR and SPECT brain images by maximization of mutual information. *Journal of Nuclear Medicine, 43* (2), 160–166.

Zitová B. & Flusser J. (2003). Image registration methods: a survey. *Image and Vision Computing*, 21(11), 977-1000.

Zuk T., Atkins S., & Booth K. (1994). Approaches to registration using 3D surfaces. *Proc. SPIE Medical Imaging: Image Processing*, 2167, 176–187.