

Medical Image Registration Using Knowledge of Adjacency of Anatomical Structures.

Derek LG Hill and David J Hawkes

Image Processing Group, Radiological Sciences,
UMDS, Guy's and St Thomas's Hospitals,
St Thomas's St, London SE1 9RT, UK
D.Hill@umds.ac.uk

Abstract

The combination of multimodality medical images is generating increasing clinical interest, but the absence of robust automatic registration algorithms is an obstacle to these techniques entering routine clinical use. In this paper we present a new algorithm based on surface fitting that makes use of anatomical knowledge of adjacency of identified anatomical structures to solve the 3D rigid body registration problem. The algorithm has been applied to the registration of MR CT and vascular images of the head, and has achieved accuracies similar to those obtained using interactive location of corresponding landmarks.

1 Introduction

Our experience of multimodality medical image combination using interactively identified corresponding point landmarks has demonstrated that there is significant clinical demand for accurately registered images, in particular, the combination of MR, CT and angiographic images for planning neuro and cranial base surgery [1], and nuclear medicine PET or SPECT images with MR or CT images of the head, neck, heart or pelvis. In these applications, the registration transformation relating the images can be adequately described by the six degrees of freedom of a 3D rigid body. A limitation of our own and other existing registration algorithms is that they require considerable user interaction, and the results can be very dependent on the skill and motivation of the user. Several authors have proposed registration algorithms based on surface fitting, but no implementation has been demonstrated to be robust and automatic in clinical use [2-5]. We propose a modification to surface fitting that makes use of anatomical knowledge to enable related but non-equivalent structures identified in different modalities to be fitted. We believe that this modification will be easier to automate than conventional surface fitting, and enables new classes of images to be registered. This paper describes the algorithm and presents preliminary results for two applications.

2 Method

2.1 Fitting Related but Non-equivalent Surfaces

It is not always possible to identify equivalent high-contrast surfaces in images that are to be registered. Two modalities may contain image features correspond-

ing to surfaces of different structures with a known relationship to each other. For example, the inner surface of the skull and the outer surface of the brain are two related but non-equivalent structures. The inner surface of the skull is clearly visible in CT images but is not a distinct feature in MR images, and whilst the outer surface of the brain is clearly visible in MR images, it is hard to delineate from CT images. These structures are not equivalent, because there are membranes, blood vessels and cerebro-spinal fluid (CSF) between the brain surface and the inner surface of the skull, and because the brain surface itself is highly folded. Points on the brain surface, however, lie within the inner surface of the skull. The concept of containment illustrated by this example cannot be used by a conventional surface fitting algorithm, but is frequently used by observers assessing the accuracy of registration: in order to confirm that images are correctly registered, an expert observer will not only check that the equivalent features used for registration (points or surfaces) overlie each other, but they will also check that known anatomical relationships between other structures are obeyed (subject to partial volume and spatial resolution limitations).

We have developed a modification to the conventional surface fitting algorithm that enables this type of information to be incorporated. The distribution of distances separating points on a pair of related but nonequivalent surfaces can be determined using preregistered images. This one dimensional distribution is then used to modify the distance transform, calculated using a 3D chamfer filter [9], to produce a cost image, in which each voxel is labelled with the cost associated with it being occupied by a point on the related surface. This algorithm has been tested on MR, CT and angiographic images of the brain, but is likely to be more generally applicable.

2.1.1 Representation of Anatomical Knowledge

Knowledge about the relationship between anatomical structures can be represented in several ways. A semantic network [6] can describe relationships between objects (eg: adjacency or containment) but provides no quantitative information about the distribution of object locations. An analogic model - such as the point distribution model [7] - can represent variability in the location of points on objects using principal component analysis. In our application, there is no need to describe the entire image, or to handle deformations of structures. Instead, it is necessary to represent the relationship between two surfaces quantitatively, which is a simpler task. For example, one surface of interest may lie inside a second surface, with the distances between points on the surfaces varying over the surface. The precise inter-surface relationship may also vary between individuals, so both the inter- and intra-individual variations in surface location must be represented. As a first approximation, this can be achieved using a one dimensional distribution of occupancy frequency as a function of distance from the surface. A distribution of this type can be acquired from measurements made on images. The occupancy distributions generated from image data will incorporate both anatomical variation and image processing artifacts. For the registration of MR and CT images this distribution was measured from patient images that had been previously registered using our interactive point landmark registration

software [8]. These preregistered images used for acquiring the distance distribution information will be called *reference* images, to distinguish them from individual patient images being registered.

An interactive intensity threshold based region growing tool (part of the software package Analyze[®]) was used to delineate the inner surface of the skull from the reference CT images. The outer surface of the brain was interactively identified from the reference MR images using our own software. The CT surface was represented as a 3D binary image in which voxels within the inner surface of the skull (object voxels) had the value 1, and all other voxels had the value 0. The MR derived brain surface was represented as a non-connected list of 3D points. The CT surface was transformed using a 3D chamfer filter (modified to work on images with non-cubic voxels) to generate a distance image in which all voxels were labelled with their distance from the inner surface of the skull. These distances were negative within the skull, and positive outside it. A histogram of distances between points on the brain surface and the inner surface of the skull was then generated by reading the voxel value from the distance image corresponding to each 3D point in the MR surface representation.

2.1.2 Use of Anatomical Knowledge for Registration of MR and CT Images of the Head

For individual patient images to be registered, the inner surface of the skull is delineated from the CT image, and the outer surface of the brain from the MR images, as described above. The distance transform from the inner surface of the skull is then computed using the 3D chamfer filter.

The distance distribution information acquired from the reference images is combined with the distance image to generate a cost image. Each voxel in the distance image has a label that represents its distance from the inner surface of the skull. This distance is negative inside the inner surface (within the skull cavity), and positive outside this volume. The distance distribution histogram acquired from the reference images relates these distances to the probability of voxels at a given distance from the inner surface of the skull being occupied by a brain surface voxel. The distance histogram is combined with the distance image in the following way. First, the distance value corresponding to the peak of the distance histogram is subtracted from all voxels in the distance image. Second, the shape of the histogram (now modified so that its peak lies at zero distance) is used to modify the distance gradients in the cost image. The higher the occupancy frequency, the greater the gradient in the cost image. The occupancy frequency values were thresholded into three bands "high", "medium" and "low", with gradients 2, 1 and 0.5 respectively. These values were incorporated into the cost image by multiplying the distances by these gradients, and adding offsets to avoid discontinuities.

2.1.3 Testing the Algorithm

The algorithm was tested in two ways for this application. Firstly, using registered images, known misregistrations were applied sequentially in each of the degrees

of freedom individually. Ideally, for each degree of freedom, the cost should be a well behaved function of misregistration, with the lowest cost solution arising when the images are correctly registered, and cost increasing monotonically with misregistration. The performance of conventional chamfer matching was compared with the algorithm incorporating anatomical knowledge to establish whether the use of anatomical knowledge leads to a more well behaved optimisation space. Secondly, the algorithm was used to register images for which the registration transformation had already been determined using point landmark registration software [8]. The solutions provided by the two algorithms were thus compared.

2.2 Registration of Blood Vessel Centrelines to Anatomical Surfaces

There is, in fact, no reason why the non-equivalent structures being fitted both need to be surfaces. The chamfer matching algorithm requires that one structure is represented as a continuous surface, but the other structure is represented as a list of discrete points. In a conventional chamfer match, these discrete points represent a surface, but they could equally well represent line segments, provided that there is a sufficient number of line segments, or the segments are sufficiently tortuous to constrain all the degrees of freedom of the registration transformation. An important clinical application of this technique is the registration of blood vessels reconstructed from bi-plane digital subtraction angiography (DSA) with MR images. DSA remains the modality of choice for many vascular imaging applications, but the angiograms obtained are projection images, and are frequently difficult to relate to structures of interest visible in tomographic modalities from the same patient. Several algorithms have been proposed for reconstructing three dimensional vascular networks from these bi-plane or stereo projections [10,11]. Such reconstructions comprise lists of blood vessel centre line coordinates with vessel radii. These representations are 3D, but only contain vascular information, and relating them to tomographic images remains difficult. The anatomical relationship between blood vessels that are large enough to be seen using DSA (greater than 0.5mm in diameter) and the surfaces of anatomical structures is often clearly defined. In the brain, for example, the larger cerebral arteries lie on the surface of the cerebrum, held in place both by small vessel branches that penetrate the surface and by the *pia mater* membrane. These arteries follow the foldings of the sulci and gyri of the cortex. Before CT and MR imaging became widely available, radiologists were able to make use of this close coupling of vessels to brain to infer anatomical information about non-vascular structure from cerebral angiograms alone [12]. Similarly, in the heart, the principal coronary arteries do not penetrate the myocardium, but are anchored to it by small vessel branches that do penetrate the myocardium. The method for registering blood vessel centre lines derived from DSA with surfaces derived from tomographic modalities is thus similar to that described in above. It was evaluated by registering segments of the major cerebral arteries to the brain surface. Registered reference data was unavailable in this application, so could not be used to measure the distribution of distances from blood vessel centreline points to the corresponding surface. This distribution was therefore inferred from the "text book" style ana-

tomical knowledge described above. Because the major cerebral arteries lie on the brain surface, their vessel centrelines should lie at a distance equal to the vessel radius from this surface. The cost image was generated from the distance image by adding a distance corresponding to the relevant vessel's radius.

2.2.1 Testing the algorithm

Data was acquired from a phantom made from a cadaver brain. The vessels of the brain were perfused with a gel that gives high signal from spin-echo MR images. Coronal T₁ weighted MR images were acquired with a slice thickness of 2mm. The surface of the brain was delineated using an interactive segmentation tool, and transformed using the chamfer filter. The centreline coordinates of the right pericallosal artery and right middle cerebral artery were selected on the same MR images, and stored as a list of 3D point coordinates. These large cerebral arteries have a diameter of approximately 2mm, so the cost image was generated from the distance image by adding 1mm to all distances. The effect of misregistration on the cost solution was simulated by applying a range of transformations individually for each degree of freedom.

2.3 Distance transforms and influence zones

An implicit assumption of surface fitting algorithms is that the transformation that generates the best surface fit corresponds to the correct registration transformation. For some surfaces, this assumption is invalid, as two very distant portions of equivalent or related surface may have very similar curvature. This problem could be overcome by labelling surface patches with an anatomical label. A high cost could then be assigned to solutions where equivalent patches were not overlying each other. This can be achieved using influence zones. Surface voxels can be given a label as well as being identified as lying on the surface. The label can identify that surface voxel as belonging to a certain surface patch. For example, brain surface voxels might be labelled with the name of the brain lobe, or even the name of the sulcus or gyrus that they lie on. When the distance transform is generated, these labels can be propagated through the distance image producing influence zones, which are volumes of the distance image containing voxels that are closer to that surface patch than to any other. The points being fitted to distance image can likewise have a label assigned to them.

2.4 Optimisation

We use the combination of multiple resolutions (which blur out many of the local minima) and stochastic optimisation algorithms (which tend to be more robust at finding the global minimum than conventional down hill methods) to overcome the problem of a large search space containing many local minima. The optimisation proceeds first at a low resolution with a large search space, then at progressively higher resolutions and smaller search spaces. At each resolution, a genetic algorithm [13] is used to determine the global minimum.

3 Results

3.1 Registration of MR and CT images of the head

Figure 1 shows the distribution of distances of points on the brain surface from the inner surface of the skull for a single set of reference MR and CT images. The peak of this distribution corresponds to the most likely location for points on the brain surface to lie (-0.7mm). For the purpose of incorporating the information in this distribution into the distance transform from an individual patient, this histogram was split into four regions: distances greater than 0 = low occupancy frequency; 0 to -2mm = high occupancy frequency; -2mm to -5mm = medium occupancy frequency; less than -5mm = low occupancy frequency.

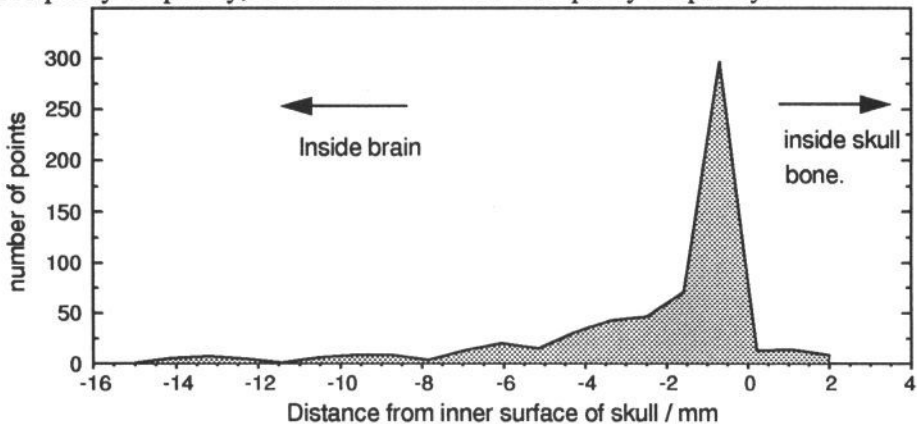


Fig 1. Distribution of distances of points on the brain surface with respect to the inner surface of the skull measured on registered MR and CT images for one patient.

Figures 2 and 3 demonstrate the effect on the behaviour of the cost function resulting from incorporating the information derived from this distribution into the distance transforms. Figure 2 demonstrates that using the conventional distance image, the minimum cost solutions for translation in the z direction (cranio-caudal) and the rotation about the x axis (a line passing through both auditor meati) are different from those found using the point landmark based registration. Figure 3, however, shows that by combining the distance image with the information derived from the distance distribution histogram, the cost gradients are steeper and the lowest cost solution corresponds to the point landmark solution for all degrees of freedom.

For two patient studies evaluated so far (one of which was the same study as that used to train the algorithm), the genetic optimisation algorithm converged to a solution within 2mm translation and 2 degrees rotation of the solution found using point landmark registration.

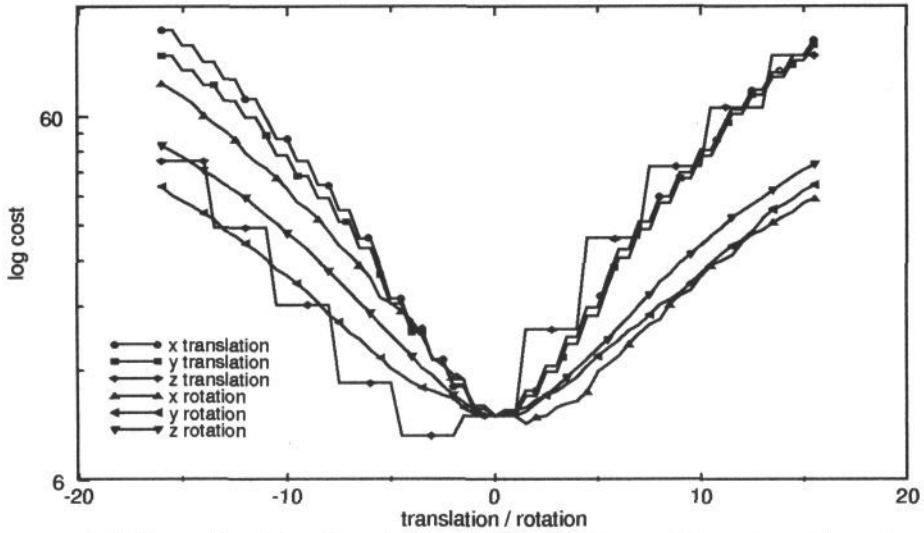


Fig 2. Conventional chamfer matching of MR and CT images. Change in cost (square of chamfer distance) resulting from translation or rotation in six degrees of freedom.

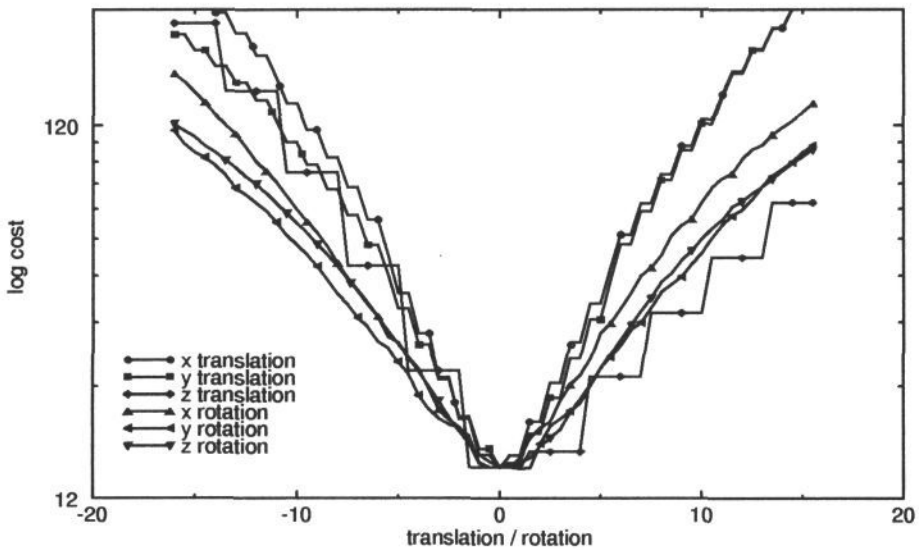


Fig 3. Chamfer matching of MR and CT images incorporating anatomical knowledge. Change in cost resulting from translation or rotations in six degrees of freedom.

3.2 Registration of MR and Vessels in the Head.

On the left of figure 4 is an example slice from a coronal MR dataset from a cadaver brain. Adjacent to it are the distance transform of the same coronal slice and the distance transform modified to incorporate the anatomical knowledge about blood vessels running along the surface of the brain.



Fig. 4. An example slice from a coronal MR dataset of a cadaver brain (left), its distance transform (centre), and modified distance transform (right).

Registration using the modified distance transform in figure 4 was evaluated using blood vessel data derived from the same MR images. Two vessel segments were used: a branch of the middle cerebral artery passing through the sylvian fissure, and the pericallosal branch of the anterior cerebral artery. Figure 5 demonstrates that the best registration transformation for these two vessels lies within a 2mm translation and a 4 degrees rotation of the known solution.

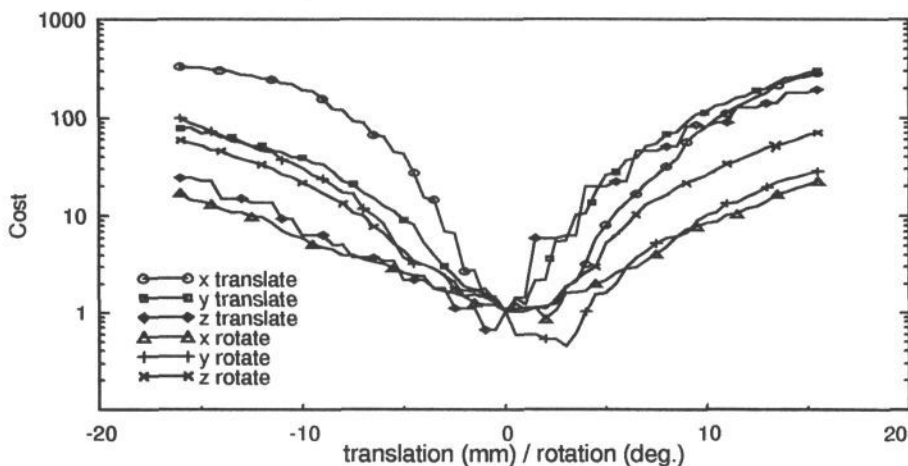


Fig. 5. Modified chamfer matching of brain surface and vessels. The minimum cost solution for all degrees of freedom is within 2mm or 4 degrees of the known solution.

4 Discussion and Conclusions

The algorithm presented in this paper potentially provides a means of automating the registration process for many applications. Anatomical knowledge combined with the use of registered reference images provides a means of customising the algorithm for particular clinical applications. We have shown how this can be done with MR, CT and angiographic images.

Different registration algorithms place different demands on pre-segmentation of the image data. In our algorithm, well defined surface structures need to be

identified from MR and CT images. Although we have done this segmentation interactively for the test data presented in this paper, automatic segmentation algorithms have been proposed for delineating these structures [eg 14]. The segmentation task is different from that required for conventional surface fitting because it is not necessary to delineate equivalent structures in the modalities being registered, and because the uncertainty in the surface delineation can be coded into the modified distance transforms. The algorithm requires the additional information on adjacency and proximity of identified structures, which in this paper was derived from registered reference images, together with very basic anatomical knowledge of the kind "cerebral blood vessels greater than 0.5mm diameter are constrained by the *pia-mater* membrane to lie on the outside of the cerebral cortex". The work presented here has used only a single influence zone. The incorporation of multiple influence zones into the distance transform would allow for the use of multiple associated features consisting of point distributions, lines or surfaces. The incorporation of anatomical knowledge into the cost function is currently very simplistic. The use of multiple influence zones and occupancy functions derived from many reference datasets would rigorous.

The registration of MR images and 3D reconstructions of vascular networks in the brain up until now has remained an unsolved problem both because of the difficulty in reconstructing the complete cerebral circulation from bi-plane angiograms, and because of the absence of equivalent structures in the two modalities that can be used for registration. We have demonstrated that, given a brain surface, and a 3D reconstruction of a small number of major cerebral artery segments, the modified chamfer matching algorithm can determine the registration transformation relating these datasets. The blood vessel segments used for this work were derived from the MR images so that the correct transformation was known a priori. It is possible to reconstruct many of the major vessels of the cerebral circulation after the user has interactively identified the vessels in bi-plane angiograms provided the x-ray projection geometry is known [10], but it is not currently feasible to reconstruct a significant proportion of the cerebral circulation [11]. If, however, a small number of vessel segments were manually reconstructed from the angiograms, these could be used for registration of those vessels to MR coordinates. This would establish the relationship between X-ray and MR coordinate systems and could provide an additional constraint for reconstruction of the vascular data.

For all non-linear optimisation problems, an algorithm that successfully avoids local minima in optimisation space is essential if the results are to be reliable. The genetic algorithm appears to be a rapid and robust method for finding an approximately correct solution (somewhere near the bottom of the global minimum, as opposed to a distant local minimum), provided that the population size is sufficiently large, but appears to be computationally expensive as a means of finding the very bottom of the global minimum. One way of speeding up the optimisation process would be to provide good starting estimates of the required registration transformation (eg: by means of interactive identification of a small number of equivalent points in the modalities being registered) thus constraining the search

space. As currently implemented, and depending on the precise parameters used, the algorithm typically took about one hour to run on a Sun SparcStation 10.

5 Acknowledgements

This work was funded by the SERC (project SMIRC).

6 References

- [1] D.L.G. Hill, S.E.M. Green, J.E. Crossman, et al: Visualisation of multi-modal images for the planning of skull base surgery. SPIE 1808, 564-573 (1992)
- [2] C.A. Pelizzari, G.T.Y. Chen, D.R. Spelbring, R.R. Weichselbaum, C-T. Chen: Accurate three dimensional registration of CT, PET and/or MR images of the brain. *Journal of Computer Assisted Tomography* 13, 20-26 (1989)
- [3] P.A. Van den Elsen, J.B. Maintz, E.D. Pol, M.A. Viergever: Image fusion using geometrical features. SPIE 1808, 172-186 (1992)
- [4] H. Jiang, R.A. Robb, K.S. Holton: New approach to 3-D registration of multi-modality medical images by surface matching. SPIE 1808, 196-213 (1992)
- [5] T.L. Faber, R.W. McColl, R.M. Opperman, J.R. Corbett, R.M. Peshock: Spatial and Temporal Registration of Cardiac SPECT and MR Images: Methods and Evaluation. *Radiology* 179, 857-861 (1991)
- [6] G.P. Robinson, A.C.F. Colchester, L.D. Griffin: Model-based recognition of anatomical objects from medical images. in *Information Processing in Medical Imaging '93*. Springer-Verlag, ed. HH Barrett 1993 pp197-211
- [7] T.F. Cootes, C.J. Taylor, D.H. Cooper, J. Graham: Training models of shape from sets of examples. *Proc. BMVC'92*, Springer-Verlag pp54-61
- [8] D.L.G. Hill, D.J. Hawkes, J.E. Crossman, et al: Registration of MR and CT images for skull base surgery using point-like anatomical features. *British Journal of Radiology* 64, 1030-1035 (1991)
- [9] G. Borgefors: Distance Transformations in Arbitrary Dimensions. *Computer Vision, Graphics and Image Processing* 27, 321-345 (1984)
- [10] D.J. Hawkes, A.C.F. Colchester, C.R. Mol: The accurate 3D reconstruction of the geometric configuration of vascular trees from X-ray recordings. In: R. Guzzardi (ed.): *Physics and engineering of medical imaging*. Holland: Martinus Nijhoff 1987, pp. 250-256
- [11] C. Smets, G. Verbeeck, P. Suetens, A. Oosterlinck: A knowledge based system for the delineation of blood vessels on subtraction angiograms. *Patt Recog Lett* 8, 113-121 (1988)
- [12] G. Szikla, G. Bouvier, T. Hori, V. Petrov: *Angiography of the human brain cortex*. Springer-Verlag, New York 1977
- [13] D.E. Goldberg: *Genetic algorithms in search optimisation and machine learning*. Mass. USA: Addison Wesley 1989
- [14] G. Gerig, J. Martin, R. Kikinis, O. Kuebler, M. Shenton, F.A. Jolesz: Automating segmentation of dual-echo MR head data. In: A.C.F. Colchester, D.J. Hawkes (eds.): *Information Processing in Medical Imaging*. Heidelberg: Springer-Verlag 1991, pp. 175-187