An Algorithm for 3D Registration of Multi-Modality Medical Images by Use of Voxel-Similarity-Based Measures

Roland Göcke
University of Rostock
rgoecke@informatik.uni-rostock.de

0. Abstract

This report describes work done within a practical training at the Philips Research Laboratories / Technical Systems Hamburg in 1995. My work was concerned with the development and implementation of a completely automatic image registration algorithm that uses a voxel-similarity-based measure as matching criterion. Both single- and multi-modality 3D medical images can be processed. Furthermore, the algorithm handles both full and truncated dataset images.

1. Introduction

In medicine, patients are often imaged with more than one tomographic radiological imaging modality (e.g. CT, MRI) for the purpose of improved diagnosis or treatment planning (e.g. operation, radiation therapy dose). Medical diagnosis benefits from the complementarity of the information in images of different modalities. It can be difficult for a clinician to mentally combine all the image information accurately because of variations in patients orientation or differences in resolution or contrast of the images. Therefore, an image registration technique is sought that transfers all the image information into a common coordinate system. The aim is to present the images in a way that makes it easier for the clinician to fuse the image information to find similarities and differences. The registration process is prior to any further image processing like segmentation of specific organs (e.g. blood vessels).

Rigid-body registration of 3D medical images is used to align two 3D scans of a patient that were taken at different times with the same modality (single-modality) or generated on different medical imgaing modalities (multi-modality). In addition, the algorithm could be used to align 3D scans of two different patients, possibly images of different modalities, to detect organic similarities. The algorithm is based on the assumption that the rigid bone of the human skull restricts the necessary transformation to the six degrees of freedom of a rigid body, i.e. three translations and three rotations.

Voxel-similarity-based matching criteria are measures of misregistration that are functions of the attributes (e.g. grey-values, gradients, textures) of all pairs of common voxels from the images to be registered at a given position of misregistration. Those criteria are based on the observation that the 2D scatter plot (or histogram) of the corresponding voxel pairs of a pair of registered images will be dispersed by misregistration. Such approaches purely based on voxel similarity can be fully automated.

Two voxel-similarity-based measures, proposed by A.Collignon et al [1], [2], were investigated as matching criteria, the entropy and the relative entropy of the joint probability distribution of the grey-values of all common voxel pairs in two images. Studholme et al [3] examined different measures and confirmed that the relative entropy as matching criterion is robust and reliable even for truncated dataset images.

In Section 2 the algorithm used is described. Section 3 shows results obtained for some test examples. Finally, in Section 4 some problems of the method are discussed.

2. Algorithm

2.1. Preliminaries

A right-hand coordinate system is assumed where the X axis is the patient left to right, the Y axis front to back and the Z axis head to feet. The images are displayed slice by slice where the x and y coordinates at a particular position on the Z axis are taken as one slice. A voxel is generally identified by integer indices. The intensity values within the images are stored as integers, as well. It is assumed that two such 3D images are given.

The coordinate systems of the images are related and the aim of the registration algorithm is to calculate the necessary transformation parameters to register the images. There exist two ways of representing the rotation parameters, either by Euclidean angles or by a vector (rotation by quaternions). The translation parameters are represented by voxels in both cases. It might be necessary to apply a scaling matrix if the images have different dimensions because of different voxel sizes or different resolutions. The elements of the diagonal of such a scaling matrix define the number of voxels per unit length and are assumed to be known.

By optimization of the voxel-similarity-based matching criterion, estimates for the transformation parameters can be derived by the algorithm described in the following.

2.2. The Algorithm - General Concept

Both matching criteria, proposed by A.Collignon et al [1], [2], i.e. the entropy and the relative entropy of the joint probability distribution of the combined grey-values of all common voxel pairs in two images, are based on information theory and the general observation that misregistration leads to a diffusion of the grey-value distribution.

In general, the entropy of a random variable is a measure of the information required on average to describe it. In my work, the entropy of the grey-value distribution defined as

$$E(X;Y) = -\sum_{x,y} p(x,y)\log(p(x,y))$$

is used where X and Y are two random variables with a joint probability distribution $p(x,y)$. Diffusion of the grey-value distribution leads to an increase in the information required to describe it. Thus, the entropy is minimal if the images are registered. However, prior experiments with a 2D implementation showed that the entropy function is only uni-modal for a small range of unregistered positions around the correct registered position (see Fig.1). If the initial guess of the transformation parameters is too far away from the true values, the algorithm will not find the correct solution but end in a local minimum or never find a solution at all.

Hier Fig.1. einfügen.

Fig.1. The leftmost image shows a slice of the reference MRI image, the following images show the dependence of the function graph of the entropy function on a shift in x or y direction, and a rotation about the z axis, respectively. The origin of the coordinate system is in the topleft corner.

Therefore, further work was concentrated on the relative entropy of the joint probability distribution of all common voxel pairs as matching criterion. The relative entropy, also known as mutual information, measures the degree of mutual dependence of one variable on another one. This measure is based on the shared information (or relative entropy) between the overlapping regions in the images. Here the mathematical definition of the relative entropy

$$S(X;Y) = \sum_{x,y} p(x,y) \log\left(\frac{p(x,y)}{p(x)p(y)}\right)$$

otherwise known in information theory as Shannon information of two random variables X and Y with marginal and joint probability distributions $p(x)$, $p(y)$ and $p(x,y)$, respectively. The function is symmetric and non-negative. It is bound above by one-to-one mappings, i.e. complete mutual dependence, and bound below by complete independence of the random variables which are the image intensity values in this case (see Fig.2). Thus, the maximization of the relative entropy is the aim of the registration algorithm.

Due to the finiteness of the voxel indices and the continuity of the transformation parameters, some form of interpolation is needed to determine corresponding voxel pairs correctly. A simple nearest neighbour interpolation was used because of its computational efficiency. Test results showed a sufficient accuracy (see Sections 3. and 4.).

Hier Fig.2. einfügen.

Fig.2. As in Fig.1. but now for the relative entropy function.

The algorithm might be summed up as follows:

Step 1 : Evaluate matching criterion at current position of misregistration
Step 2 : Optimize matching criterion
Step 3 : Calculate corresponding transformation parameters
Step 4 : Images registered?
If yes, Stop, else go back to Step 1.

2.3. Optimization

The general aim for the optimization process is to keep the number of function evaluations as small as possible in order to register the images as fast as possible. Several optimization techniques were investigated for their usability, like separate optimization of each transformation parameters, optimization by Golden Section, optimization by bisection, and optimization by a gradient descent algorithm. Unfortunately, the maximum of the relative entropy function is a tall but narrow peak and, hence, the function is not differentiable at the maximum. Therefore, many faster optimization methods will not yield the correct result or even work at all. On the other hand, the change of the gradient towards the maximum is relatively small if the current position is far away from the registered position. That means that progress in the objective function is small in that case.

The gradient descent method performed best in tests under these circumstances. After determining the gradient, a Golden Section search is started to obtain the optimal step length.

3. Results

For the test examples in this section two real CT and MRI images of a human head were used. The images were transformed by a rotation about (1, -2, 10) degrees and a translation of (5, 4, 2) voxels. Then the transformed and original images were used to recalculate the transformation parameters by the registration algorithm given above (using the relative entropy as matching criterion). The results are shown in Tab.1. The initial guess of the transformation parameters was the zero vector. It would be very helpful and probably speed up the algorithm if the user could supplement the algorithm with any better initial guess.

Experiment	Size / type of image	Accuracy level	Estimates for rotations	Estimates for translations	No. of steps
1	256*256*120 MR	0.1	1.002046 -2.002490 9.997544	4.979126 4.052769 1.509748	6
2	320*320*87 CT	0.1	1.008468 -2.029081 10.001192	4.166164 4.699414 1.509748	6
3	256*256*120 MR	0.01	0.999562 -2.007441 10.001244	4.109394 4.041059 1.410396	7
4	320*320*87 CT	0.01	1.009090 -2.005226 9.995352	4.181391 4.709747 1.606722	8

Tab.1. Parameters and results of test eaxmples

The experiments show that sub-voxel accuracy can be reached. They were performed on a SUN 4/80 Sparcstation 10 (60 MHz SuperSparcCPU, 192 MB RAM). In experiments 1 and 2 the deviation from the correct transformation parameters was of order 0.03 degrees for the rotation parameters and 0.95 voxels for the translation parameters. Generally, the execution time of the whole registration process of images like those used in experiments 1 and 2 with an accuracy level of 0.1 was approximately between 1.5 and 2 hours.

Experiments 3 and 4 were performed with a higher accuracy level of 0.01. Small improvements in most rotation parameters

Hier Fig.3. einfügen.

Fig.3. The left two images show a particular slice in the reference and retransformed CT image, respectively. The third image visualizes the difference between original and retransformed image at an intermediate stage. The last image shows this difference after finishing the registration process.

were obtained but on the cost of a distinctly higher calculation time. Even a slight deterioration for the translation parameters (except for one) could be noticed but the error was still less than the voxel size. In these experiments the deviation from the true transformation parameters was of order 0.009 degrees for the rotation parameters and 0.96 voxels for the translation parameters. Here, the execution time was approximately between 2 and 3 hours.

4. Discussion of Method and Problems

The algorithm is completely automatic and does not need any user interactivity or pre-segmentation. It works for both complete and truncated images. Both CT and MRI single-modality images can be processed as well as multi-modality images. The algorithm is based on the original grey-values only and, hence, more robust than surface-based registration algorithms. For a large range of misregistered positions no differences between the original and retransformed images could be seen by visual inspection.

The disadvantage is the high execution time which is due to the nature of the voxel-similarity-based measure. As for any voxel-similarity-based measure, the realtive entropy takes account of all (!) voxels in the images. Such measures are, therefore, less prone to image noise, image truncation, image artefact and small tissue deformations than surface-based methods. On the other hand, it means that the complete amount of data of both images must be checked at each evaluation of the objective function.

Two approaches seem to improve the speed. Firstly, one might try to find a faster optimization method that needs fewer function evaluations. Unfortunately, many faster methods cannot be used because of the function's shape (see Section 2.3. Optimization). Secondly, a multi-resolution approach might be possible to reduce the amount of data from the images (e.g. if the size of an image is halved in each dimension, the amount of data processed is reduced by a factor of eight).

Occasionally, problems can occur if the images contain very severe noise. The more noise is on the images, the flatter is the graph of the relative entropy as objective function, i.e. progress in the optimization process is very slow and the algorithm may stop at a local minimum and, hence, lead to an incorrect solution. In this case, filter or thresholds should be used before registration to reduce the influence of noise.

In the current implementation a simple nearest neighbour interpolation was used and test examples showed that sub-voxel accuracy can be reached. Nevertheless, a more sophisticated interpolation method like trilinear interpolation could presumably improve the accuracy of the solution. Generally, a trade-off between speed and accuracy has to be found.

5. References

[1] Collignon A., Vandermeulen D., Suetens P., Marchal G.: 3D Multi-Modality Medical Image Registration Using Feature Space Clustering. CVRMed'95, 3-6 April, 1995, Nice (France)

[2] Collignon A., Maes F., Delaere D., Vandermeulen D., Suetens P., Marchal G.: Automated Multi-Modality Image Registration based on Information Theory. IPMI'95, 26-30 June, 1995, Ile de Berder (France)

[3] Studholme C., Hill D.L.G., Hawkes D.J.: Automated 3D Registration of Truncated MR and CT Images of the Head. Proceedings of the British Machine Vision Conference, Eds. David Pycock, Sept. 1995, vol. 1, pp 27-36