OPTIMISATION STRATEGIES APPLIED TO GLOBAL SIMILARITY BASED IMAGE REGISTRATION METHODS

Martin Čapek
Institute of Physiology, Academy of Sciences of CR,
Vídeňská 1083, Praha 4, 142 20, Czech Republic
e-mail: capek@biomed.cas.cz

ABSTRACT

This paper investigates the best possible optimisation strategy for determining global extremes of objective functions in the framework of image registration. We register serial optical slices captured by a confocal laser-scanning microscope (CLSM). We experiment with following objective functions: the Sum of Absolute Valued Differences (SAVD), Normalised Correlation Coefficient (NCC) and Mutual Information (MI). Full search of the parametric space for the global extreme is tedious, therefore, we study following optimisation strategies: the genetic algorithm, downhill simplex method, Powell's method, and simulated annealing, all with the use of multiresolution approach to decrease the computational load.

Keywords: image registration, optimisation strategy, global similarity evaluation method, serial optical slices.

INTRODUCTION

We understand image registration as determination of coefficients of a rigid transformation (translations and rotations) between two images accomplishing alignment of their pixel matrices. This definition can be expressed by the following equation,

$$T' = arg \min_{T} F(\, u(.), v(T(.))\,), \qquad (1)$$

where T', T stand for geometrical transformations; F for a global similarity evaluation function; $u(.)$ for a pixel of a reference image; $v(T(.))$ for a pixel of an image to be registered which is related to the corresponding pixel of the reference image u by transformation T. As most of optimisation techniques deal with minimum of functions, *min* is put in Eq. 1.

We applied registration to images captured by CLSM attempting to render and visualise in 3D a large biological tissue specimen. We chose global similarity based image registration methods that exploit information contained in whole images not requiring segmentation, because we aim at an automatic procedure. We experimented successfully [Čapek98] with following global similarity evaluation methods: the Sum of Absolute Valued Differences (SAVD), Normalised Correlation Coefficient (NCC) and Mutual Information (MI), see APPENDIX. Image registration requires finding the global extreme of these functions over the space of parameters of geometrical transformation by the use of optimisation methods. In this paper we compare the full search strategy with the genetic algorithm, downhill simplex method, Powell's method, and simulated annealing with regard to their convergence speed and their reliability to find the global extreme. All strategies were applied to images with using multiresolution approach to decrease the computational load.

DATA

We register peripheral optical slices of physical stacks of biological tissue. These slices, however, may be mechanically deformed due to their physical sectioning. Moreover, the distance between consecutive peripheral optical slices is unknown. Thus the slices may not match perfectly. Optical slices are grey-scale 512x512 pictures (see Fig. 1). Studied misregistrations of two optical slices to be registered consist of a rotation and translations along both x and y axis restricted to approx. 100 pixels from the registered position.

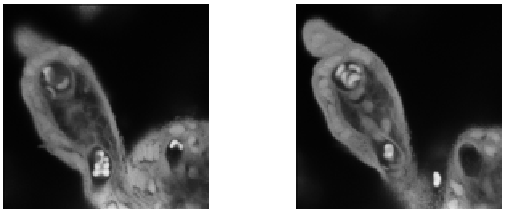

Fig. 1: Optical slices of biological tissue (terminal villus of human placenta) captured by CLSM (the distance between slices is 4 μm).

Our practical experiments with the similarity evaluation methods shown that SAVD is the fastest one, however, NCC is more robust, though it is more computationally demanding. The most robust but also far more time-consuming approach (see Table 1) is represented by MI which is rather an ″overkill″ for registration of CLSM images, but we pay attention to this criterion because of its suitability to register inter-modality images [Maes97].

Function	SAVD	NCC	MI
Time [s]	32,26	44,56	306,65

Table 1: CPU time elapsed for generation of graphs in Fig. 2 representing similarity methods regarding misregistration of images in Fig. 1. The images were re-sampled to 64x64 pixels. A portion of the above times were spent on computation of the input pictures.

Optimisation techniques are used to find global extremes of the objective functions. From the point of view of the optimisation strategies applied to similarity evaluation methods as the objective functions, the broader the extreme of these functions, the easier is to localise it in a n-dimensional plane. On the other hand, the sharper the extreme, the more precise registration is achievable. Fig. 2 shows graphs of extremes of three similarity functions for translation misregistration in range of ±32 pixels around the registered position of optical slices in Fig. 1 re-sampled to 64x64 pixels. Both SAVD and NCC have broader extremes when compared with that of MI which is sharp and narrow. For the sake of a uniform view of all graphs in Fig. 2, both SAVD and NCC are depicted as multiplied by –1 because of their positive valued extremes.

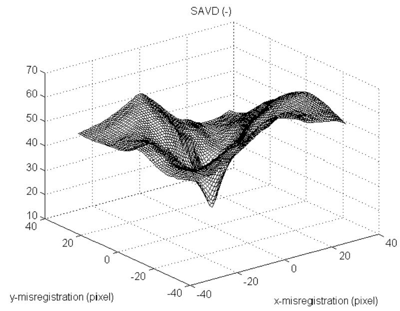

A)

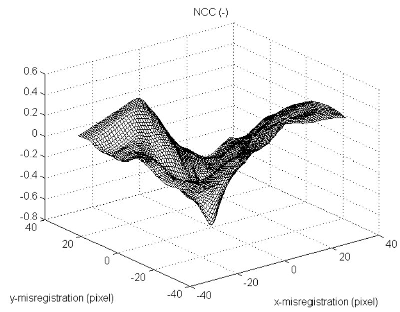

B)

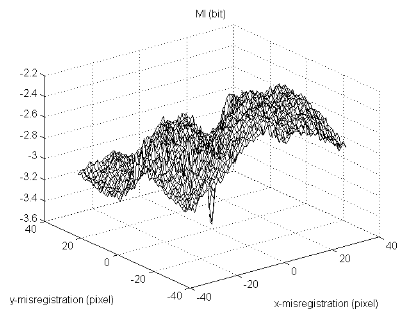

C)

Fig. 2: Graphs of different similarity evaluation methods for translation misregistration of ±32 pixels: A) SAVD, B) NCC, C) MI.

METHODS AND RESULTS

The below described optimisation strategies, i.e. full search, the genetic algorithm, amoeba, Powell's method and simulated annealing, were applied to optical slices and the objective functions were SAVD, NCC and MI, respectively. These strategies were evaluated according to their ability to find the global extreme and their convergence speed (i.e. the number of objective function evaluations necessary for finding the global extreme). All methods were implemented in C language and ran on PC (MMX 266MHz, 96MB RAM).

To speed up all the calculations we used **multiresolution pyramidal approach** [Burt84], introducing picture pyramids with original pictures (e.g. 512x512) at their bottoms and pictures of lower resolutions (e.g. 256x256, 128x128, 64x64,…) at their higher levels. Image registration starts at the highest level that has minimal number of pixels and proceeds down to the original pictures. Optimal transformation parameters are searched only in the neighbourhood of the extreme found at the previous (higher) level. If the number of evaluated geometrical transformation parameters is low (up to three), this technique maintains an acceptable computational load. In case of the multiresolution approach, it is vitally important to detect the global extreme on the highest level as a prerequisite of correct registration of the original images. Therefore, the optimisation methods were applied to images in Fig. 1 resampled to 64x64 pixels, which is the size of the smallest image matrices in the pyramids. The images are registered in the position depicted, however, they are not identical but only similar in contents. The searches were restricted to ±12 pixels, and a general rotation was considered.

The simplest optimisation technique is the **full search** [Tekal95] requiring search space division into discrete steps and evaluation of the objective function over all this set of these discrete parameters. This technique is time-consuming and computational time increases significantly with the number of parameters of geometrical transformation. Convergence to global extremes is quite sure for this strategy, however, it requires the highest number of function evaluations (t_{SAVD}=1236,56 s; t_{NCC}=1438,52 s; t_{MI}=8223,14 s; N=112500 f.e. /function evaluations/). In practice, full search enables sampling of the parameter space only with predefined steps of the parameters, which does not allow sub-pixel accuracy. This would demand very small steps increasing highly the computational burden.

The **genetic algorithm** [Parke97] is a stochastic optimisation algorithm that represents possible solutions to a problem by position of individuals in the n-dimensional search space, and this position is expressed by a bit string called a chromosome. The chromosome comprises a concatenation of genes, each of which is a binary representation of one of the n parameters of the search space. The search space is defined for each degree of freedom, and an initial estimate of the solution is provided. The starting population of possible solutions includes the starting estimate and other random individuals generated within the search space. The objective function specified for the optimisation problem is used to assign each member of the population with a fitness value. The optimum solution is found by reproduction of the individuals in the population. Reproduction can be asexual (a new individual is generated with the same genetic material as one of its parents) or sexual (in which case the offspring's chromosome is derived from the genetic material of both its parents). The probability of sexual reproduction is predefined and is normally greater than 50%. In either case there is a small probability that cosmic radiation will cause a mutation, thus introducing new genetic material into the population. Cosmic radiation is implemented by random inversion of randomly selected bit of the chromosome.

The genetic algorithm does not guarantee to find global functional optima because the length of binary strings of the individuals can substantially reduce the solution accuracy and the search process does not ergodically cover the parametrical space. Therefore, the accuracy of the global extreme determination is restricted. We used real translation parameters encoded as binary strings by 6 bits and real rotation parameter encoded by 9 bits. The genetic algorithm was easily able to find global extremes when used SAVD and NCC (population size=100, number of generations=5, preservation of the last generation=5%; t_{SAVD}=6,49 s, t_{NCC}=7,75 s, N=500 f.e.). However, the algorithm can take much more additional generations to get very close to extreme peaks (even 25 generations extra), because it does not include any internal force driving for the extreme. The extreme of MI is much narrower when compared with the ones of SAVD and NCC, hence the algorithm required about 10 generations to find it (t_{MI}=76.76 s, N=1000 f.e.) and about 12 generations extra to determine it more precisely.

Another approach is ″**amoeba**″ or downhill simplex method [Press92]. An amoeba is a shape with n+1 vertices in n dimensional space. For the purposes of n dimensional optimisation, the amoeba can be considered as a creature with n+1 legs, which tries to crawl into the global minimum. Its legs should be spread around the global minimum. The objective function is evaluated for each leg, and the creature moves its worst leg to a better position. Because the amoeba is the simplest creature that can exist in n-dimensional space, all of its remaining legs lie in a n-1 dimensional plane. The creature first tries to reflect the worst placed leg across this plane. If its new position is better than the previous one, it proceeds in that direction. Alternatively, it places the leg in an intermediate position. If none of these moves result in an improvement, it moves all of its legs closer to the position of the best one. It again selects its worst leg, and repeats the described cycle. The amoeba stops when all its legs are within a predefined distance of each other.

Powell's method [Press95] in n dimensions is based on a sequence of one-dimensional minimisations along a set of directions represented by unit vectors

e_1, e_2,…e_n, which starts from a predefined point. It moves along the first direction to its minimum using a line minimisation routine, then the second direction is minimised, and so on, cycling through the whole set of directions as many times as necessary, until the objective function stops decreasing. This method consists of prescriptions for updating the set of directions avoiding endless cycling through this set of directions. This is ensured by the following condition: during the movement along an actual direction, the gradient of the objective function must stay perpendicular to the previous direction. If this condition holds pairwise for all members of the set of directions, they are said to be a conjugate set. There is no need to redo any of those directions during successive line minimisation of the objective function along the conjugate set of linearly independent directions, because minimisation along one vector does not spoil minimisation along previous one.

In practice, amoeba alone and Powell's method gave very often only local extremes. Therefore, a combination of the genetic algorithm and these approaches leading to **″stochastic amoeba″** [Hill93] and **″stochastic Powell's method″** gives a more robust and more efficient optimisation strategies improving the search for the global extreme hidden among local ones. The stochastic amoeba, but not the stochastic Powell's method, was able to find global extreme in case of MI registration very precisely when we used a randomly generated population of amoebas of population size of 10, and 3 generations requiring CPU time 412.68 seconds and 4657 function evaluations. For both SAVD and NCC the stochastic amoeba worked well with only 4 re-starts without genetic recombinations with random initialisation of positions of amoeba legs (t_{SAVD}=8.68 s, N_{SAVD}=707 f.e.; t_{NCC}=11.65 s, N_{NCC}=814 f.e). The stochastic Powell's method proved to be trapped into a local extreme more frequently than the stochastic amoeba, therefore, 10 re-starts was used (t_{SAVD}=95.93 s, t_{NCC}=81,64 s). Using these methods the sub-pixel accuracy was achieved.

The method of s**imulated annealing** [Szu87], [Press92] is a technique suitable for optimisation problems where a global extreme is masked by many local extremes. This method is an analogy with thermodynamics, specifically with the way that, for example, metals cool and anneal or liquids freeze and crystallise. If a liquid is cooled slowly, its atoms are often able to line themselves up and form a pure crystal that represents the state of minimum energy of the system. On the contrary, if a liquid is cooled quickly, it does not reach this state but rather ends up in a polycrystalline or amorphous state having somewhat higher energy. In physics it is known that the system in thermal equilibrium at temperature T has its energy probabilistically distributed among all different energy states E. This is expressed by the so-called Boltzmann probability distribution,

$$\mathrm{Prob}(E) \sim \exp(-E/kT). \qquad (2)$$

Even at low temperature, there is a chance, albeit very small, of a system being in a high-energy state. Therefore, there is a corresponding chance for the system to get out of a local energy minimum in favour of finding a better, more global, one. The Boltzmann's constant k relates temperature to energy. In other words, the system sometimes goes uphill as well as downhill; but the lower the temperature, the less likely is any significant uphill excursion. To use this method for other then thermodynamic systems, the following items must be defined: a description of possible system configurations, a generator of random changes in the configuration, an objective function E (analog of energy) whose minimisation is the goal of the annealing, a control parameter T (analog of temperature) and an annealing schedule that governs its lowering from high to low values, e.g. after how many random changes in configuration follows each downward step in T, and how large is that step. Simulated annealing was implemented by a modification [Press92] of the above amoeba method: a positive, logarithmically distributed random variable, proportional to the temperature T was added to the stored objective function value associated with every leg of the amoeba, and a similar random variable was subtracted from the function value of every new position of legs. Thus the method always accepts a true downhill step, however, sometimes accepts an uphill one. At a given T, the amoeba covers the size of the region that can be reached at this temperature, and then executes a stochastic, tumbling Brownian motion within that region, sampling new, approximately random, points. If the temperature is reduced sufficiently slowly, it becomes highly likely that the amoeba will shrink into that region containing the minimum value of the objective function.

Simulated annealing is the technique that is able to successfully register images without using repeated re-starts or the genetic approach. The legs of initial amoebas were put to span maximum of the parametrical space, initial temperatures were set to 10, 0.6 and 1 for SAVD, NCC and MI, respectively, and numbers of amoeba iterations for each temperature were set to 30 for SAVD and NCC, and 20 for MI. Number of functional evaluations attained was 1097 for SAVD (t_{SAVD}=16,04 s), 1227 for NCC (t_{NCC}=20,50 s), and 581 for MI (t_{MI}=60,22 s).

CONCLUSIONS

Three similarity criteria were studied – SAVD, NCC, and MI, respectively – in conjunction with optimisation methods – full search, the genetic algorithm, downhill simplex method, Powell's method, simulated annealing – in order to find a relatively quick, but also robust technique suitable for registration of serial optical slices captured by CLSM.

From the point of view of convergence speed of the optimisation strategies and their precision to find the global extreme, the best results were based on the genetic algorithm. The genetic algorithm itself converges to the peak of the extreme rather slowly, therefore, its combinations with downhill simplex method (stochastic amoeba) or Powell's method (stochastic Powell's method) were also tested. These combinations lead to faster going uphill the global extreme and to statistically higher probability of its finding. Simulated annealing worked well too, however, success of finding the global extreme depended on initial position of amoeba legs and determination of the initial temperature that varies with the function. Hence, the version of simulated annealing tested in this paper is not so general as the stochastic methods.

MI was studied here due to its ability to register inter-modality images. This was not examined in our tests, but we only looked for the most optimal strategy that allows determination of its global extreme which is rather sharp, narrow and surrounded by local ones. The stochastic amoeba seems to be the most convenient technique for this purpose.

ACKNOWLEDGEMENTS

This study was supported by grant of Ministry of Education of CR No. VS 97033 "Laboratory of Computer Supported Medical Diagnostics" and grant of the Czech Ministry of Health No. 4178-3.

Thanks go to Dr. Kubínová from Institute of Physiology of Academy of Sciences of the Czech Republic for kind providing optical slices of biological tissue.

REFERENCES

[Čapek98] Čapek, M., Krekule, I., Kubínová, L.: Practical Experiments with Automatic Image Registration of Serial Optical Slices of Thick Tissue. Conf. Proc., 8-th Intern. IMEKO Conf. on Measurement in Clinical Medicine, Dubrovnik, Croatia, p. 10-6–10-9, September 1998.

[Tekal95] Tekalp, A., M.: Digital Video Processing. Prentice Hall, New York, 1995.

[Burt84] Burt, P. J.: The Pyramids as a Structure for Efficient Computation. In: Rosenfeld, A.: Multiresolution Image Processing and Analysis. Springer-Verlag, New York 1984, p. 6-35.

[Parke97] Parker, J. R.: Algorithms for Image Processing and Computer Vision. John Wiley & Sons, Inc., New York, 1997.

[Press92] Press, W. H., et al.: Numerical Recipes in C: the Art of Scientific Computing. 2nd ed., Cambridge University Press, New York, p. 412-420, 1992.

[Hill93] Hill, D., L., G., et al.: A Strategy for Automated Multimodality Image Registration Incorporating Anatomical Knowledge and Imager Characteristic. Conf. Proc. Information Processing in Medical Imaging. IPMI'93, p. 182-196.

[Szu87] Szu, H., Hartley, R.: Fast Simulated Annealing. Physics Letters A, Vol. 122, No. 3,4, p. 157-162, June 1987.

[Brown92] Brown, L. G.: A Survey of Image Registration Techniques. ACM Computing Surveys, Vol. 24, No. 4, p. 325-376, December 1992.

[Fras86] Fraser, A. M., Swinney, H. L.: Independent Coordinates for Strange Attractor from Mutual Information. Physical Review A, Vol. 33, No. 2, p. 1134-1140, February 1986.

[Maes97] Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P.: Multimodality Image Registration by Maximization of Mutual Information. IEEE Trans. Med. Imag., Vol. 16, No. 2, p. 187-198, April 1997.

APPENDIX: GLOBAL SIMILARITY EVALUATION METHODS

SAVD [Brown92] is defined by

$$SAVD(T) = \frac{1}{N} \sum_{x \in I_1} \left| I_1(x) - I_2(T(x)) \right|, \qquad (3)$$

$I_{1,2}$ being a pair of images to be registered, x a vector of co-ordinates of an image pixel and N the number of pixels of picture matrices. SAVD is based on subtraction of values of corresponding pixels, and, therefore, is simple and computationally fast. It attains values between zero and maximal possible value of pixel intensity. We introduced a normalisation factor $1/N$, thus the function range is independent on the picture size. Image registration is achieved by minimum of SAVD.

NCC [Brown92] is expressed by

$$NCC(T)=\frac{\sum_{x\in I_1}[I_1(x)-I_1^*].[I_2(T(x))-I_2^*]}{\{\sum_{x\in I_1}[I_1(x)-I_1^*]^2.\sum_{x\in I_1}[I_2(T(x))-I_2^*]^2\}^{\frac{1}{2}}}, \quad (4)$$

where $I^*_{1,2}$ represent mean intensity values of the corresponding images. It attains values in [-1,1]; the higher the value, the greater is the similarity and the better is registration.

MI [Fras86] [Maes97] is based on estimations of marginal and joint image intensity distributions $p_{I_1}(x), p_{I_2}(T(x))$, and $p_{I_1,I_2}(x,T(x))$, and measures the degree of dependence between I_1 and I_2 by the distance between the joint distribution $p_{I_1,I_2}(x,T(x))$ and the distribution associated to the case of complete independence $p_{I_1}(x).p_{I_2}(T(x))$ by

$$MI(T)=\sum_x p_{I_1,I_2}(x,T(x)).\log\frac{p_{I_1,I_2}(x,T(x))}{p_{I_1}(x).p_{I_2}(T(x))} \quad (5)$$

MI is related to Shannon entropies by the formula

$$MI(T) = H(I_1) + H(I_2,T) - H(I_1,I_2,T), \quad (6)$$

where $H(I_1)$ and $H(I_2,T)$ are the entropy of I_1 and I_2 (mapped by transformation T on I_1), respectively, and $H(I_1,I_2,T)$ their joint entropy

$$H(I_1)=-\sum_x p_{I_1}.\log p_{I_1}(x) \quad (7)$$

$$H(I_2,T)=-\sum_x p_{I_2}.\log p_{I_2}(T(x)) \quad (8)$$

$$H(I_1,I_2,T)=-\sum_x p_{I_1,I_2}(x,T(x)).\log p_{I_1,I_2}(x,T(x)) \quad (9)$$

The bounds of MI are

$$0\le MI(T)\le \min\{H(I_1),H(I_2,T)\}. \quad (10)$$

It attains the lower bound if I_1 and I_2 are independent, and the upper bound $H(I_2,T)$ is reached if I_2 follows uniquely from I_1. If the log is taken to the base two, the volume similarity evaluated by MI is expressed in bits. Images are registered when MI reaches its maximum value.