General Multimodal Elastic Registration Based on Mutual Information

J. B. Antoine Maintz, Erik H.W. Meijering, and Max A. Viergever

Image Sciences Institute, Utrecht University, P.O.Box 80089, 3508 TB, Utrecht, the Netherlands

ABSTRACT

Recent studies indicate that maximizing the mutual information of the joint histogram of two images is an accurate and robust way to rigidly register two mono- or multimodal images. Using mutual information for registration directly in a local manner is often not admissible owing to the weakened statistical power of the local histogram compared to a global one. We propose to use a global joint histogram based on optimized mutual information combined with a local registration measure to enable local elastic registration.

Keywords: multimodal elastic registration, matching, mutual information

1. INTRODUCTION

Medical imaging has become an increasingly vital component of a large number of clinical applications during the past decades. These applications not only concern diagnostic studies, but also planning, performing, and evaluating radiotherapeutical and surgical procedures. A large number of digital imaging modalities is now being used on a routine basis in normal clinical practice. Such modalities include CT, CTA, MRI, MRA, DSA, US*, digital radiography, portal imaging, and video imaging (*e.g.*, laparascopy or laryngoscopy), which constitute anatomical modalities –*i.e.*, primarily depicting patient morphology– and SPECT, PET, fMRI†, which constitute functional modalities, *i.e.*, depicting primarily information on the metabolism of the underlying anatomy. Some other techniques, such as EEG and MEG‡, are strongly image related: even though these are not *direct* imaging techniques, information obtained using these modalities can be presented in a pictorial fashion.

Since information obtained from two images acquired in a clinical track of events is usually of a complementary nature, proper integration of the separate images is often desired. A first step in this integration process is a registration step to ensure that the images of interest are in spatial alignment. In previous work,[1,2] we classified medical image registration methods according to nine criteria, *viz.:*

- Dimensionality: the number of spatial and time dimensions of the images to be registered.
- Registration basis: *extrinsic*, *i.e.*, based on artificial objects introduced into the image; *intrinsic*, *i.e.*, based on patient generated image information only; or using *calibrated scanner coordinate systems*.
- Nature of transformation: rigid, affine, projective or curved.
- Domain of transformation: global or local.
- Degree of user interaction.
- Optimization procedure.
- Modalities involved.
- Subject: intrasubject or intersubject.
- Part of the body imaged.

Corresponding author: J.B.A. Maintz, address see above, email: twan@cs.ruu.nl, phone: +31-30-2533899, fax: +31-30-2513791, www: http://www.cs.ruu.nl/people/twan/

*Respectively Computed Tomography, Computed Tomography Angiography, Magnetic Resonance Imaging, Magnetic Resonance Angiography, Digital Subtraction Angiography, and UltraSound.

†Respectively Single Photon Emission Computed Tomography, Positron Emission Tomography, and functional MRI.

‡Electro EncephaloGraphy and Magneto EncephaloGraphy.

Classifying individual registration methods is a useful exercise in the sense that it allows for an assessment of several method-specific aspects such as extensibility to other applications and problems to be expected in verification of method accuracy. It also provides us with a framework to quickly look up and compare similar methods occurring in the literature.

For practical purposes, the most significant classification criterion is the registration basis. The extrinsic methods can be subdivided into invasive methods –*e.g.*, making use of a stereotactic frame[3,4] or screw markers[5–8]– and non-invasive methods, using a mould, non-invasive frame or dental adapter,[9–13] or fiducial skin markers.[13–18] Intrinsic methods can be based on landmarks,[13,14,18–24] on segmented structures (mostly surfaces)[8,24–33] or on measures computed directly from the grey values; *i.e.*, voxel properties. In the latter case, many different approaches are possible, and a large number of references could be cited. We will however restrict ourselves to references pertaining to mutual information based registration methods.[34–43]

2. METHODS

In this paper, we present a method which is classified as follows:

- Dimensionality: two images having an equal number n of spatial dimensions.
- Registration basis: intrinsic. The original images are employed without any pre-processing, or including any marking devices. Extensions of the method with pre-processing are discussed.
- Nature and domain of the transformation: local curved/elastic transformations.
- Optimization procedure: finding a maximum of a multi-variable function The actual optimization procedure does not seem to be critical (unless the registration application is time-constrained), and we tested methods employing quasi-exhaustive searching and series of one-dimensional hill climbing.
- Modalities: the method was designed for general multimodal images of the same scene.
- Subject: intrasubject, possibly extensible to intersubject depending on the application.
- Part of the body imaged: theoretically any.

2.1. Mutual information based registration

Registration of two multimodal images of the same scene would be an easy task if we had a priori knowledge of a one-to-one relationship f between the grey values of those images, *i.e.*, if we knew that if a voxel has grey value m in image M, this implies that the corresponding voxel in image N has grey value $n = f(m)$, and vice versa. In practice, this never occurs. Even assuming ideal sampling and an absence of noise and any image artifacts, there is no ideal one-to-one relationship, owing to the fact that multimodality imaging generally implies measuring a different physical quantity for each modality.

On the other hand, it is in many cases reasonable to assume that *some* relationship exists. For instance, figure 1 shows a 2D histogram of registered CT and MR images of the human head. The y-axis corresponds to CT grey values, and the x-axis to MR grey values. Although the relationship between the grey values is far from ideal, is will be clear that the occurrence of some MR grey value is *more likely* to have certain CT grey values corresponding to it than others. This likelihood will decrease if the images are brought out of registration.

If we consider the grey values occurring in two images to be registered as stochastic variables, their *mutual information* gives us a measure of the strength of the dependence between these variables. The mutual information I of two random variables M and N is defined by

$$I_{M,N} = \sum_{(m,n)} p_{MN}(m,n) \ \log \left(\frac{p_{MN}(m,n)}{p_M(m) p_N(n)} \right), \tag{1}$$

where p_M and p_N are the distributions of M and N respectively, and p_{MN} is the joint distribution of M and N. Occurrences of $0 \log 0$ are set to zero. In the case of images, we can approximate the distributions from the marginal and joint grey value histograms. The frequency entries can be converted to probabilities by simple normalization. The domain of voxels used in computing the histogram can be the overlapping part of the images.

The mutual information measures the distance (the so-called Kullback-Leibler distance) between the case of complete independence of M and N –with associated distribution $p_M(m)p_N(n)$– and the current situation, with associated distribution $p_{MN}(m,n)$. The limiting cases of complete dependence and complete independence of M and N have mutual information values $I = H$ and $I = 0$ respectively, where H is the entropy of either one of the images M or N.

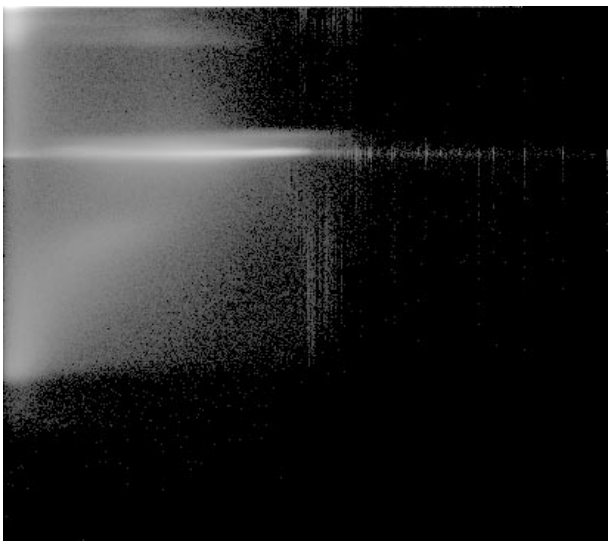

Figure 1. 2D histogram of the grey values of a registered MR (x-axis) and CT image (y-axis).

2.2. Optimization of mutual information: global transformations

Under the assumption that the mutual information of the two images is maximum when the images are in registration, registration can be performed by maximizing the mutual information as a function of a geometric transformation t of one of the input images (say N):

$$\begin{cases} T = t(N) \\ I_{\text{MN}}(t) = \sum\limits_{(m,n)} p_{\text{MT}}(m,n) \ \log\left(\dfrac{p_{\text{MT}}(m,n)}{p_{\text{M}}(m)p_{\text{T}}(n)}\right) \\ t_{\text{reg}} = \arg\max\limits_{t} I_{\text{MN}}(t), \end{cases}$$

where t_{reg} is the transformation that will bring the images into registration. The value of I can as before be computed from the normalized marginal and joint histograms. The histograms are constructed taking only the voxels occurring in the overlapping part of the (transformed) images into account. For many applications –including such applications as rigid or affine registration of CT, MR and PET brain images– it is essential that the adverse influence of random effects is diminished by binning or low-pass filtering the joint histogram prior to computing the mutual information.

This approach to finding the registration transformation is suited for those cases where the number of parameters involved in the global transformation t is limited, *e.g.,* for rigid or affine transformations. For global elastic transformation the number of parameters to be optimized is generally too large to be feasible in practice. A common approach to overcome this is to use a local approach: Usually, the image is divided in windows, a local transformation is found in each window, and the global elastic transformation is found by assimilating all of the local transformations into a continuous transformation. The local transformations are frequently very simple, commonly only translations.

2.3. Local elastic registration

The mutual information approach to global registration described above cannot be used directly locally for most applications: for the local joint histogram to have sufficient statistical power, the window size should be chosen so as to be far too large to yield an acceptable elastic transformation. Choosing overlapping windows can only partially solve this problem.

Our approach to overcome this is to make use of the *global* histogram in finding the local transformations. The local registration is performed by the following procedure:

1. The input images are registered rigidly by optimizing the global mutual information of their joint histogram.
2. From the joint histogram of the rigidly registered input images M and N we extract the conditional probability densities $p(n|m)$, *i.e.,* given a grey value m from image M we compute the probability that this value corresponds to any grey

value n occurring in image N. This can be achieved by normalizing the values of each row of the joint histogram parallel to the N-axis.

3. We perform the elastic registration by dividing the input images into windows and optimizing the local grey value correspondence probability c as a function of translations t of each window. The grey value correspondence $c(t)$ of a window W is defined by
$$c(t) = \sum_{w \in W} p(N(t(w))|M(w)), \tag{2}$$
where $N(w)$ and $M(w)$ are the grey values at voxel w of the images N and M respectively.

4. The found vector field of local displacements is translated into a global transformation and applied to the appropriate input image.

An implicit assumption we have made here is that after step (1), the images are registered up to small local elastic deformations, and that the conditional probabilities computed from the joint histogram after rigid registration are adequate approximations of the "real" probabilities after elastic registration.

The registration process is not symmetrical in the order of input images. In other words, we may in general expect different results when registering an MR to a CT image than when registering a CT to an MR image.

2.4. Implementation issues

The rigid registration step was carried out by optimization of the mutual information using Powell's method.[44] To avoid introducing spurious grey values in the process by interpolation, the joint histogram of the first input image and the transformed second input image is computed using D-linear partial volume interpolation,[39] where D is the dimensionality of the input images. This process entails that, given a voxel w from image M and its transformed counterpart $t(w)$ in N, instead of updating the histogram with the grey-value pair $(M(w), N(t(w)))$, where $N(t(w))$ will generally have to be interpolated from the nearest grid grey values, we fractionally update all the histogram entries $(M(w), N(w_i))$, where w_i ($i = 0, \ldots, 2^D$) is the neighborhood comprising the nearest neighbors of $t(w)$. Each fraction is determined by the weight normally associated with D-linear interpolation.

The number of iterations and tolerances in Powell's method are set to values so large c.q. small that the Powell's optimization is no influential source of inaccuracy in the process given reasonable input images. The bin size, *i.e.*, the number of grey values pooled in one slot of the histogram is set so that the histogram contains 256 bins per image (or less if the image contains less than 256 grey values).

The local optimization of the grey value correspondence c (see equation 2) was done by either

1. exhaustive searching, or
2. iterated 1-D exhaustive searches (along each of the translation axes separately), or
3. iterated 1-D hill-climbing.

In all cases, only integer-valued translations within a limited region were considered. If the optimal value of c occurred at a number of (connected) locations, the center of gravity of these locations was chosen as the optimal location. Since exhaustive searching did not significantly increase the accuracy of our results, we use the (fast) third method by default. A number of aspects were avoided in the implementation:

- Sub-voxel translations were not used, so no interpolation artifacts can occur. Sub-voxel translations were however estimated by parabolic fitting around the optimum found and its nearest neighbors.
- Elastic tissue properties and spatial connectivity were not modeled, *i.e.*, no additional terms were added to equation 2, and the local translations found are truly local and independent of the displacements found nearby.

Note that including these aspects in the implementation may prove valuable, but in the present phase they are hampering the study of the pure behavior of the local registration measure. A disadvantage of this particular *implementation* is that we cannot expect it to perform with sub-voxel accuracy everywhere.

The translation of the vector field of local translations to a global deformation is carried out using special graphics hardware (programmed using the Open Graphics Library). Given a 2D vector field, the displacement at an arbitrary location is determined by bilinear interpolation of the displacements at the nearest grid points. The grey value is established by nearest neighbor interpolation.

3. RESULTS

3.1. Consistency check

Before carrying out actual local registration experiments it is essential to establish that the correspondence measure c is consistent with the mutual information measure used in the rigid registration. To investigate this we computed c taking the entire image for a window W, and examined its behavior around the optimum as established by the mutual information criterion: for consistency, c should have its optimum at the same location. We examined this for a number of artificial and real images (presented later on), and found no inconsistencies. The behavior of c in terms of smoothness was found to be dependent on the chosen bin size of the histogram, see for example figure 2. Regardless of the bin size chosen, artificial shifts of the input

Figure 2. Example of the behavior of the local registration measure c around the optimum as predicted by mutual information. In each figure, the values of c at the translations $(0,0), (1,0), (-1,0), (0,1)$, and $(0,-1)$ are shown graphically, where $(0,0)$ is the correct optimum. The input images used are a CT and an MR image of the head. The left diagram shows the result with a bin size of 1 grey value, the right diagram with a bin size of 50 grey values. The original images have a grey value range of 2500.

images (relative to the mutual information optimum) ranging from 1 to 10 pixels were faultlessly recovered by optimizing c as described before.

3.2. Experiments on artificial images

These experiments were done using the following scheme:

- Create an artificial image I_1.
- Create a second image I_2 by elastically deforming I_1.
- Randomize the grey values of I_2.
- Elastically register I_2 to I_1 based on the joint histogram of I_1 and I_2 as described in section 2.3.
- Evaluate the results by visual inspection and comparison to the applied elastic deformation.

The experiments were carried out using increasingly complex images. In the simplest case, we used a 2D binary image of a square for I_1, and a rectangle for I_2. In this case, the deformation was found *exactly* by the algorithm: the difference image of I_1 and the registered I_2 showed only zeros. The next series of experiments was done using simple geometrical shapes, see *e.g.*, figure 3. After registration, the registration error was sub-voxel throughout the images when visually inspecting difference images. When comparing the applied and recovered deformation vector fields however, there is a significant error, which is undetectable by visual inspection of the result. This is an important observation that we will come back to in the discussion.

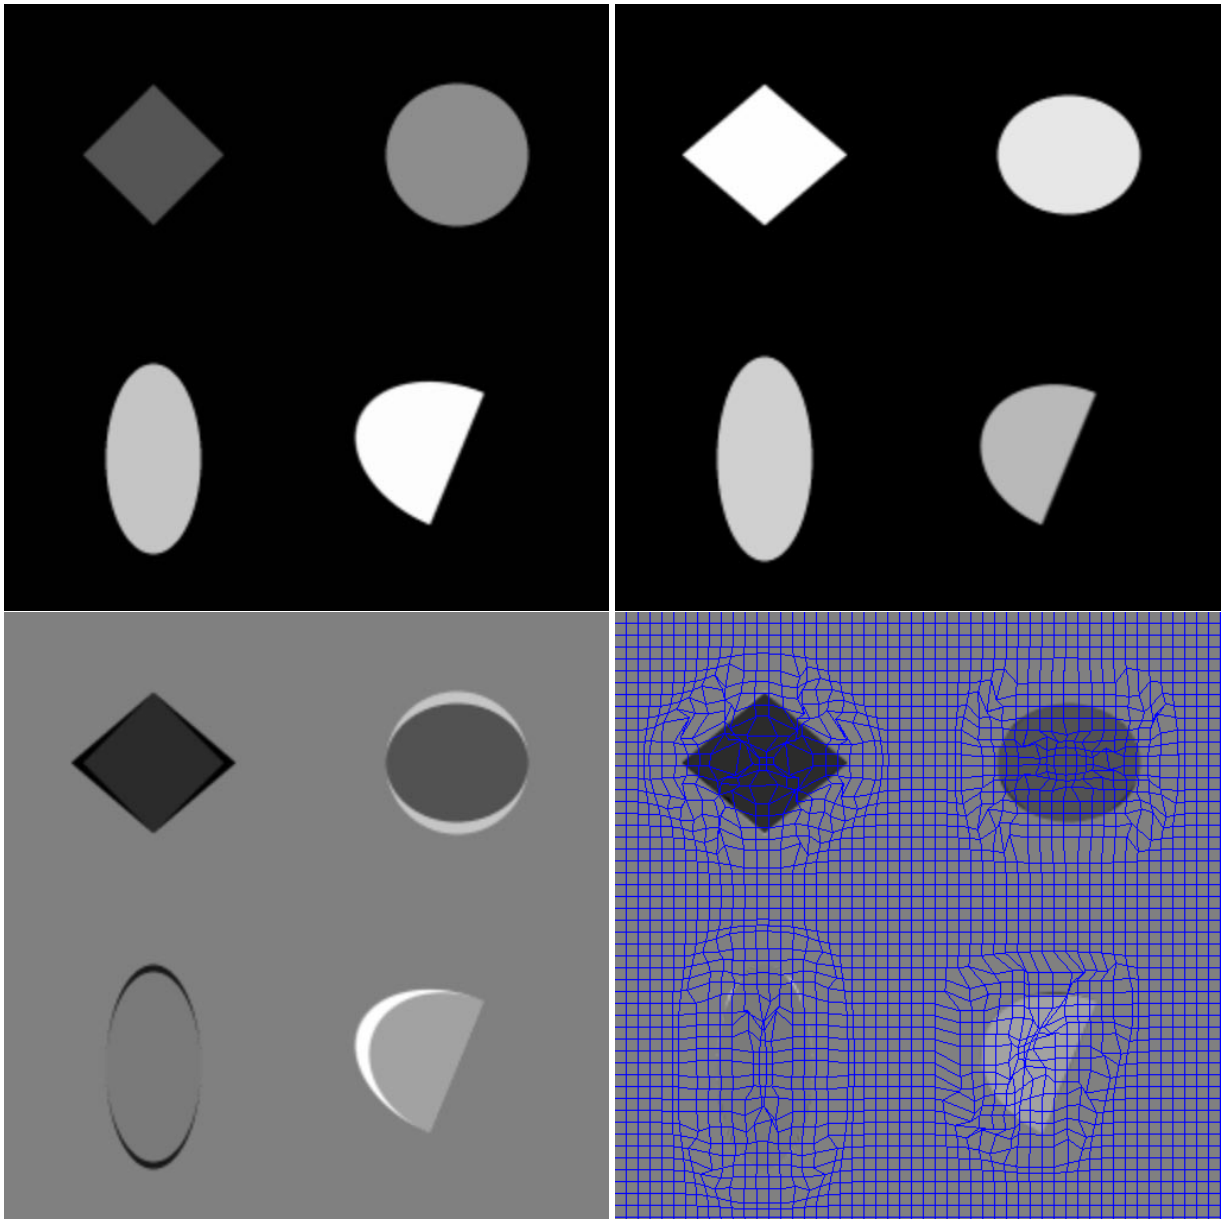

Figure 3. Example of elastic registration of simple geometrical shapes. Top left: original image. Top right: after elastic deformation and grey value randomization. Bottom left: the difference image of the first two images, showing clearly the deformations applied. Bottom right: the difference image after elastic registration, with the registration deformation grid overlaid.

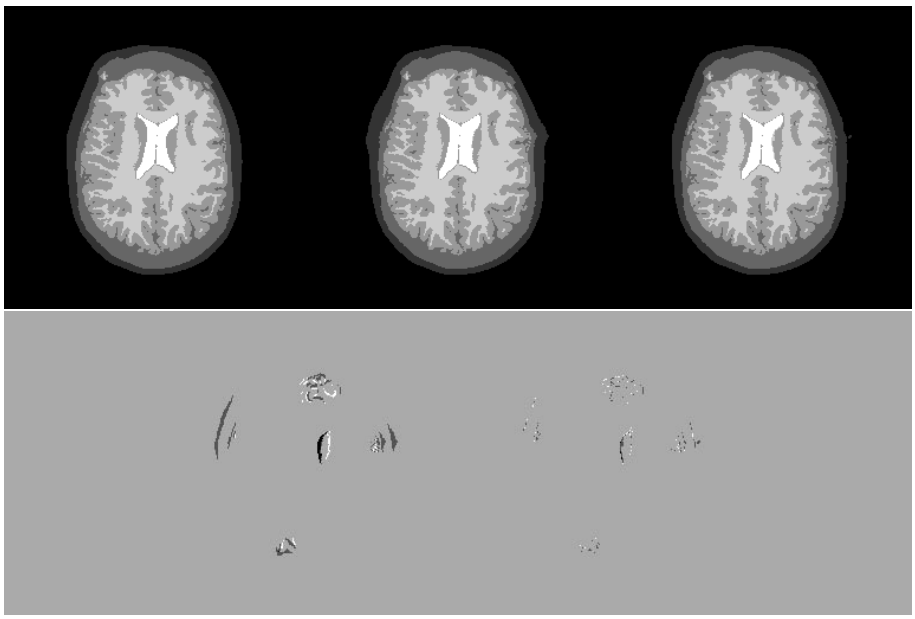

Figure 4. Example of registration using a 6-segment MR image. Top left: original image. Top middle: after manual elastic deformation. Top right: after registration. Bottom left: the difference of the original and the deformed image. Bottom right: the difference of the original and the registered image.

3.3. Experiments using a real and an artificially deformed image

In this series of experiments we adapted our procedure slightly in the sense that I_1 is now a real image. The rest of the procedure as described in the previous section remains unaltered.

This series of experiments was carried out using real images with the number of grey values reduced in such a way that they contained little partial volume effects. The images I_1 included both CT and MR brain images, as well as segmented versions of these. A typical result is presented in figure 4. In this example, the original image contained 256×256 pixels, the deformed image contained 1.06% voxels that were misclassified, after registration using windows of 10×10 pixels, this was reduced to 0.28%. It is however inappropriate to attach much value to these quantifications for two reasons:

- Considering the convoluted nature of the segment boundaries, the nature of the deformation, and the "window" implementation of the registration, we can not expect perfect results given reasonable (relatively large) window sizes.
- We *can* reduce the number of misclassified voxels in the registration to zero by simply choosing a window size of 1, which results in each pixel getting "pushed" to the nearest correct segment, but this would violate the spatial connectivity of the image.

Enlarging the window size reduces the occurrence of the latter effect greatly, but in the current implementation reduces the ability of the system to correct for elastic deformations that cannot be captured using local translations.

3.4. Experiments on real CT and MR images

In these experiments we used real high resolution CT and T1 weighted MR brain images without any pre-processing. The images were rigidly registered in 3D by optimization of mutual information as described in section 2.2. This registration was verified by visual inspection of the CT bone contours overlaid onto the registered MR images. The "fit" of these contours was evaluated with the contours in registered position, and shifted one pixel along any image axis. This process was carried out in a number of original transversal slices as well as in reformatted coronal and sagittal slices. It was established that no one-pixel shift improved the visual quality of the match at the bone contours. There were clear visual differences elsewhere in the images, notably so at the back of the head, where the supports touch the patient's head.

The results after local elastic registration were ambiguous. The registration at the back of the head was now improved, as figure 5 shows, but incorrect window translations were found at places were the grey value correspondence locally differs significantly from the global one, *e.g.,* around the eyes. This problem can be remedied here by setting all histogram entries

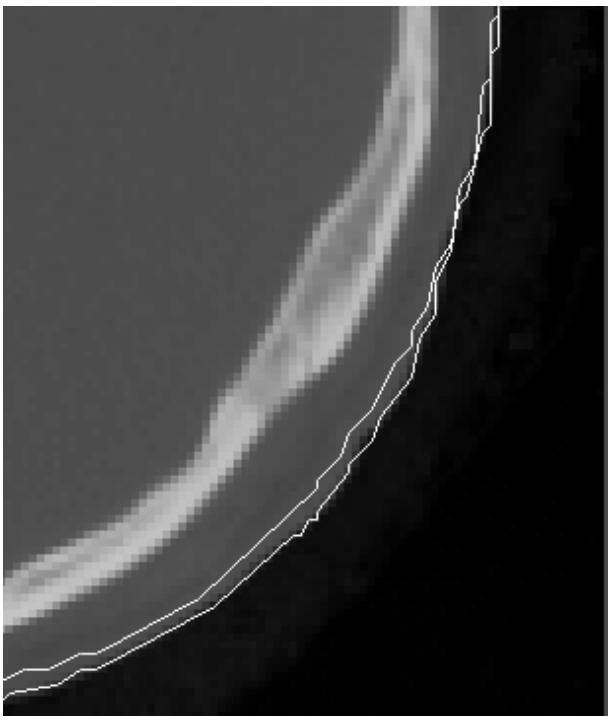

Figure 5. Example of elastic registration (using 10×10 windows) after rigid registration of a CT and an MR image. Shown is a zoom of part of the CT image. The upper overlaid contour represents the MR skin after rigid registration, the bottom contour after elastic registration.

below a certain threshold to zero, but such an application dependent approach is not generally recommended in this phase of the research. Interestingly, the problem can *not* be solved by allowing only those elastic transformations where the measure c exceeds some threshold.

4. CONCLUSIONS AND DISCUSSION

We have presented a method for general multimodal elastic registration. It is based on an extension of mutual information, although the mutual information itself is only used explicitly in a pre-processing step which entails rigid registration of the images at hand. The subsequent elastic registration step locally optimizes the grey value correspondence probabilities as approximated from the joint histogram of the images registered rigidly by optimization of mutual information. The preliminary experimental results presented here show that the method has potential for a number of applications. On the other hand, it also demonstrates a number of pitfalls and problems that need to be addressed.

An interesting observation is that the deformation grid found by the registration process may *not* be physically correct, even though it apparently "solves" the registration problem to –in this case– a sub-voxel level of precision when inspecting difference images of original and registered images. It is reasonable to assume that this phenomenon will also occur to some degree in real medical images, therefore two observations can be made: (1) modeling of the deformations using elastic properties of tissues to some degree is mandatory, and (2) caution is warranted when elastic deformation fields obtained by registration procedures are used for surgical (or otherwise interventional) localization procedures.

For the modeling of elastic properties we intend to consider the following approach: after the rigid registration step:

- Perform the elastic registration step, but restrict the deformation to very small deformations locally.

- Check the resultant vector field for inconsistencies and correct outliers.
- Repeat this process until no significant deformations are found anymore.

A further improvement –which for many medical applications may prove mandatory– is to model the elastic tissue properties more elaborately, by assigning a rigidity/elasticity variable to each tissue type. This approach however requires a segmentation to be carried out.

It has been noted that the smallest registration window (1×1) produces trivial and physically incorrect –if visually perfect!– results. This can be attributed to the fact that a 1×1 window does not take the spatial context of a voxel into account. Enlarging the window increases the likelihood of a realistic solution, since more spatial information is being considered.

Our method implicitly assumes that the grey value correspondence obtained from the global histogram adequately represents the local situation in the images. As our last experiment shows, this assumption may break down for some types of images. A range of grey values may locally correspond to a different range than the one it corresponds globally to. Unfortunately such problems are unavoidable to some extent in *any* method of local registration: local registration is only possible if there is local information available, either shape information or –as in our case– by a grey value correspondence. If absent, a local registration may at best be inferred from neighboring transformations, but even then it may be incorrect.

In this paper, we have suggested a general approach to local registration of arbitrary multimodal images. We have shown that the method has potential for elastic registration. For many applications however, it is vital that the method be adapted to the specifics of the images at hand, in that the physical properties in terms of tissue elasticity and tissue-associated grey values are considered.

ACKNOWLEDGMENTS

We would like to thank Frederik Maes, André Collignon, and Dirk Vandermeulen (Laboratory for Medical Imaging Research, KU Leuven, Belgium) for providing us with their software.

REFERENCES

1. J. B. A. Maintz and M. A. Viergever, "A survey of medical image registration," *Medical image analysis* **2**(1), p. ?, 1998. In press.
2. J. B. A. Maintz, *Retrospective registration of tomographic brain images*. PhD thesis, Utrecht University, 1996.
3. L. D. Lunsford, *Modern stereotactic neurosurgery*, Martinus Nijhoff, Boston, MA, 1988.
4. T. M. Peters, J. A. Clarck, and A. Olivier, "Integrated stereotaxic imaging with CT, MR imaging, and digital subtraction angiography," *Radiology* **161**, pp. 821–826, 1986.
5. K. P. Gall and L. J. Verhey, "Computer-assisted positioning of radiotherapy patients using implanted radioopaque fiducials," *Medical physics* **20**(4), pp. 1152–1159, 1993.
6. C. R. Maurer, J. J. McCrory, and J. M. Fitzpatrick, "Estimation of accuracy in localizing externally attached markers in multimodal volume head images," in *Medical imaging: image processing*, M. H. Loew, ed., vol. 1898, pp. 43–54, SPIE Press, (Bellingham, WA), 1993.
7. C. R. Maurer, J. M. Fitzpatrick, R. L. Galloway, M. Y. Wang, R. J. Maciunas, and G. S. Allen, "The accuracy of image-guided neurosurgery using implantable fiducial markers," in *Computer assisted radiology*, H. U. Lemke, K. Inamura, C. C. Jaffe, and M. W. Vannier, eds., pp. 1197–1202, Springer-Verlag, (Berlin), 1995.
8. D. A. Simon, M. Hebert, and T. Kanade, "Techniques for fast and accurate intra-surgical registration," *Journal of image guided surgery* **1**(1), 1995.
9. T. Greitz, M. Bergström, J. Boëthius, D. Kingsley, and T. Ribbe, "Head fixation system for integration of radiodiagnostic and therapeutic procedures," *Neuroradiology* **19**, pp. 1–6, 1980.
10. L. V. Laitinen, B. Liliequist, M. Fagerlund, and A. T. Eriksson, "An adapter for computer tomography guided stereotaxis," *Surgical neurology* **23**, pp. 559–566, 1985.
11. L. R. Schad, R. Boesecke, W. Schlegel, G. H. Hartmann, G. H. Sturm, L. G. Strauss, and W. Lorenz, "Three dimensional image correlation of CT, MR, and PET studies in radiotherapy treatment of brain tumors," *Journal of computer assisted tomography* **11**, pp. 948–954, 1987.
12. D. J. Hawkes, D. L. G. Hill, and E. C. M. L. Bracey, "Multi-modal data fusion to combine anatomical and physiological information in the head and heart," in *Cardiovascular nuclear medicine and MRI*, J. H. C. Reiber and E. E. van der Wall, eds., pp. 113–130, Kluwer, (Dordrecht, the Netherlands), 1992.

13. A. C. Evans, S. Marrett, J. Torrescorzo, S. Ku, and L. Collins, "MRI-PET correlation in three dimensions using a volume of interest (VOI) atlas," *Journal of cerebral blood flow and metabolism* **11**, pp. A69–A78, 1991.
14. G. Q. Maguire, M. Noz, H. Rusinek, J. Jaeger, E. L. Kramer, J. Sanger, and G. Smith, "Graphics applied to medical image registration," *IEEE Computer graphics and applications* **11**(2), pp. 20–28, 1991.
15. M. Y. Wang, J. M. Fitzpatrick, C. R. Maurer, and R. J. Maciunas, "An automatic technique for localizing externally attached markers in MR and CT volume images of the head," in *Medical imaging: image processing*, M. H. Loew, ed., vol. 2167, pp. 214–224, SPIE Press, (Bellingham, WA), 1994.
16. M. Y. Wang, J. M. Fitzpatrick, and C. R. Maurer, "Design of fiducials for accurate registration of CT and MR volume images," in *Medical imaging*, M. Loew, ed., vol. 2434, pp. 96–108, SPIE, 1995.
17. R. D. Bucholz, K. R. Smith, J. Henderson, and L. McDurmont, "Intraoperative localization using a three dimensional optical digitizer," in *Medical robotics and computer assisted surgery*, pp. 283–290, 1994.
18. P. J. Edwards, D. L. G. Hill, D. J. Hawkes, R. Spink, A. C. F. Colchester, A. Strong, and M. Gleeson, "Neurosurgical guidance using the stereo microscope," in *Computer vision, virtual reality, and robotics in medicine*, N. Ayache, ed., vol. 905 of *Lecture notes in computer science*, pp. 555–564, Springer-Verlag, (Berlin), 1995.
19. A. C. Evans, S. Marrett, L. Collins, and T. M. Peters, "Anatomical-functional correlative analysis of the human brain using three dimensional imaging systems," in *Medical imaging: image processing*, R. Schneider, S. Dwyer III, and R. Jost, eds., vol. 1092, pp. 264–274, SPIE press, (Bellingham, WA), 1989.
20. D. L. G. Hill, D. J. Hawkes, J. E. Crossman, M. J. Gleeson, T. C. S. Cox, E. C. M. L. Bracey, A. J. Strong, and P. Graves, "Registration of MR and CT images for skull base surgery using pointlike anatomical features," *British journal of radiology* **64**(767), pp. 1030–1035, 1991.
21. Y. Ge, J. M. Fitzpatrick, J. R. Votaw, S. Gadamsetty, R. J. Maciunas, R. M. Kessler, and R. A. Margolin, "Retrospective registration of PET and MR brain images: an algorithm and its stereotactic validation," *Journal of computer assisted tomography* **18**(5), pp. 800–810, 1994.
22. D. Vandermeulen, A. Collignon, J. Michiels, H. Bosmans, P. Suetens, G. Marchal, G. Timmens, P. van den Elsen, M. Viergever, H. Ehricke, D. Hentschel, and R. Graumann, "Multi-modality image registration within COVIRA," in *Medical imaging: analysis of multimodality 2D/3D images*, L. Beolchi and M. H. Kuhn, eds., vol. 19 of *Studies in health, technology and informatics*, pp. 29–42, IOS Press, Amsterdam, 1995.
23. I. G. Zubal, S. S. Spencer, J. S. Khurseed Imam, E. O. Smith, G. Wisniewski, and P. B. Hoffer, "Difference images calculated from ictal and interictal technetium-99m-HMPAO SPECT scans of epilepsy," *Journal of nuclear medicine* **36**(4), pp. 684–689, 1995.
24. A. C. Evans, D. L. Collins, P. Neelin, and T. S. Marrett, "Correlative analysis of three-dimensional brain images," in *Computer-integrated surgery*, R. H. Taylor, S. Lavallée, G. C. Burdea, and R. Mösges, eds., Technology and clinical applications, ch. 6, pp. 99–114, MIT Press, Cambridge, MA, 1996.
25. C. Chen, C. A. Pelizzari, G. T. Y. Chen, M. D. Cooper, and D. N. Levin, "Image analysis of PET data with the aid of CT and MR images," in *Information processing in medical imaging*, pp. 601–611, 1987.
26. D. N. Levin, C. A. Pelizzari, G. T. Y. Chen, C. Chen, and M. D. Cooper, "Retrospective geometric correlation of MR, CT, and PET images," *Radiology* **169**(3), pp. 817–823, 1988.
27. D. L. G. Hill, D. J. Hawkes, N. A. Harrison, and C. F. Ruff, "A strategy for automated multimodality image registration incorporating anatomical knowledge and imager characteristics," in *Information processing in medical imaging*, H. H. Barrett and A. F. Gmitro, eds., vol. 687 of *Lecture notes in computer science*, pp. 182–196, Springer-Verlag, (Berlin), 1993.
28. H. Rusinek, W. Tsui, A. V. Levy, M. E. Noz, and M. J. de Leon, "Principal axes and surface fitting methods for three-dimensional image registration," *Journal of nuclear medicine* **34**, pp. 2019–2024, 1993.
29. A. Collignon, D. Vandermeulen, P. Suetens, G. Marchal, A. Baert, and A. Oosterlinck, "Automatic registration of 3D images of the brain based on fuzzy objects," in *Medical imaging: image processing*, M. H. Loew, ed., vol. 2167, pp. 162–175, SPIE Press, (Bellingham, WA), 1994.
30. W. E. L. Grimson, T. Lozano-Pérez, W. M. Wells III, G. J. Ettinger, S. J. White, and R. Kikinis, "An automatic registration method for frameless stereotaxy, image guided surgery, and enhanced reality visualization," in *Computer vision and pattern recognition*, pp. 430–436, IEEE Computer Society press, (Los Alamitos, CA), 1994.
31. F. Betting, J. Feldmar, N. Ayache, and F. Devernay, "A new framework for fusing stereo images with volumetric medical images," in *Computer vision, virtual reality, and robotics in medicine*, N. Ayache, ed., vol. 905 of *Lecture notes in computer science*, pp. 30–39, Springer-Verlag, (Berlin), 1995.

32. Y. Ge, C. R. Maurer, and J. M. Fitzpatrick, "Surface-based 3-D image registration using the iterative closest point algorithm with a closest point transform," in *Medical Imaging: Image processing*, M. H. Loew and K. M. Hanson, eds., vol. 2710, pp. 358–367, SPIE, (Bellingham, WA), 1996.
33. W. E. L. Grimson, G. J. Ettinger, S. J. White, T. Lozano-Pérez, W. M. Wells III, and R. Kikinis, "An automatic registration method for frameless stereotaxy, image guided surgery, and enhanced reality visualization," *IEEE Transactions on medical imaging* **15**(2), pp. 129–140, 1996.
34. A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal, "Automated multimodality image registration using information theory," in *Information processing in medical imaging*, Y. Bizais and C. Barillot, eds., pp. 263–274, Kluwer Academic Publishers, Dordrecht, 1995.
35. P. Viola and W. Wells III, "Alignment by maximization of mutual information," in *International conference on computer vision*, pp. 16–23, IEEE computer society press, (Los Alamitos, CA), 1995.
36. P. A. Viola, *Alignment by maximization of mutual information*. PhD thesis, Massachusetts institute of technology, artificial intelligence laboratory, 1995. A. I. technical report No. 1548.
37. W. M. Wells III, P. Viola, and R. Kikinis, "Multi-modal volume registration by maximization of mutual information," in *Medical robotics and computer assisted surgery*, pp. 55–62, Wiley, 1995.
38. F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multi-modality image registration by maximization of mutual information," in *Mathematical methods in biomedical image analysis*, pp. 14–22, IEEE computer society press, (Los Alamitos, CA), 1996.
39. F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multi-modality image registration by maximization of mutual information," *IEEE Transactions on medical imaging* **16**, pp. 187–198, April 1996.
40. P. Pokrandt, "Fast non-supervised matching: a probabilistic approach," in *Computer assisted radiology*, H. U. Lemke, M. W. Vannier, K. Inamura, and A. G. Farman, eds., vol. 1124 of *Excerpta medica - international congress series*, pp. 306–310, Elsevier, (Amsterdam), 1996.
41. C. Studholme, D. L. G. Hill, and D. J. Hawkes, "incorporating connected region labelling into automated image registration using mutual information," in *Mathematical methods in biomedical image analysis*, pp. 23–31, IEEE computer society press, (Los Alamitos, CA), 1996.
42. P. Viola, N. N. Schraudolph, and T. J. Sejnowski, "Empirical entropy manipulation for real-world problems," *Advances in neural information processing systems* **8**, 1996. In press.
43. W. M. Wells III, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis, "Multi-modal volume registration by maximization of mutual information," *Medical image analysis* **1**(1), pp. 35–51, 1996.
44. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical recipes in C*, Cambridge university press, second ed., 1992.