Available online at www.sciencedirect.com

SCIENCE DIRECT®

Medical Image Analysis 10 (2006) 432–439

MEDICAL IMAGE ANALYSIS

www.elsevier.com/locate/media

Normalized mutual information based registration using *k*-means clustering and shading correction

Z.F. Knops [a,*], J.B.A. Maintz [a], M.A. Viergever [a,b], J.P.W. Pluim [b]

[a] *Utrecht University, Department of Computer Science, P.O. Box 80089, NL-3508 TB Utrecht, The Netherlands*
[b] *University Medical Center Utrecht, Image Sciences Institute, Q05.459, P.O. Box 85500, NL-3508 GA Utrecht, The Netherlands*

Received 12 October 2004; received in revised form 26 January 2005; accepted 4 March 2005
Available online 18 August 2005

Abstract

In this paper the influence of intensity clustering and shading correction on mutual information based image registration is studied. Instead of the generally used equidistant re-binning, we use *k*-means clustering in order to achieve a more natural binning of the intensity distribution. Secondly, image inhomogeneities occurring notably in MR images can have adverse effects on the registration. We use a shading correction method in order to reduce these effects. The method is validated on clinical MR, CT and PET images, as well as synthetic MR images. It is shown that by employing clustering with inhomogeneity correction the number of misregistrations is reduced without loss of accuracy thus increasing robustness as compared to the standard non-inhomogeneity corrected and equidistant binning based registration.

Keywords: Image registration; Normalized mutual information; *k*-Means clustering; Shading correction; Robustness

1. Introduction

In clinical practice, detailed knowledge of anatomical structures aids the physician with patient diagnosis or treatment. Three-dimensional (3D) imagers, for example MRI or CT scanners, provide the physician with detailed knowledge in image form of the patient. Since different scanners make information available about different tissue properties, it is possible that for proper assessment of all information the physician needs the use of several scanner types. The combination of the images from the different scanners provides the physician with all the information requested. For many representations of the fused data, spatial alignment, usually denoted by image registration, of the images has to be carried out.

In order to find the correct spatial alignment of images, a measure of the quality of the registration is necessary. In this paper we use normalized mutual information (NMI) (Collignon, 1998; Collignon et al., 1995; Studholme et al., 1999; Viola, 1995; Viola and Wells III, 1995; Wells et al., 1995) as a measure for rigid 3D registration of clinical and artificial images. For a number of modalities NMI has proven to be a robust and accurate similarity measure in both mono- and multimodality image registration. With the given quality measure we use rigid transformations and Powell's method to find the optimal spatial alignment.

The spatial transformation yielding the highest NMI value is assumed to be the optimal registration of two images. An estimate of the intensity distribution of the images is necessary to compute the NMI value. It is

* Corresponding author. Tel.: +31 30 253 1454; fax: +31 30 251 3791.
E-mail addresses: zeger@cs.uu.nl (Z.F. Knops), twan@cs.uu.nl (J.B.A. Maintz), max@isi.uu.nl (M.A. Viergever), josien@isi.uu.nl (J.P.W. Pluim).

1361-8415/$ - see front matter © 2005 Published by Elsevier B.V.
doi:10.1016/j.media.2005.03.009

customary to use the histogram of intensities for this purpose. To reduce noise effects neighboring intensities are usually grouped together in a so-called bin. The number of bins is chosen a priori and each intensity is mapped in a rounded linear fashion to the bins of the histogram. The fact that histogram binning does not take the intensity distribution of the image into account implies that natural image segments may be distributed over different bins. This splitting of segments is likely to have adverse effects on the registration procedure in terms of convergence and probability of misregistration. In this paper the binning method *k*-means clustering (MacQueen, 1967) is compared to the standard equidistant binning. *k*-Means clustering uses a variable bin size for each bin in order to achieve a more natural clustering.

An additional complication for registration can be intensity inhomogeneities that occur in MR images. The images may exhibit slow intensity variations within an anatomical structure. These variations can have adverse effects on the registration process owing to increased dispersion of this structure in the intensity distribution. Many solutions to counter the effects of inhomogeneities have been proposed. We use the method as proposed by Likar et al. (2000, 2001), which is based on entropy minimization. Likar et al. describe the image degradation by a linear model consisting of a multiplicative and an additive component modeled by a linear image formation model. Optimization of these components is done by reducing the image entropy while the global intensity statistics are preserved. Likar's method is well suited for NMI registration: both are entropy based. Furthermore, it is likely that *k*-means clustering will be enhanced by the inhomogeneity correction because of the reduction of histogram dispersion (Knops et al., 2003a,b).

In the experiments discussed in this paper it was found that the optimal NMI value in general corresponds to an accurate registration. The process of locating this optimum is, however, not fully robust, resulting in misregistrations. In order to compensate for misregistrations and improve robustness several methods have been proposed (Pluim, 2001). In this paper we combine a new binning approach with an inhomogeneity correction in order to reduce the number of misregistrations without loss of accuracy or a large increase in computation time.

In Section 2 we will discuss the registration process with NMI and describe the binning method and the inhomogeneity correction. In Section 3 the image data sets that will be used in this study are introduced. In Section 4 we will present the experimental results of registration with our binning approach combined with inhomogeneity correction with respect to robustness, accuracy and speed of convergence. Our conclusions are summarized and discussed in Section 5.

2. Methods

2.1. Normalized mutual information

For a discrete random variable A, where $p_A(a)$ is the probability that A has value a, the Shannon entropy H is defined as

$$H(A) = -\sum_{a} p_A(a) \log p_A(a).$$

If the entropy of an image intensity distribution is computed, the entropy measures how well we are able to predict the intensity at an arbitrary point in the image. If there is no uncertainty about the intensity, the entropy is zero, and the image is completely homogeneous. On the other hand, if the image consists of a large number of intensities which all have the same probability, the entropy will be high. Note that the above definition of entropy, if applied to images, does not take any spatial information into account.

For two discrete random variables A and B the Shannon entropy of their joint distribution is defined as

$$H(A,B) = -\sum_{a,b} p_{AB}(a,b) \log p_{AB}(a,b).$$

A joint histogram, which represents the distribution of the intensity couples of corresponding voxels in images A and B, can be used to compute the Shannon entropy.

The *mutual information*, MI, of two images expresses how much the uncertainty on one of the image decreases when the other one is known. It is assumed to be maximum when the images are registered. Studholme et al. (1999) have shown that the registration quality might decrease despite an increasing MI value. For example, if the images are misaligned such that only one voxel overlaps the MI is at a maximum but it is clear the quality of the registration is not optimal. To counter the effect of increasing MI with decreasing registration quality we use normalized mutual information, NMI, in this study. NMI is a well-established registration quality measure which can be defined in terms of image entropies (Collignon et al., 1995; Studholme et al., 1999):

$$\mathrm{NMI}(A,B) = \frac{H(A) + H(B)}{H(A,B)}.$$

2.2. Intensity inhomogeneity correction in MR images

In MR images inhomogeneities can occur, which are partially due to the overall patient anatomy and position, and partially to technical aspects such as poor radio frequency (RF) coil uniformity, static field inhomogeneity, RF penetration and gradient-driven eddy currents. The resulting inhomogeneities occur as slow intensity variations of the same tissue class throughout the image. Moreover, there are spurious intensity variations which may reach up to 30% of the image amplitude (Likar et al., 2000, 2001).

Likar et al. (2000, 2001) assume that an image with inhomogeneities, P, has higher entropy than the same image, A, without inhomogeneities:

$$H(P) = H(f(A)) > H(A),$$

where f models intensity inhomogeneity. Under the constraint of constant mean intensity a correction f^{-1} that minimizes the entropy of the corrected image $f^{-1}(P)$ is pursued. The function f is modeled by a combination of a multiplicative and an additive component

$$P = f(A) = f_{\mathrm{Mul}}A + f_{\mathrm{Add}}.$$

Note that in the above equations we have left out spatial coordinates for legibility, i.e. the equation is to be carried out on a voxel-wise basis. Both f_{Mul} and f_{Add} are five-parameter quadratic functions which have been shown to model the observed intensity variations well. The optimal parameters are found using Powell's method (Press et al., 1992).

2.3. Non-equidistant binning by clustering

In order to compute the NMI measure, a joint histogram of image intensities has to be estimated. In contrast to equidistant binning, k-means clustering takes the entries of the histogram itself into account in order to get a better estimation of the intensity distribution. The length of each interval, the bin width, is adjusted according to minimal variance of intensities in that interval of the image histogram. During registration the images remain unaltered thus the intensity distribution has to be estimated only once at the beginning of the registration process.

The main drawback of equidistant binning is the possibility that the intensities of a single structure end up in different bins. For example, an anatomical structure whose representation in the image histogram is a clear and distinct peak could be divided over two bins. If we vary the bin size in order to minimize the sum of variances of intensities for all bins, a more natural clustering is achieved and as a result the joint histogram is less dispersed.

A k-means clustering problem is generally solved using an iterative procedure (Jain and Dubes, 1988). The clustering thus obtained by this procedure is dependent on the initialization, and it is possible a suboptimal solution is found. We use dynamic programming (Bellman, 1962) to find the optimal clustering for an image. Computation time is kept reasonable, $\sim$1 min on a conventional PC, by intelligent subdivision of the clustering problem (Knops et al., 2004).

In Fig. 1 the results of k-means clustering and equidistant binning on an MR image are shown. The number of bins was eight and contrast has been adjusted in all three images to emphasize the differences. Dissimilarities between the equidistantly binned image and the original MR image are noticeable, for example the segment denoted by the white arrow, which clearly is larger in the original image. There is a clear correspondence between segments in the k-means clustered image and the original image. For each modality used in this study the background is composed of only a few bins with a large bin width, regardless of the clustering structure in the remainder of the intensity distribution differs per modality.

k-Means clustering is sensitive to peaks in the histogram; the bin width near peaks is decreased compared to more dispersed areas of the histogram in order to minimize the bin variance. Inhomogeneities disperse the histogram and broaden peaks, which results in less accurate bin widths for anatomical structures that are represented as peaks in the histogram. By using an inhomogeneity correction method based on lowering image entropy, such as Likar's method, the histogram dispersion is reduced. The reduction in dispersion is likely to enhance the k-means clustering of the image histogram and capture the anatomical structures more

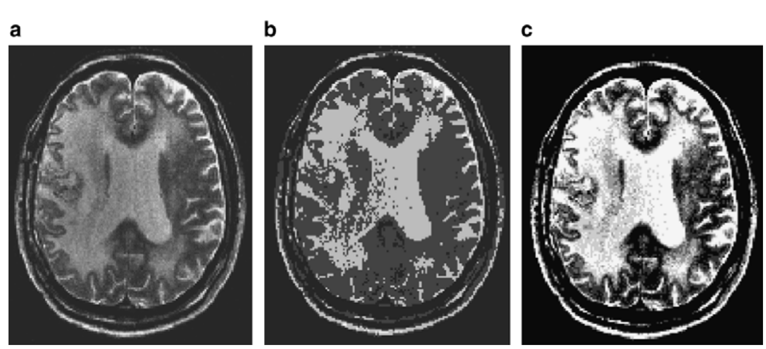

Fig. 1. (a) shows a slice of a 3D MRI dataset. Effects of equidistant binning with eight bins are visible in image (b). Segments that are clearly visible in (a) have been split over several bins. Image (c) is the result of the k-means clustering binning method. Segments that are clearly visible in (a) are not significantly altered. Contrast has been adjusted in all images to emphasize the differences.

accurately. This in turn may enhance the registration process with respect to the number of misregistrations and computational speed.

3. Image datasets

Three sets of multimodality volumetric data consisting of 15 CT-MRI and seven PET-MRI and two artificial MRI brain volumes have been used in this study. The patient studies were taken from the RREP data (http://www.vuse.vanderbilt.edu/~image/registration; West et al., 1997), and the artificial brain volumes were created using the Brainweb simulator (http://www.bic.mni.mcgill.ca/brainweb).

All patient data come with a gold standard for registration based on screw markers with a 0.4 mm accuracy for the CT-MR studies and a 1.7 mm accuracy for the PET-MR studies (West et al., 1997). Furthermore, the RREP dataset is an internationally accepted standard dataset for comparing rigid registration methods. The data characteristics are shown in Table 1.

4. Results and experiments

Registration results using equidistant binning and k-means clustering were compared with and without inhomogeneity correction. NMI was optimized as a function of rigid 3D geometric transformations using Powell's method (Press et al., 1992). The number of optimization steps taken to reach registration can be used as a measure of computational speed. Using the known gold standard the quality of the solution found was assessed. The registrations were performed for a variety of numbers of bins and starting positions. Multiple starting positions were used to investigate registration robustness.

To compare our solution with the gold standard, a sphere was fitted approximately around the head (i.e. centered on the midpoint of the image) with 10,000 points on the boundary. For each point the Euclidean distance between the gold standard position and the position after transformation with our solution was computed. The median value of these distances was taken as the error measure.

A result was considered a misregistration if the median error exceeded the largest voxel dimension of the two registered images. For PET-MRI a registration was considered a misregistration if the median error was above 8 mm. A CT-MRI registration was considered a misregistration if the median error was above 4 or 3 mm depending on patient data. Registrations done with the artificial MRI volume were considered a misregistration if the median error was above 1 mm.

4.1. PET-MRI

For all experiments the gold standard transformation was used as initial transformation. An offset transformation was added by using a rotation offset of $-5°$ or $+5°$ or -10 or $+10$ mm translation for each coordinate axis. All offset combinations were used resulting in 64 different starting positions for each patient registration totaling 448 experiments for each bin size. Initially we intended to use 4, 8, 16, 32, 64, 128 and 256 bins for our experiments. However, since PET has a relatively small intensity range, using 256 bins results in a large number of misregistrations due to increased noise effects. Hence, no experiments were done using 256 bins.

The number of optimization steps needed to reach registration, the average of the median errors and the number of misregistrations for PET-MRI with and without inhomogeneity correction for all starting positions are compared for equidistant binning and k-means clustering in Fig. 2.

Note that the average number of optimization steps, and the average median error are for correct registrations only. Also note that a specific offset transformation that yields a misregistration for one method could yield a good registration for another method. Although including all registrations in the computation of these measures gives an unbiased picture of the overall performance of the methods, the inclusion of successful and

Table 1
The RREP data characteristics

CT	The CT slices are 512 pixels square with a slice thickness of 4.0 or 3.0 mm. The number of slices varies between 27 and 34, respectively, 40 and 49. The pixel size is 0.65 mm^2, respectively, varies between 0.40 and 0.45 mm.
MRI	The MR T2 weighted slices have a resolution of 256 pixels square, the slice thickness is 4.0 or 3.0 mm. The number of slices varies between 20 and 26 or was 52. The pixel size is 0.82, 0.86 or 0.78 mm^2, respectively, 1.25 mm^2
PET	For PET scanning each patient was injected with 10 mCi of 18F-fluorodeoxyglucose. Scanning was started 40–50 min after injection and continued for 25 min. The slice thickness is 8.0 mm and the pixels are 2.59 mm^2. All slices have a resolution of 128 pixels square; the number of slices was 15.
Artificial MRI	The artificial brain volumes consist of 1.0 mm thick slices with a pixel size of 1.0 mm^2. All the slices have a resolution of 181 by 217 pixels and the number of slices was 181. Two volumes were constructed: a T1 weighted volume without inhomogeneity and noise and a T2 weighted volume with 40% intensity non-uniformity and 30% noise.

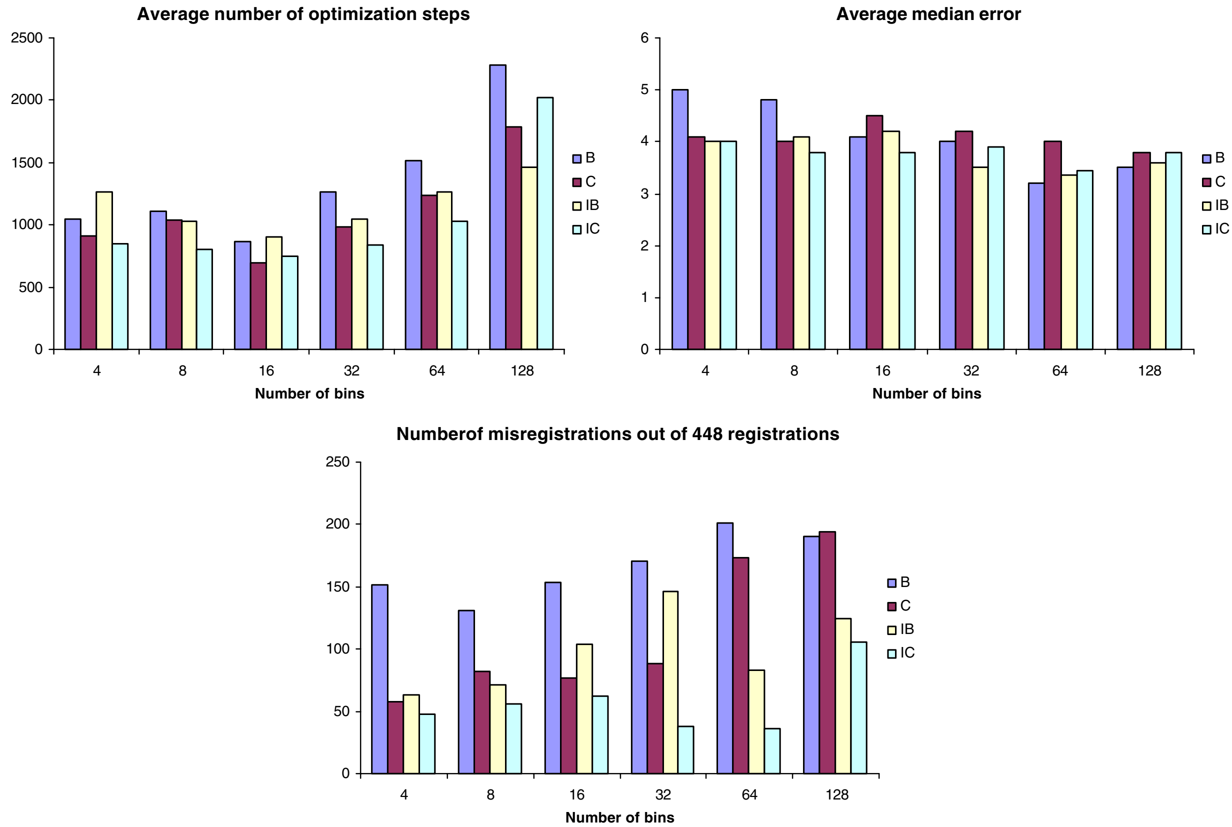

Fig. 2. PET-MRI registration results. Comparison of the number of optimization steps, median error and the number of misregistrations of binning without inhomogeneity correction (B), *k*-means clustering without inhomogeneity correction (C), binning with inhomogeneity correction (IB) and *k*-means clustering with inhomogeneity correction (IC). The errors are given in mm and both the errors and the average number of optimizations steps are computed for the non-misregistrations only.

unsuccessful registrations makes it hard to draw conclusions on accuracy related to robustness of the methods; including misregistrations would make the measures highly variant because of the less-than-robust performance of some methods (Fig. 2, the number of misregistrations). The present presentation allows conclusions on the robustness of the methods used, and whether the change in robustness has any impact on the accuracy and speed of computation of successful registrations.

For all methods and bin sizes the number of bins does not greatly influence the average median error, but for the average number of optimization steps there is a clear difference. There is a local minimum when using 16 bins and from 32 to 128 bins there is an increase in the average number of optimization steps. Using 16 bins and clustering without inhomogeneity correction yields the best results with respect to the average number of optimization steps.

For all bin sizes except 128 the use of clustering yields a reduction in the number of misregistrations regardless of the inhomogeneity correction and with comparable accuracy or increase in computational time. The use of inhomogeneity correction decreases the number of misregistrations regardless of the binning method for all bin sizes. The combined use of inhomogeneity correction with clustering yields the lowest number of misregistrations for any bin size. Using 32 or 64 bins with inhomogeneity correction and *k*-means clustering yields the smallest number of misregistrations.

4.2. CT-MRI

For all experiments the gold standard transformation was used as initial transformation. An offset transformation was added by using a rotation offset of $-5°$ or $+5°$ or -10 or $+10$ mm translation for each coordinate axis. All offset combinations were used resulting in 64 different starting positions for each patient registration totaling 960 experiments for each bin size.

The number of optimization steps needed to reach registration, the average of the median errors and the number of misregistrations for CT-MRI with and without inhomogeneity correction for all starting positions are compared for equidistant binning and *k*-means clustering in Fig. 3.

There is no discernable pattern in the average number of optimization steps with respect to the bin size. However, note that the number of misregistrations influences the average number of optimization steps. The use of clustering regardless of the inhomogeneity correction

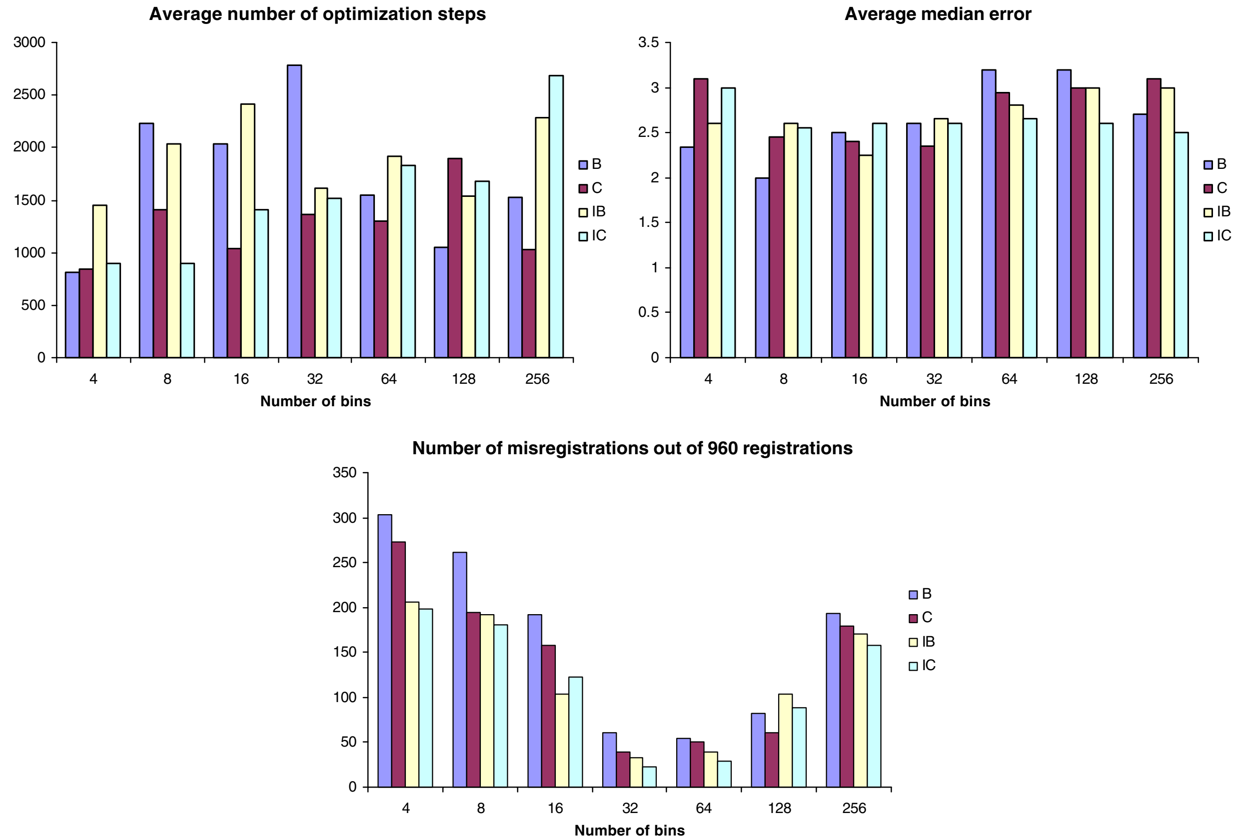

Fig. 3. CT-MRI registration results. Comparison of the number of optimization steps, median error and the number of misregistrations of using binning without inhomogeneity correction (B), using *k*-means clustering without inhomogeneity correction (C), using binning with inhomogeneity correction (IB) and *k*-means clustering with inhomogeneity correction (IC). The errors are given in mm and both the errors and the average number of optimizations steps are computed for the non-misregistrations only.

yields a decrease in the number of optimization steps when using 8, 16, 32 or 64 bins.

The errors were comparable if a correct CT-MRI registration was found. The use of clustering instead of equidistant binning without inhomogeneity correction yields a decrease in the number of misregistrations for all bin sizes. When using inhomogeneity correction clustering decreases the number of misregistrations for all bin sizes except 16. When using both inhomogeneity correction and clustering the number of misregistrations is minimum for all bin sizes except 16. The number of misregistrations is lowest when combining clustering and inhomogeneity correction at 32 bins.

Moreover there is a clear relation between the number of bins and the number of misregistrations which are minimal at 32 or 64 bins for all methods. For all methods the number of misregistrations decreases when increasing the number of bins from 4 to 32 bins and increases when increasing the number of bins for 64 to 256 bins.

4.3. Artificial T1 MRI–T2 MRI

To further investigate the effects of inhomogeneity correction, bin size and binning method with respect to misregistrations and optimization steps, experiments were performed on a T1 and a T2 artificial MRI volume. By using artificial images (Brainweb, (http://www.bic.mni.mcgill.ca/brainweb/)) we are able to control the amount of image noise and inhomogeneity. When using low amounts, all methods perform well in terms of accuracy and robustness. We therefore use relatively high levels to show the differences between the methods.

For all experiments the gold standard transformation was used as initial transformation. An offset transformation was added by using a rotation offset of $-10°$, $0°$ or $+10°$ or -20, 0 or $+20$ mm translation for each coordinate axis. All offset combinations were used resulting in 729 different starting positions. An artificial T1 image was registered to an artificial T2 image with 40% intensity non-uniformity and also to the same image after inhomogeneity correction. In both cases the T2 image had 30% noise.

The number of optimization steps needed to reach registration with and without inhomogeneity correction for all starting positions is compared for equidistant binning and *k*-means clustering in Fig. 4. If a correct registration was found the error was below 0.1 mm for all cases.

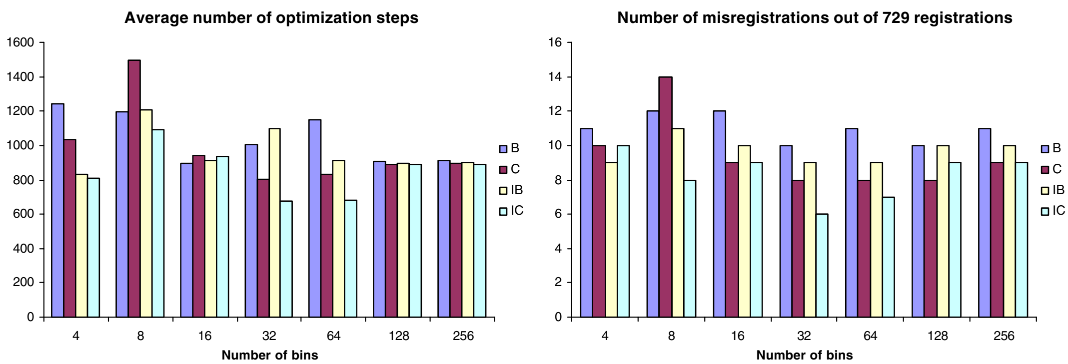

Fig. 4. Comparison for MRI–MRI registration of the number of optimization steps and the number of misregistrations using binning without inhomogeneity correction (B), k-means clustering without inhomogeneity correction (C), binning with inhomogeneity correction (IB) and k-means clustering with inhomogeneity correction (IC). For correct registrations the median error was below 0.1 mm.

Clustering and inhomogeneity correction have no effect on the registration with respect to the registration quality since, if a solution was found the error was below 0.1 mm. However, this does not hold for the number of misregistrations and optimization steps.

The effects of clustering and inhomogeneity correction on the number of optimization steps depend on the number of bins used. When using 16, 128 or 256 bins there are no clear differences between the four methods. When using 32 or 64 bins clustering yields the best results.

The use of clustering instead of binning yields a reduction of the number of misregistrations with or without inhomogeneity correction for all bin sizes except 4 and 8 bins. Using 32 bins with clustering and inhomogeneity corrections yields the best results with respect to the number of misregistrations.

5. Conclusions and discussion

In summary, the combined use of k-means clustering and inhomogeneity correction improves the robustness of PET-MRI and CT-MRI brain image NMI-based registration for all histogram bin sizes tested.

For most bin sizes, the use of either enhancement, k-means clustering or inhomogeneity correction, also results in robustness improvement. The effect is most pronounced, i.e. the least number of misregistrations occur, when using both methods using 32 or 64 bins.

The same results hold for artificial MRI–MRI image registration: the combined use of inhomogeneity correction and k-means clustering when using 32 or 64 bins yields a reduction in the number of misregistrations without loss of accuracy.

Mutual information based registration assumes that similar tissues correspond to similar within-image intensities and these intensities correspond to regions of similar within-image intensities (but probably with another between-images intensity value) in another image modality. Inhomogeneity disrupts this relation; due to a gradient similar tissue can correspond to a wide range of different intensities. We have shown that by correcting for inhomogeneities the registration process can be enhanced. Furthermore, the generally used equidistant binning method can also undermine the assumption by dividing similar tissue over different intensities.

Using the same number of bins regardless of the modalities seems illogical: a good intensity estimation of a CT image should typically contain fewer bins than a good intensity estimation of an MR image. We suspect a low dynamic range to be a factor and so is specific anatomy. Based on the current findings we were unable to establish an a priori rule for the number of bins with respect to anatomy and modality. However, experiments indicate choosing a specific number of bins for each patient does not improve the registration any further.

The method presented in this paper is an example of combining segmentation with registration, an approach previously investigated by several other authors (Studholme, 1997; Seghers et al., 2004). Experiments have shown improvements, however a theoretical approach of the combination of both registration and segmentation as shown in (Studholme, 1997) could help further research along this line.

Acknowledgements

Several of the images were provided by the project, "Retrospective Image Registration Evaluation", National Institutes of Health, Project No. 8R01EB00-2124-03, Principal Investigator, J. Michael Fitzpatrick, Vanderbilt University, Nashville, TN, United States.

References

Bellman, 1962. Dynamic programming treatment of the travelling salesman problem. Journal of Association for Computing Machinery 9 (1), 61–63.

Collignon, A., 1998. Multi-modality medical image registration by maximization of mutual information. Ph.D. Thesis, Catholic University of Leuven, Leuven, Belgium.

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

Cocosco, C.A., Kollokian, V., Kwan, R.K.-S., Evans, A.C., 1997. BrainWeb: Online Interface to a 3D MRI Simulated Brain Database. In: Proceedings of the Third International Conference on Functional Mapping of the Human Brain, NeuroImage, 5(4), S425. Available from: <http://www.bic.mni.mcgill.ca/brainweb/>.

http://www.vuse.vanderbilt.edu/~image/registration.

Jain, A.K., Dubes, R.C., 1988. Algorithms for Clustering Data. Prentice Hall Reference Series.

Knops, Z.F., Maintz, J.B.A., Viergever, M.A., Pluim, J.P.W., 2003. Normalized mutual information based registration using *K*-means clustering based histogram binning. In: Sonka, M., Fitzpatrick, J.M. (Eds.), Proceedings of SPIE Medical Imaging, Image Processing, SPIE Press, Washington, 4(23), 1072–1080.

Knops, Z.F., Maintz, J.B.A., Viergever, M.A., Pluim, J.P.W., 2003b. Normalized mutual information based PET-MR registration using *k*-means clustering and shading correction. In: Gee, J.C., Maintz, J.B.A., Vannier, M.W. (Eds.), Proceedings of the Second International Workshop on Biomedical Image Registration, Lecture Notes in Computer Science, vol. 2717. Springer-Verlag, Berlin, pp. 31–39.

Knops, Z.F., Maintz, J.B.A., Pluim, J.P.W., Viergever, M.A., 2004. Optimal one-dimensional *k*-means clustering using dynamic programming. Internal report. Utrecht University, Department of Computer Science, UU-CS-2004-041.

Likar, B., Maintz, J.B.A., Viergever, M.A., Pernuš, F., 2000. Retrospective shading correction based on entropy minimization. Journal of Microscopy 197 (3), 285–295.

Likar, B., Viergever, M.A., Pernuš, F., 2001. Retrospective correction of MR intensity inhomogeneity by information minimization. IEEE Transactions on Medical Imaging 20 (12), 1398–1410.

MacQueen, J., 1967. Some methods for classification and analysis of multivariate observations. In: Le Cam, M., Neyman, Jerzy (Eds.), Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability. University of California Press, pp. 281–297.

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

Press, W.H., Flannery, B.P., Teukolsky, S.A., Vetterling, W.T., 1992. Numerical Recipes in C. Cambridge University Press, Cambridge, UK.

Seghers, D., D'Agostino, E., Maes, F., Vandermeulen, D., Suetens, P., 2004. Construction of a brain template from MR Images using state-of-the-art registration and segmentation techniques. In: Barillot, C., Haynor, D.R., Hellier, P. (Eds.), Proceedings of Medical Image Computing and Computer-assisted Intervention – MICCAI, Lecture Notes in Computer Science, vol. 3216. Springer-Verlag, Berlin, pp. 696–703.

Studholme, C., 1997. Measures of 3D medical image alignment, Ph.D. Thesis, University of London, London, UK.

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

Viola, P., 1995. Alignment by maximization of mutual information, Ph.D. Thesis, Massachusetts Institute of Technology, Boston, MA, USA.

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

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

West, J., Fitzpatrick, J.M., Wang, M.Y., Dawant, B.M., Maurer Jr., C.R., Kessler, R.M., Maciunas, R.J., et al., 1997. Comparison and evaluation of retrospective intermodality image registration techniques. Journal of Computer Assisted Tomography 21 (4), 554–566.