Multimodal Image Registration for Efficient Multi-resolution Visualization

Joerg Meyer

Department of Electrical Engineering and Computer Science
644E Engineering Tower, Irvine, CA 92697-2625
jmeyer@uci.edu

Summary. Arising from the clinical need for multimodal imaging, an integrated system for automated multimodal image registration and multi-source volume rendering has been developed, enabling simultaneous processing and rendering of image data from structural and functional medical imaging sources. The algorithms satisfy real-time data processing constraints, as required for clinal deployment.

The system represents an integrated pipeline for multimodal diagnostics comprising of multiple-source image acquisition; efficient, wavelet-based data storage; automated image registration based on mutual information and histogram transformations; and texture-based volume rendering for interactive rendering on multiple scales.

Efficient storage and processing of multimodal images as well as histogram transformation and registration will be discussed. It will be shown how the conflict of variable resolutions that occurs when using different modalities can be resolved efficiently by using a wavelet-based storage pattern, which also offers advantages for multi-resolution rendering.

1 Introduction

This chapter describes a set of algorithms that form an integrated system for automated image processing, multimodal image registration and multi-source volume rendering. The system has been developed to enable simultaneous processing and rendering of image data from multiple medical imaging sources.

A complete system (figure 1) consists of an integrated pipeline for multimodal diagnostics. The components include multiple-source image acquisition; automated image registration based on histogram transformations and optimization of normalized cross-correlation; efficient, wavelet-based data storage and transmission; and 3D-texture-based, multi-resolution volume rendering for interactive visualization.

The focus of this chapter is on histogram transformation and registration. A short section in the end will discuss efficient storage, processing and rendering of multimodal images. It will be shown how different modalities, even when they express various tissue types differently or when their resolution differs, can be combined

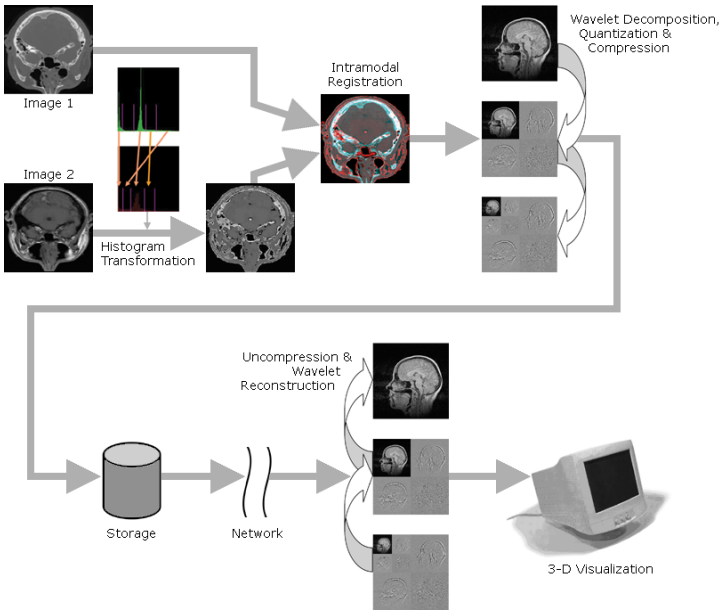

Fig. 1. Rendering pipeline for multimodal images.

and stored efficiently by using a histogram transformation and a wavelet-based storage pattern. The latter also offers advantages for network-based data transmission, and for visualization on multiple scales. In the end, references will be given for further study.

This chapter is structured as follows. Section 2 describes the difference between intermodal and intramodal registration and provides examples of both techniques. Section 3 describes the pre-processing step of segmenting the images in each modality and correlating the scalar values for each cluster. Section 4 outlines the registration process using cross-correlation as a similarity measure, and section 5 shows the results obtained using this method. Section 6 explains how multi-resolution storage patterns and three-dimensional, texture-based volume rendering can be used to display multimodal scans in real time. Section 7 summarizes the findings and outlines future research goals and potential directions.

2 Background

2.1 Intermodal vs. Intramodal Registration

This section addresses the problem of automatic registration of CT and MR images. The task is non-trivial due to the fact that particular tissue types are represented in dissimilar ways in different modalities and may appear in a different order in the gray level spectrum. For example, bone is visually represented as bright white (high X-ray absorption) in CT scans (figure 9(a)), and as a black band (low proton

density) in MRI scans (figure 9(b)). Hence, the correlation between corresponding pixels in each scan is not linear (figure 9(c)). This makes the task of aligning them, without the use of stereotactic frames or markers, difficult.

As an alternative approach to intermodality registration, a method that uses feature segmentation based on intensity histograms [Gon02] of the two images is proposed here. The idea is that if one of the scans is transformed to match the gray value distribution of the other scan, then the similarity measures used in *intramodality* registration techniques can be used for *intermodality* registration. Section 3 outlines techniques to achieve a reasonable segmentation of the various tissue types and correlation of the same to achieve a robust registration. The fact that the algorithm in the current implementation uses several different tissue types to control the alignment process and the correlation computations adds additional robustness to the approach.

Combining multiple scans of the same specimen into a single image or volume has numerous advantages in diagnostics, pre-surgical planning and radiation therapy planning, especially when multiple modalities depicting various physical properties of the scanned sample are involved. Some tissue types may be better visible in one type of scan, others may only appear in another scan.

Similar scans (e.g. multiple CT scans) can be registered by using a simple differential function as a similarity measure (*intramodal registration*). Combining images taken from different modalities (*intermodal registration*) is a much harder problem. The main reason is the fact that the various tissue types that occur in the images are represented in different ways in different modalities and may appear in a permutated order in the gray level spectrum. For example, a linear transfer function maps the high X-ray absorption values of bone to an intensity value that is visually represented as a bright white, whereas the same linear transfer function would produce a black pixel for the same location in an MRI scan due to the low proton density of bone matter. Hence, the correlation between corresponding pixels in each scan is not linear over the entire range of values. Moreover, the intensity ranges associated with a particular tissue type do not appear in the same order in a histogram. This makes the task of aligning two images from different modalities difficult.

2.2 Registration Methods

Three different, classical registration methods with different prerequisites and application areas are presented here. The *Global Difference* method is based on the assumption that intensities in the two participating images are linearly correlated. This assumption makes the method not suitable for intermodal registration, because absolute values, the spacing between these values in an intensity histogram, and the order of clusters in such a histogram vary due to the difference in the way material properties and the physical quantities measured by a particular scanning method differ.

The second method is called *Geometric Features*. It compares outlines of objects, which can be identified by extracting intensity gradients using a Laplacian or high-pass filter. This method is based on the assumption that tissue gradients are expressed similarly in both images. Unfortunately, this is not the case for the same reason as for the absolute values, i.e., because of the differences in the physical quantities being measured by the two incompatible image acquisition devices.

A third method is known as *Mutual Information*. This technique computes the statistical dependence or information redundancy between image intensities of corresponding pixels in both images. No specific assumptions about feature expression are being made. This method is suitable for intermodal registration and has been widely studied in the literature [Mae97].

Mutual information (MI) works for different modalities that contain the same information, but expressed differently. It is typically computed as illustrated in equation 1.

$$MI(I_1, I_2) = \sum_{i_2(x,y) \in I_2} \sum_{i_1(x,y) \in I_1} h(i_1(x,y), i_2(x,y)) \cdot \log \frac{h(i_1(x,y), i_2(x,y))}{h(i_1(x,y)) \cdot h(i_2(x,y))} \quad (1)$$

I_1 and I_2 describe the two images, and i_1 and i_2 the respective pixels in these images. $h(i)$ refers to the respective histograms, while $h(i_1, i_2)$ refers to the mutual histogram (products of pixel counts per intensity value). The goal is to maximize MI.

The idea of the method presented in this chapter is that methods suitable only for *intramodal registration* (same imaging device type) can be adapted to work also for *intermodal registration* (different imaging device types). The method can be efficiently implemented using a look-up table and, depending on the implementation of the intramodal registration algorithm, is therefore potentially faster than the computationally expensive MI method.

2.3 Advanced Methods

Numerous attempts to accelerate the alignment process, ranging from manual to semi-automatic to fully automatic methods [Els94, Els95, Stu96, Pen98, Hil01, Jen02, LoC03], have been proposed in the literature. These methods include histogram matching [Hua98], texture matching [Ash95], intensity cross-correlation [Mai98], kernel-based classification methods [Cri00], information divergence minimization [Sto98] and optical flow matching [Lef01]. A good survey can be found in [Bro92].

More recently developed fully automated methods essentially revolve around entropy [Mey97, But01, Zhu02] and mutual information (MI) [Vio95, Wel96, Plu03, Wel96]. Even though they are more convenient than manual or semi-automatic techniques, they are still time consuming and computationally expensive. The general consensus on MI-based techniques is that they provide the best results in terms of alignment precision and are generally consistent with manual, expert-based alignment procedures. However, the loss of spatial coherence during the registration, which is due to the different physical quantities being measured and the resulting variations in the feature boundaries and gradients, is a potential cause of inaccurate registration.

2.4 Using Intramodality Techniques for Intermodality Registration

As described before, we propose an alternate method for registering CT and MRI scans, usually requiring *intermodality* techniques. The method described here employs *intramodality* methods to accomplish the alignment task. Our approach is

rather intuitive and inexpensive. We hypothesize that if we convert one of the scans (e.g. MRI) to match the other (e.g. CT) in terms of the scalar value representation of various tissue types that occur in the images, then by applying a similarity measure based on the correlation between corresponding features and globally optimizing this parameter, we can achieve alignment of the two different modalities at a relatively cheaper cost, providing results comparable in quality to those of intermodality registration techniques.

The task of converting the representation of the various tissue types present in one scan into the representation given by another scanning technology or modality involves the use of feature segmentation followed by histogram matching [Hua98]. By sorting the intensities of each image into clusters corresponding to the intensity ranges of characteristic features, we can approximately segment the scans into the various compounds or tissue types. These individual segmentation clusters can then be mapped onto relevant bins from the second scan. A histogram matching of these segmented images, which includes a permutation and adaptation of the bin widths in the histogram, generates a transfer function that defines a mapping between the original scan and the alternate scan.

3 Segmentation Using Histogram Binning

In order to correlate the individual representations of the various anatomical structures in each scan, we use intensity histograms. Figure 10(a) shows the intensity profile of a CT scan (data specification: 512 x 512 grayscale image). Figure 10(b) shows the intensity profile of an MRI scan (data specification: T1-weighted, 256 x 256 grayscale image) of a human head. All images and histograms were stored in an 8-bit format (value range [0...255]). The adjustment of the resolution (scaling factor: 2) is accomplished as a by-product of the wavelet reconstruction described in section 6, which is not carried out to the highest detail level. Other scaling factors can be obtained using bi-linear interpolation of the next higher resolution step. The non-standard Haar wavelet decomposition [Mey03] provides resolution steps in multiples of two in each dimension.

Histogram binning is an approximation technique based on known expressions of particular tissue types in various modalities. Depending on the number of expected distinct features (i.e. tissue types), we create bins of intensity clusters. For example, we have one bin to store all the high intensity pixels in a CT image which are characteristic of skull bone (high X-ray absorption). This would map onto the bin containing the low intensities from the skull in the MRI scan (low hydrogen proton concentration in bone). Similarly, the gray matter, white matter and cerebral spinal fluid (CSF) can be stored in one bin corresponding to the uniform gray region of brain matter in the CT.

The bins thus correspond to a partitioning of the image into characteristic tissue expressions. Initial bins were obtained and refined using user interaction, and a table of mapping values was generated (table 1).

Having thus created the desired bins, we then generate a transfer function to carry out the mapping between each segment (bin) of the MRI and the CT scan based on table 1, using the principles of histogram matching [Hua98]. This process of selective histogram matching ensures a good approximation.

Table 1. Mapping of MRI intensities to corresponding CT intensities for each feature space.

CT intensity	MR intensity	Feature
0 – 9	0 – 1	Background
10 – 60	175 – 255	Adipose
61 – 140	56 – 119	Soft tissue
141 – 255	2 – 55	Bone

Figure 11 shows a graphical representation of the bins and a typical permutation. The arrows indicate the mapping shown in table 1. It becomes obvious that bone, for instance, is represented differently in different modalities, and that the width of the bins varies. Therefore, the mapping is not linear. A piece-wise linear transfer function is designed to perform the permutation and scaling of the histogram bins.

After applying the transfer function, the MRI image now contains intensity values that correspond to those found in the CT scan. This means that the various tissue types are now represented by a similar intensity value. Figure 12(a) shows the original superimposed histograms of the CT and the MRI scan before application of the transfer function. The image shows that there is minimal overlap, which makes image registration of features that are expressed differently difficult. Figure 12(b) shows the superimposed histograms of the CT and the transformed MRI scan. Obviously, there is a much better match in the overall shape of the histogram, indicating that the two images, even though generated by two different modalities, will yield a much better registration of individual features. Figures 2(a) and (b) show the corresponding images, i.e., the original MRI image, and the new MRI image after application of the transfer function. Figure 2(b) is now similar to the original CT image, which is shown in figure 2(c) for comparison.

The quality of the transformation can be evaluated by subtracting either the histograms or the images. Using a low-pass filter on the images or a quantization on the histogram is necessary to obtain a reasonable measure of the quality, because due to the difference in the image acquisition process, the individual pixels or bin distributions will still be very different, but the average distribution is similar.

Our registration algorithm takes as input a file that contains the segmentation information, i.e., the number of clusters or bins followed by the intensity range for each cluster. The algorithm then optimizes a global similarity measure based on correlation, as described in the following section.

4 Registration

The registration was done using an intensity-based similarity measure and an optimization algorithm for various rotation angles and translation vectors. Optimization of a normalized cross-correlation (NCC) function using Powell's multi-dimensional line maximization [Jac77] formed the basis of the similarity measure.

$$R = \frac{\sum_{k=0}^{n}[(r(k) - \mu(r)) - (f(k) - \mu(f))]}{\sqrt{\sum_{k=0}^{n}(r(k) - \mu(r))^2 - \sum_{k=0}^{n}((f(k) - \mu(f))^2}} \tag{2}$$

Equation 2 shows the NCC function. Here r and f represent the reference and the floating image, respectively. $\mu(r)$ and $\mu(f)$ represent the mean intensity values of the reference and the floating image. In our case, the CT scan is the reference image and the transformed MRI is the floating image. R represents the metric used to determine the similarity between the CT and the transformed MR image. $R = 1$ implies maximum correlation, $R = 0$ implies no correlation, and $R = -1$ implies an opposing relation between the intensities from the two modalities. Powell's multi-dimensional line maximization algorithm, which employs Brent's one-dimensional optimization algorithm, was used to compute different variants and to determine the optimum value for R in the X-Y plane.

A "sum of squared differences" technique may also be used for the registration process, since specific optimization algorithms work very efficiently for the same. However, the choice of NCC as a similarity measure was dictated by the fact that tissue-type-based clustering is not an exact science due to large variations in individual tissue expressions across patients. In addition, noise, if not removed prior to the transformation, is also transformed by the histogram matching process, has a strong effect on the "sum of squared differences" method and interferes with the registration, especially when the noise is not uniformly distributed or Gaussian. The NCC function together with a low-pass filter ensure the robustness of the algorithm.

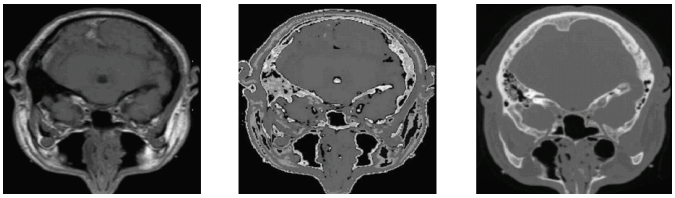

Fig. 2. (a) Original MRI image. (b) Transformed MRI image. (c) Original CT image. The transformed MRI (center) matches the geometry of the original MRI (left) and the intensity profile of the CT scan (right).

5 Results

Figure 13(a) shows the superimposed images of the CT and the original MRI scans in their initial, unregistered position. Note that the bone (bright red or bright cyan) is not aligned (white, composition of red and cyan). Figure 13(b) shows the registered set, again with no good match (few white pixels). The hollow space (transparent) in the MRI is completely filled by the skull (cyan) from CT. This means that the bone is hardly visible in the given MRI scan.

Figure 13(c) shows the same results for a transformed MRI image. Note that the bone (white, composition of cyan and red) is much better visible now. Especially after registration (figure 13(d)), the bone appears mostly white. This means that it is visible in both images (CT and transformed MRI). This property is most obvious for tissues that appear bright in the current representation (bone), but it also applies to other tissue types, as is shown in the following images. The assignment of intensity

ranges (histogram bin mapping) is arbitrary and typically chosen to match one of the two modalities, in this case the CT scan.

Figures 3 and 4 were created using thresholding of the pixels that indicate the highest structural similarity (bone only). As stated before, not only the pixels with the highest values (typically bone) can be used for a correspondence analysis. Other tissue types can be used as well, as shown in figures 5 and 6 (all tissue types).

The last image in each series (figures 4(b) and 6(b)) shows the largest number of pixels usable for correspondence analyses. An analysis of the pixel count indicates that the transformation using transfer functions and histogram permutation increases the number of pixels that can be used for correspondence analyses, therefore increasing the accuracy and robustness of the method.

The algorithm can be very efficiently implemented using look-up tables so that it was possible to run it on a laptop computer. In comparison to other intensity-based techniques for intermodality registrations that use entropy or MI calculations, NCC computations are less expensive. The pre-processing step is also inexpensive and is independent of the alignment routine. Performance data of the various algorithms is difficult to compare due to large variations in the implementations and applications. Therefore, absolute timing measurements were taken, which may serve as a guideline for similar implementations.

The computation of the transformation of the original MRI to its CT-look-alike took approximately 10 milliseconds, and the alignment algorithm took 3.25 seconds for a single slice (approx. 6 minutes for a typical data set of 113 slices) on an IBM R40 Thinkpad with 256MB RAM and 1.4 GHz PentiumM processor. The graphics card is an ATI Mobility Radeon with 16MB memory.

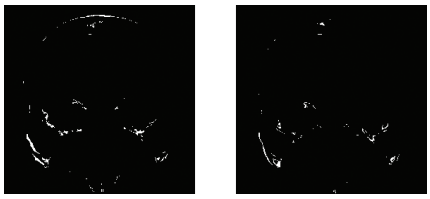

Fig. 3. CT and original MRI. (a) unregistered, (b) registered. (Bone only.)

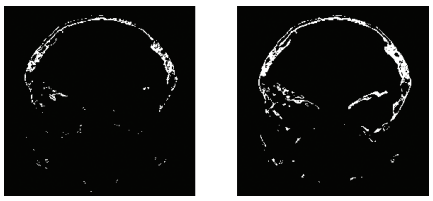

Fig. 4. CT and transformed MRI. (a) unregistered, (b) registered. (Bone only.) White regions indicate matched bone regions expressed similarly in the CT and the transformed MRI scan.

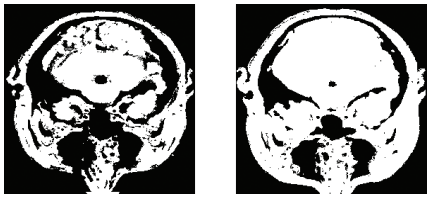

Fig. 5. CT and original MRI. (a) unregistered, (b) registered. (All tissue types.)

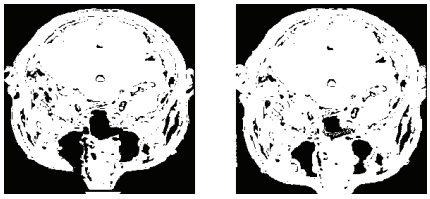

Fig. 6. CT and transformed MRI. (a) unregistered, (b) registered. (All tissue types.) White regions indicate matched regions of all tissue types expressed similarly in the CT and the transformed MRI scan. The right image shows an almost perfect match (all tissue types).

6 Multimodal Image Acquisition and Multi-resolution Visualization

This section will give a brief overview of the image acquisition, storage, processing and rendering pipeline for multimodal images. The registration algorithm described in the previous sections will be used in this context.

When images are acquired using two different medical scanners, registration of the three-dimensional volume data sets is always required. Even if the scanners are perfectly calibrated, if the geometry of the scanned area is similar and the scanned specimen is the same, there is still a margin of error and deviation due to positioning of the specimen in the scanner, shifting of tissues (e.g., movements of a brain inside a skull), and geometric distortions of the scanner. Some of these deviations and deformations require linear, others non-linear transformations. Also, the resolution must be adjusted using an interpolation or multi-resolution technique (see section 3).

The image registration described earlier in this book chapter applies to the two-dimensional case, which requires at least the body axes to be aligned in the scanners. This constraint can be dropped if the algorithm is translated into the three-dimensional domain, which is pretty straightforward. In this case, the normalized cross-correlation function (equation 2) must be computed for three-dimensional images r and f.

Since three-dimensional data sets, especially when obtained from two independent storage devices, can become very large, efficient data compression is required for data transmission over limited-bandwidth networks and archival storage. Since there is no correlation in the absolute data values obtained from two different modalities, the compression algorithm cannot take advantage of similarities in the features. Data compression usually becomes more efficient when smaller values are involved,

increasing the likelyhood of finding similar or identical values. Quantization helps to increase the number of identical values and therefore improves the compression ratio by sacrificing some of the image quality.

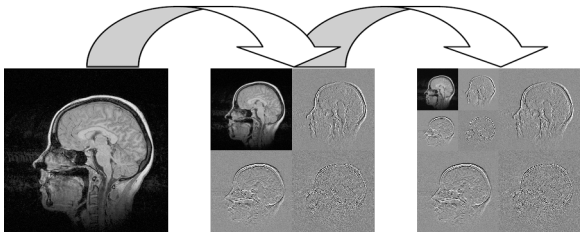

Fig. 7. Wavelet decomposition.

Two-dimensional wavelet decomposition (figure 7) works by separating an image into low-frequency and high-frequency coefficients [Mey03]. Three-dimensional wavelet compression algorithms have been shown to provide an efficient way to compress volume data, as they take advantage of the similarity of neighboring voxels in three spatial directions. Wavelet decomposition algorithms become more efficient in higher dimensional data sets if the data is correlated in each dimension. By extending the three-dimensional wavelet compression into the fourth dimension, where the fourth dimension represents the various modalities (transformed into a common representation of a reference modality), the compression algorithm that follows the wavelet decomposition can take advantage of the improved correlation between the two modalities after histogram transformation (figure 8).

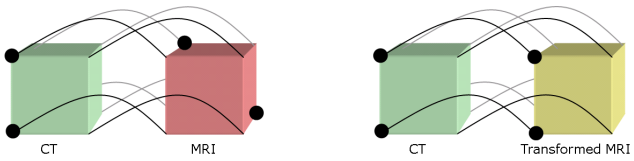

Fig. 8. Four-dimensional, non-standard Haar Wavelet decomposition. Detail coefficients (differences between adjacent voxel values) in the fourth dimension (modality number) are typically smaller for CT–transformed MRI pairs (right) than for CT–original MRI pairs (left) and therefore yield better compression ratios.

In this figure, the dots at the voxel corners indicate high values, while the other corners represent low values. Please note that the voxel corners in the CT–transformed MRI pair on the right are more similar to each other than the voxel corners in the CT–original MRI pair on the left. This is due to the histogram transformation and turns out to be an advantage when a non-standard Haar wavelet decomposition is applied. Since the two voxels are similar, the average of the two voxels (low-frequency component) is similar to both voxel values, while the difference between the two voxels (high-frequency component) is small. Detail coefficients

(high-frequency components) can be quantized and stored with fewer bits, and small coefficients can be ignored.

The efficiency increases with the dimensionality of the data set, because the average (low-frequency component) is computed between two elements in 1-D, between $2 * 2$ (four) elements in 2-D, $2 * 2 * 2$ (eight) elements in 3-D, and so forth. If the modality number is considered the fourth dimension, an average of sixteen elements can be computed and represented as a single number. The detail coefficients (high-frequency components) are typically small if the modalities are represented similarly and therefore compress well in a subsequent compression algorithm (run-length encoding with quantization, entropy encoding, Lempel-Ziv-Welch, etc.)

A wavelet decompression and reconstruction algorithm (the inverse of the wavelet decomposition algorithm) can be used to obtain volumetric images at multiple levels of detail (figure 7). The advantage of such a multi-resolution reconstruction is that a low resolution image can be displayed instantly, while increasing levels of detail can be reconstructed sequentially as they become available (either as streaming data from a network or as high-resolution data from an archive). The details of the wavelet decomposition and reconstruction algorithm as well as the three-dimensional, texture-based rendering algorithm are described in [Mey03].

7 Summary

This chapter explained how histogram information of a grayscale image can be adapted based on another image using previous knowledge about the bins of the reference image histogram. By transforming one image into the value range of the reference image using a non-linear mapping, *intramodal* registration techniques, as opposed to computationally more expensive *intermodal* techniques, can be used to align the images. Optimization of a normalized cross-correlation function (NCC) is an efficient way to implement an intramodal image registration algorithm. Histogram transformation also benefits file storage and data compression without changing the information content in the image, allowing for simultaneous storage and rendering of multiple modalities.

8 Acknowledgements

This project was funded in part by the National Institute of Mental Health (NIMH) through a subcontract with the Center for Neuroscience at the University of California, Davis (award no. 5 P20 MH60975), and by the National Partnership for Advanced Computational Infrastructure (NPACI), Interaction Environments (IE) Thrust (award no. 10195430 00120410).

The author would like to thank Ruparani Chittineni for contributing some of the code, images and descriptions.

References

[Ash95] Ashley J., Barber R., Flickner M., Lee D., Niblack W., and Petkovic D.: Automatic and semiautomatic methods for image annotation and retrieval in qbic. In: Proc. SPIE Storage and Retrieval for Image and Video Databases III, 24–35 (1995)

[Bro92] Brown L. G.: A survey of image registration techniques. ACM Computing Surveys 24, **4**, 325–376 (1992)

[But01] Butz T. and Thiran J.: Affine registration with feature space mututal information. In: Lecture Notes in Computer Science 2208: MICCAI 2001, Springer-Verlag Berlin Heidelberg, 549–556 (2001)

[Cri00] Cristiani N. and Shaw-Taylor J.: Suport Vector Machines and other kernel-based learning methods. Cambridge U. Press (2000)

[Els94] Van den Elsen P.A., Pol D. E., Sumanaweera S.T., Hemler P. F., Napel S., Adler J. R.: Grey value correlation techniques used for automatic matching of CT and MR brain and spine images. Proc. Visualization in Biomedical Computing, SPIE 2359, 227–237 (1994)

[Els95] Van den Elsen P. A., Maintz J. B. A., Pol D. E., Viergever M. A.: Automatic registration of CT and MR brain images using correlation of geometrical features. IEEE Transactions on Medical Imaging 14, **2**, 384–396 (1995)

[Gon02] Gonzalez R. C., Woods R. E.: Digital Image Processing. Prentice Hall (2002)

[Hil01] Hill D., Batchelor P., Holden M., and Hawkes D.: Medical image registration. Phys. Med. Biol., **26**, R1–R45 (2001)

[Hua98] Huang J., Kumar S., Mitra M., and Zhu W.: Spatial color indexing and applications. In: Proc. of IEEE International Conf. Computer Vision ICCV '98, Bombay, India, 602–608 (1998)

[Jac77] Jacobs D. A. H.: The state of the art in numerical analysis. Academic Press, London (1977)

[Jen02] Jenkinson M., Bannister P., Brady M., and Smith S.: Improved methods for the registration and motion correction of brain images. Technical report, Oxford University (2002)

[Lef01] Lefébure M. and Cohen L.: Image registration, optical flow and local rigidity. J. Mathematical Imaging and Vision, (14) 2, 131–147 (2001)

[LoC03] Lo C. H., Guo Y., Lu C. C.: A binarization approach to CT-MR registration using Normalized Mutual Information. Proc. IASTED Signal and Image Processing, 399 (2003)

[Mae97] Maes F., Collignon A., Vandermeulen D., Marchal G., Suetens P.: Multimodality image registration by maximization of mutual information. IEEE Transactions on Medical Imaging 16, **2**, 187–198 (1997)

[Mai98] Maintz J. B. and Viergever M.: A survey of medical image registration. Medical Image Analysis, (2) 1, 1–36 (1998)

[Mey97] Meyer C. R., Boes J. L., Kim B., Bland P. H., Zasadny K. R., Kison P. V., Koral K. F., Frey K. A., and Wahl R. L.: Demonstration of accuracy and clinical versatility of mutual information for automatic multimodality image fusion using affine and thin-plate spline warped geometric deformations. Medical Image Analysis, (1) 2, 195–206 (1997)

[Mey03] Meyer J., Borg R., Takanashi I., Lum E. B., and Hamann B.: Segmentation and Texture-based Hierarchical Rendering Techniques for Large-scale Real-color Biomedical Image Data. In: Post F. H., Nielson G. H., Bonneau G.-P., eds., Data Visualization - The State of the Art, Kluwer Academic Publishers, Boston, 169–182 (2003)

[Pen98] Penney C., Weese J., Little J., Hill, D., and Hawkes, D.: A comparison of similarity measures for used in 2-D-3-D medical image registration. IEEE Trans. on Medical Imaging, (17) 4, 586–595 (1998)

[Plu03] Pluim J. P. W., Maintz J. B. A., Viergever M. A.: Mutual Information based registration of medical images: a survey. IEEE Transactions on Medical Imaging 22, **8**, 896–1004 (2003)

[Sto98] Stoica R., Zerubia J., and Francos J. M.: The two-dimensional wold decomposition for segmentation and indexing in image libraries. In: Proc. IEEE Int. Conf. Acoust., Speech, and Sig. Proc., Seattle (1998)

[Stu96] Studholme C., Hill D. L. G., Hawkes D. J.: Automated 3-D registration of MR and CT images of the head. Medical Image Analysis 1, **2**, 163–175 (1996)

[Vio95] Viola P. and Wells III W. M.: Alignment by maximization of mutual information. In: Proceedings of IEEE International Conference on Computer Vision, Los Alamitos, CA, 16–23 (1995)

[Wel96] Wells W. M., Viola P., Atsumi H., Nakajima S., Kikinis R.: Multi-modal volume registration by maximization of mutual information. Medical Image Analysis 1, **1**, 35–51 (1996)

[Zhu02] Zhu Y. M.: Volume image registration by cross-entropy optimization. IEEE Transactions on Medical Imaging 21, 174–180 (2002)

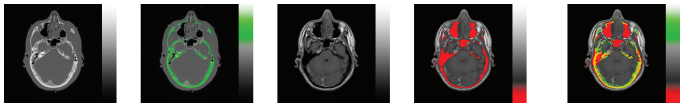

Fig. 9. Bone. (a) CT scan (green), (b) MRI scan (red), (c) superimposed (yellow).

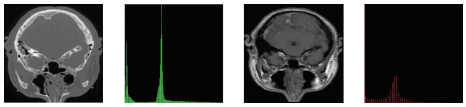

Fig. 10. (a) CT scan and (b) MRI scan with intensity histograms.

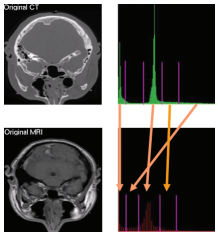

Fig. 11. Mapping of CT histogram (green) onto MRI histogram (red). Bins (purple).

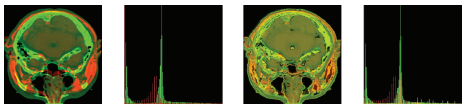

Fig. 12. CT (green) with MRI (a) original (b) transformed. Histogram in (b) shows both original (red) and transformed (yellow) MRI. Yellow is a better match.

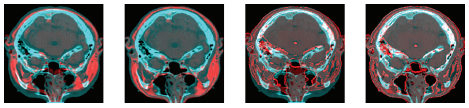

Fig. 13. (a), (c) unregistered; (b), (d) registered CT (cyan) and MRI (red). Left pair: Original MRI. Right pair: Transformed MRI. White: Matched bone.