A Novel Image Similarity Measure for Registration of 3-D MR Images and X-Ray Projection Images

Torsten Rohlfing and Calvin R. Maurer, Jr.

Image Guidance Laboratories, Department of Neurosurgery, Stanford University
300 Pasteur Drive, MC 5327, Room S-012, Stanford, CA 94305-5327, USA
{rohlfing,calvin.maurer}@igl.stanford.edu

Abstract. Most existing methods for registration of three-dimensional tomographic images to two-dimensional projection images use simulated projection images and either intensity-based or feature-based image similarity measures. This paper suggests a novel class of similarity measures based on probabilities. We compute intensity distributions along simulated rays through the 3-D image rather than ray sums. Using a finite state machine, we eliminate background voxels from the 3-D image while preserving voxels from air filled cavities and other low-intensity regions that are part of the imaged object (e.g., bone in MRI). The resulting tissue distributions along all rays are compared to the corresponding pixel intensities of the real projection image by means of a probabilistic extension of histogram-based similarity measures such as (normalized) mutual information. Because our method does not compute ray sums, its application, unlike DRR-based methods, is not limited to X-ray CT images. In the present paper, we show the ability of our similarity measure to successfully identify the correct position of an MR image with respect to a set of orthogonal DRRs computed from a co-registered CT image. In an initial evaluation, we demonstrate that the capture range of our similarity measure is approximately 40 mm with an accuracy of approximately 4 mm.

1 Introduction

Most current methods for registering three-dimensional (3-D) tomographic images to two-dimensional (2-D) projection images (e.g., X-ray fluoroscopy, electronic portal images (EPIs) in radiation therapy) make use of digitally reconstructed radiographs (DRR) computed from CT images. The physical foundations of 3-D CT and 2-D X-ray projection imaging are very similar [1]. Therefore, by casting virtual rays through a CT image, one can compute simulated projection images that resemble actual X-ray images (likewise for EPI) of the same patient in the appropriate pose. These simulated projections are compared to the real projections using standard intensity-based image similarity measures in order to achieve registration of projections and 3-D volume [2,3]. Other approaches use geometrical features, such as edges [4] or point-based landmarks

T. Dohi and R. Kikinis (Eds.): MICCAI 2002, LNCS 2489, pp. 469–476, 2002.
© Springer-Verlag Berlin Heidelberg 2002

(either anatomical or artificial) that are back-projected from the 2-D projections into 3-D space and registered using 3-D point-based algorithms [5,6]. Methods based on artificial landmarks (fiducials) are necessarily invasive. Anatomical landmarks, on the other hand, are hard to identify reliably, especially in multimodal images and 2-D projections.

Our group has recently introduced a third class of approaches to the registration of 3-D volumetric images and 2-D projections that is based neither on intensities nor on features, but instead on probabilities [7]. Using probabilistic DRRs (pDRR) and a probabilistic extension of histogram-based image similarity measures, we were able to preserve the spatial information present in volumetric images and use it during registration. An additional advantage is that pDRR computation is not based upon the physical interpretation of voxel intensities as X-ray attenuation coefficients. The method can therefore be applied in a meaningful way to tomographic images other than X-ray CT.

In the present paper, we apply our probabilistic similarity measure based on pDRR to the registration of 3-D MR images to standard DRR projection images computed from CT. Here, the deterministic DRR images serve as a model for real X-ray projection images, but with a highly accurate known pose, thanks to CT-to-MR co-registration. We also describe a method of distinguishing bone in MRI from image background on-the-fly while iterating over the voxel samples along a ray. In summary, this work is, as far as we are aware, the first to introduce a direct way of registering MR images with projection images without any segmentation or other pre-processing.

2 Methods

Probabilistic DRR. For the ray associated with the detector position $\mathbf{x}_{\text{det}}$ we define the probabilistic DRR (pDRR) as the distribution P of intensities μ sampled discretely at N uniformly-spaced locations $\mathbf{x}_i$ along this ray:

$$\text{pDRR}(\mathbf{x}_{\text{det}}, c) = P[\mu(\mathbf{x}_i) = c \mid 0 \leq i < N]. \tag{1}$$

In order to save computation time, the range of samples visited along the ray is restricted to the actual intersection of ray and 3-D image. This is achieved by computing the index I_{in} of the entry point of the ray into the volume and the index I_{out} of the exit point. This is efficiently achieved by solving a system of inequalities, originally described in an algorithm for 3-D line clipping on viewport boundaries by Liang and Barsky [8]. The probabilistic DRR can thus be equivalently rewritten as

$$\text{pDRR}(\mathbf{x}_{\text{det}}, c) = P[\mu(\mathbf{x}_i) = c \mid 0 \leq I_{\text{in}} \leq i \leq I_{\text{out}} < N]. \tag{2}$$

For a particular pose (position and orientation) of a CT image, we compute a pDRR by generating a histogram of CT intensities along each projection ray. Each pixel in the pDRR image therefore corresponds to a distribution of CT values along the ray that resulted in the projection value at that pixel. In order

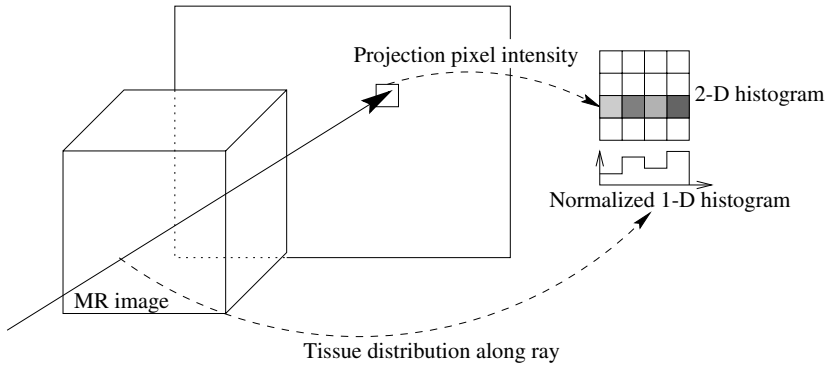

Fig. 1. Registration of MRI and projection image (e.g., fluoroscopy, EPI, DRR) using pDRR and pMI. For each projection pixel, the distribution of MR intensities along the corresponding ray is computed. The histogram is normalized to total mass 1, and added to the row in the 2-D histogram that is indexed by the intensity of the current pixel in the projection image.

to avoid interpolation, the original intensities along the ray are entered into the histogram. The proximity of each voxel to the ray is taken into account by adding to the respective histogram bin only a fractional value between 0 and 1 identical to the weight that would otherwise be used for this voxel in the interpolation (partial volume integration [9]). We will later in this paper apply the same principle in order to handle non-scalar, in our particular case probabilistic, data during the computation of histogram-based similarity measures.

Probabilistic Mutual Information. The mutual information (MI) image similarity measure [9] has been used with great success in the registration of 3-D to 3-D images [10] (single or multi modality). Based on our previous experience, we usually apply the normalized mutual information [11] (NMI) image similarity measure, which is derived from MI and appears to be less susceptible to changes in mutual image overlap. Both measures are usually computed from discrete 2-D histograms.

A 2-D histogram is a matrix $\mathbf{H}$ for which each row corresponds to a range of voxel intensities of one of the two images, and each column corresponds to a range of voxel intensities of the other image. A pair of corresponding voxels under the current coordinate transformation therefore indexes one of the matrix fields. The 2-D histogram defined by two images and a particular transformation is the matrix for which every entry has the value that equals the number of corresponding voxel pairs indexing this entry.

In 3-D to 3-D image registration, the voxel intensities of one of the two images (the "floating" or "interpolation image") need to be determined at the voxel locations of the other image "reference image"). Different methods can be used to enter the resulting voxel pairs into the 2-D histogram. The most straightforward techniques involve computing an interpolated intensity value

from the intensities of the eight voxels enclosing the respective location. Let for example r be the intensity of a particular reference image voxel and f_i for $i = 0, \ldots, 7$ the intensities of the eight enclosing voxels in the floating image. Then one may increment the histogram bin indexed by r and the interpolated floating voxel intensity f as follows, producing an updated histogram $\mathbf{H}'$ as follows:

$$\mathbf{H}'_{r,f} = \mathbf{H}_{r,f} + 1 \text{ where } f = \sum_i w_i f_i. \tag{3}$$

The two most commonly used interpolation schemes, nearest neighbor and tri-linear interpolation, are both special cases of the expression, each with a specific way of computing the interpolation coefficients w_i. However, Maes et al. [9] suggest a technique called partial volume integration that completely avoids intensity interpolation. Instead of applying an interpolation scheme such as the one outlined above to the voxel intensities, each of them is entered into the histogram with a weight that is determined by the tri-linear interpolation coefficients that would be applied in the particular situation. As the histogram is actually 2-D, this means that each of the values is actually paired with the single value taken from the other image, and all pairs are entered into the histogram with the respective weights.

$$\mathbf{H}'_{r,f_i} = \mathbf{H}_{r,f_i} + w_i \text{ for all } i. \tag{4}$$

This behavior can be understood as adding to the matrix $\mathbf{H}$ the result of the outer product of two vectors as follows. One of the vectors is the unit column vector $\mathbf{d}_r^T$ indexing the r-th row of $\mathbf{H}$ while the second vector is the distribution of weights assigned to the columns of $\mathbf{H}$:

$$\mathbf{H}' = \mathbf{H} + \mathbf{d}_r^T \left(\sum_{i=0}^{7} w_i \mathbf{d}_{f_i} \right) \tag{5}$$

Here and in all following equations we assume that the respective vector dimensions match the number of rows and columns of $\mathbf{H}$, respectively. The interpolation weights w_i are all between 0 and 1 with a total sum of 1. They can therefore easily be re-interpreted as probabilities in a distribution of discrete values (see Fig. 1). We refer to the similarity measures MI and NMI computed from the histograms thus generated as probabilistic MI (pMI), and probabilistic NMI (pNMI), respectively.

Background and Air vs. Bone Detection. Clinical images usually show the region of interest of the patient's body embedded in air. This is useful to ensure that the image boundaries do not crop the presentation of the patient, which would lead to incorrect computation of projections due to missing data. From an image processing point of view, the object of interest is thus surrounded by more or less extended regions of image background, easily detected by its low pixel intensities. For standard DRR computation, values close to zero have no substantial effect on the result.

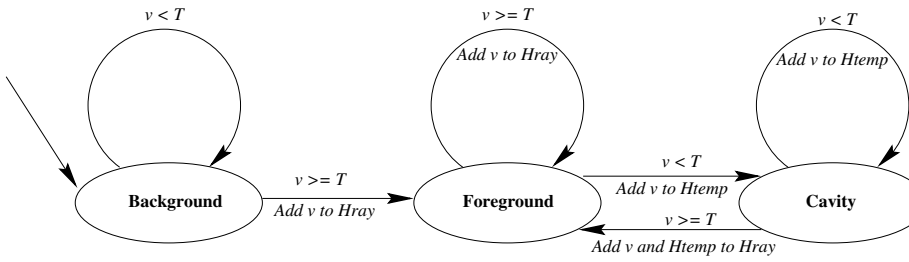

Fig. 2. Finite state machine to distinguish image background from air-filled cavities and surface folds. The lower object voxel threshold is denoted by T, the intensity of the next voxel along the ray is denoted by v. The inequalities over each arrow indicate the condition that leads to the respective state transition. The textual description under the arrows is the operation performed upon this transition.

However, when computing the distribution of intensities along a ray, the background pixels do have a substantial impact on the result. On the other hand, one cannot ignore all voxels identified as background by an intensity threshold, as this would also remove voxels that represent air-filled cavities inside the patient's body or surface folds. Both obviously carry important information about the shape and distribution of tissues inside the patient. When considering MR images, not abandoning voxels below a certain threshold becomes even more essential, since bony structures, from which most information in X-ray projection images originates, would also be removed by such an operation.

Instead of simple thresholding, we have implemented a finite state machine (FSM) to distinguish between air-filled cavities and bone inside the patient, voxels from which are included in the resulting tissue distribution, and image background, voxels from which are discarded. The FSM is illustrated in Fig. 2. Its fundamental principle of operation is to enter voxels encountered along the ray into either the ray histogram ("Hray") or a temporary histogram ("Htemp"), depending on which state the FSM is in. The temporary histogram temporarily stores below-threshold voxels which are moved to the main histogram when the next above-threshold voxel is encountered.

3 Results

We have computed the pNMI image similarity measure between probabilistic DRRs computed from a 3-D MR image and a DRR computed from a co-registered CT image[1]. The results are visualized in Fig. 4. For translations of up to 40 mm in either direction along the x, y, and z axes, we found a peak of the similarity measure at the known correct pose (translation in x and z direction), or at least close to it (within 4 mm in y direction).

[1] The registration transformation between CT and MRI was computed using an intensity-based algorithm based on NMI [12]. Our algorithm has been validated to achieve better than 1 mm accuracy for CT to MR registration using the Vanderbilt image data [10].

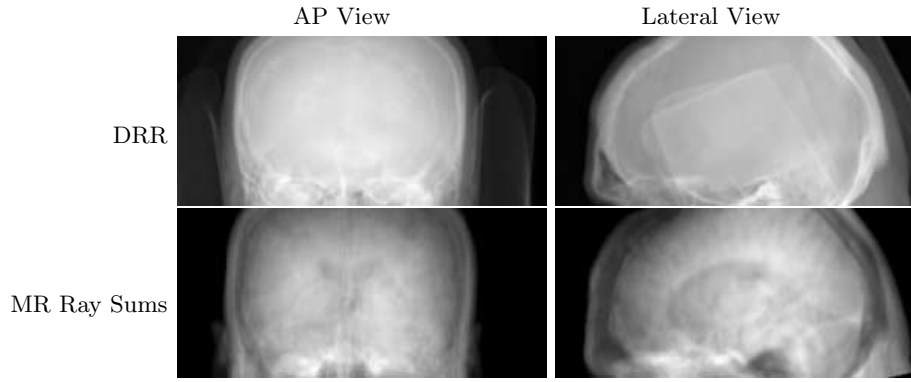

Fig. 3. DRR images (*top row*) and spatially equivalent MR ray sum images (*bottom row*). The 3-D CT and MR image were aligned using an intensity-based rigid-body image registration algorithm.

4 Discussion

This paper has presented a novel approach to the registration of 3-D tomographic images with 2-D projections. Our method is based on probabilities rather than intensities or geometric features. We have described an extension and reinterpretation of histogram-based similarity measures that allows us to compute these between probabilistic, non-scalar images. We have also introduced a probabilistic extension to DRR computation that is not based on the physical interpretation of voxel intensities as X-ray attenuation. Therefore, this extension and the subsequent computation of entropy-based similarity measures can be applied to other imaging modalities than CT. In particular, we have demonstrated the capability of our similarity measure to identify the correct pose of an MRI volume with respect to two orthogonal DRR images.

It is worth noting that the described procedure of computing pMI (pNMI) from pDRR is fundamentally equivalent to back-projecting the real projection image into 3-D space and computing standard MI (NMI) between the 3-D image and this back-projection. This observation may provide some justification for our method and explain to some extent how and why it works. In comparison, however, our method avoids problems resulting from the non-orthogonal grid of the back-projected data when working with the common projection geometries. Furthermore, our approach allows for an easy detection of background vs. bone and air-filled cavities along each ray, and the integration of fuzzy-segmented X-ray projections [13] is straight forward.

Obviously, the problem of registering MRI to real, especially intraoperative, X-ray projections is substantially harder than registering to DRR due to noise, presence of surgical instruments, and possibly geometrical distortions. We are therefore currently acquiring multi-modal 3-D image data (CT and MRI) and 2-D flat-panel X-ray images of patient anatomy with implanted markers that will provide for gold-standard pose information to validate our similarity measure against.

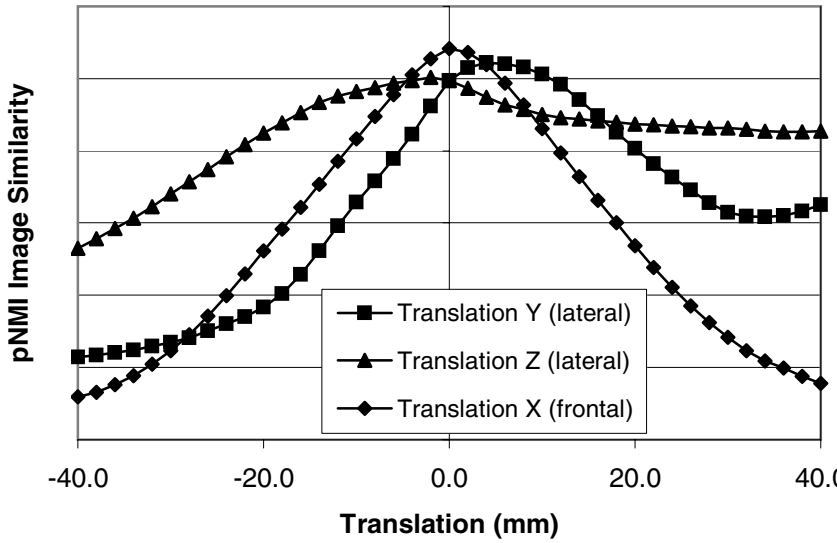

Fig. 4. Probabilistic normalized mutual information (pNMI) image similarity measure. Probabilistic DRRs were computed from MRI for different poses and compared to a single DRR image computed from a co-registered CT image. The similarity measure was plotted for translations. For translations along the x axis, image similarity was computed from the AP (frontal) projection image, since due to the near-parallel projection geometry there was no sufficient perspective scaling of the lateral projection images.

Acknowledgments

TR was supported by the National Science Foundation under Grant No. EIA-0104114. We acknowledge support for this research provided by CBYON, Inc., Mountain View, CA.

References

1. G. T. Herman, *Image Reconstruction from Projections*, Academic Press, 1980.
2. G. P. Penney, P. G. Batchelor, D. L. G. Hill, D. J. Hawkes, and J. Weese, "Validation of a two- to three-dimensional registration algorithm for aligning preoperative CT images and intraoperative fluoroscopy images," *Med Phys* **28**, pp. 1024–1032, June 2001.
3. G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. G. Hill, and D. J. Hawkes, "A comparison of similarity measures for use in 2D-3D medical image registration," *IEEE Trans Med Imaging* **17**, pp. 586–595, Aug. 1998.
4. D. Tomaževič, B. Likar, and F. Pernuš, "Rigid 2D/3D registration of intraoperative digital X-ray images and preoperative CT and MR images," in *Medical Imaging: Image Processing*, M. Sonka and J. M. Fitzpatrick, eds., vol. 4684 of *Proceedings of SPIE*, Feb. 2002. In print.

5. M. J. Murphy, J. R. Adler, M. Bodduluri, J. Dooley, K. Forster, J. Hai, Q.-T. Le, G. Luxton, D. Martin, and J. Poen, "Image-guided radiosurgery for the spine and pancreas," *Comput Aided Surg* **5**, pp. 278–288, 2000.
6. J. R. Adler, M. J. Murphy, S. D. Chang, and S. L. Hancock, "Image-guided robotic radiosurgery," *Neurosurgery* **44**, pp. 1299–1307, June 1999.
7. T. Rohlfing, D. B. Russakoff, M. J. Murphy, and C. R. Maurer, Jr., "An intensity-based registration algorithm for probabilistic images and its application for 2-D to 3-D image registration," in *Medical Imaging: Image Processing*, M. Sonka and J. M. Fitzpatrick, eds., vol. 4684 of *Proceedings of SPIE*, Feb. 2002. In print.
8. Y.-D. Liang and B. Barsky, "A new concept and method for line clipping," *ACM Transactions on Graphics* **3**, pp. 1–22, Jan. 1984.
9. F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximisation of mutual information," *IEEE Trans Med Imaging* **16**(2), pp. 187–198, 1997.
10. J. B. West, J. M. Fitzpatrick, M. Y. Wang, B. M. Dawant, C. R. Maurer, Jr., R. M. Kessler, R. J. Maciunas, C. Barillot, D. Lemoine, A. Collignon, F. Maes, P. Suetens, D. Vandermeulen, P. A. van den Elsen, S. Napel, T. S. Sumanaweera, B. Harkness, P. F. Hemler, D. L. G. Hill, D. J. Hawkes, C. Studholme, J. B. A. Maintz, M. A. Viergever, G. Malandain, X. Pennec, M. E. Noz, G. Q. Maguire, Jr., M. Pollack, C. A. Pelizzari, R. A. Robb, D. Hanson, and R. P. Woods, "Comparison and evaluation of retrospective intermodality brain image registration techniques," *J Comput Assist Tomogr* **21**(4), pp. 554–566, 1997.
11. C. Studholme, D. L. G. Hill, and D. J. Hawkes, "An overlap invariant entropy measure of 3D medical image alignment," *Pattern Recognit* **33**(1), pp. 71–86, 1999.
12. T. Rohlfing, J. B. West, J. Beier, T. Liebig, C. A. Taschner, and U.-W. Thomale, "Registration of functional and anatomical MRI: Accuracy assessment and application in navigated neurosurgery," *Comput Aided Surg* **5**(6), pp. 414–425, 2000.
13. D. B. Russakoff, T. Rohlfing, and C. R. Maurer, Jr., "Fuzzy segmentation of fluoroscopy images," in *Medical Imaging: Image Processing*, M. Sonka and J. M. Fitzpatrick, eds., vol. 4684 of *Proceedings of SPIE*, Feb. 2002. In print.