

An intensity-based registration algorithm for probabilistic images and its application for 2-D to 3-D image registration

Torsten Rohlfing¹, Daniel B. Russakoff¹, Martin J. Murphy², and Calvin R. Maurer, Jr.¹

¹Image Guidance Laboratories, Department of Neurosurgery, Stanford University, Stanford, CA

²Department of Radiation Oncology, Stanford University, Stanford, CA

ABSTRACT

Registration of 2-D projection images and 3-D volume images is still a largely unsolved problem. In order to register a pre-operative CT image to an intra-operative 2-D x-ray image, one typically computes simulated x-ray images from the attenuation coefficients in the CT image (Digital Reconstructed Radiograph, DRR). The simulated images are then compared to the actual image using intensity-based similarity measures to quantify the correctness of the current relative pose. However, the spatial information present in the CT is lost in the process of computing projections. This paper first introduces a probabilistic extension to the computation of DRRs that preserves much of the spatial separability of tissues along the simulated rays. In order to handle the resulting non-scalar data in intensity-based registration, we propose a way of computing entropy-based similarity measures such as mutual information (MI) from probabilistic images. We give an initial evaluation of the feasibility of our novel image similarity measure for 2-D to 3-D registration by registering a probabilistic DRR to a deterministic DRR computed from patient data used in frameless stereotactic radiosurgery.

Keywords: Probabilistic DRR, intensity-based image registration, mutual information, x-ray projection, CT, fuzzy segmentation.

1. INTRODUCTION

Registration of two-dimensional (2-D) projection images to three-dimensional (3-D) volume images is still a largely unsolved problem. However, there is a variety of clinical problems that can benefit from successful 2-D to 3-D registration. One example is fluoroscopic navigation where the relative pose of the patient during surgery and the pre-operative CT image is computed by registering the CT to one or several intra-operative fluoroscopic x-ray projection images with known projection geometry parameters and calibrated image position. Another example is radiation therapy, where the 2-D projection image is an electronic portal image (EPI)¹ generated by the therapy beam and typically detected using flat panel detectors.

In order to register a CT image to a 2-D projection image one typically computes a projection image from the attenuation coefficients in the CT image (Digital Reconstructed Radiograph, DRR).^{2,3} This simulated image is then compared to the actual image using intensity-based similarity measures to quantify the correctness of the current relative pose. One of the major problems that make registration of 2-D projections to 3-D images more difficult than 3-D-to-3-D image registration is that, due to the nature of projection imaging, spatial information is lost in the imaging process. The resulting projection pixels are the superposition of different tissues along the respective rays, and their exact spatial distribution cannot in general be recovered from the projection image.

However, using fuzzy segmentation techniques,⁴ pixels in projection images can be assigned probability distributions of tissues that the respective ray transverses. Hard segmentation techniques that assign a unique label to each pixel are obviously not useful here since there is no single tissue that corresponds to a projection pixel. While one cannot expect the assigned distributions to be correct for each pixel, there is hope that they model the overall projection process quite well.

Further author information: (Send correspondence to T.R.)

T.R., D.B.R., and C.R.M.: Image Guidance Laboratories, Department of Neurosurgery, Stanford University, 300 Pasteur Drive, MC 5327, Room S-008, Stanford, CA 94305-5327, phone: +1 (650) 498-7958, fax: +1 (650) 724-4846; T.R.: rohlfiing@stanford.edu; D.B.R.: russakof@robotics.stanford.edu; C.R.M.: calvin.maurer@igl.stanford.edu; M.J.M.: Department of Radiation Oncology, Stanford University, 300 Pasteur Drive, MC 5304, Stanford, CA 94305-5304, phone: +1 (650) 725-8380, fax: +1 (650) 498-4015; martin.murphy@stanford.edu;

Other than hard segmentation data, fuzzy segmented images do not have scalar pixel intensities but are vector-valued. In order to use this kind of image information in intensity-based image registration, existing algorithms can be modified to handle the probabilistic data. This paper describes an image similarity measure for probabilistic images, which is an extension of existing algorithms that optimize similarity measures computed from the joint probability distribution function. For a particular pose (position and orientation) of a CT image, we compute a probabilistic DRR (pDRR) by generating a histogram of CT intensities along each projection ray. Each pixel in the probabilistic DRR image therefore corresponds to a distribution of CT values along the ray that resulted in the projection value at that pixel. For the corresponding projection pixel, we compute a similar distribution using a fuzzy segmentation algorithm developed in related work of our group.⁴ The outer product of both distributions is computed, resulting in a 2-D joint histogram for each projection ray. The superposition of these histograms for all rays defines a combined joint histogram from which we compute mutual information by discrete entropy calculation.

We present in detail the methodology of our novel probabilistic similarity criterion and give a first brief evaluation of its accuracy and robustness properties. The similarity measure is computed for a range of translations and rotations of a 3-D CT image in the neighborhood of a known registration with a 2-D x-ray projection, determined using implanted markers. From the resulting plots, we estimate the expected capture range and registration accuracy using our criterion.

2. METHODS

2.1. Deterministic and probabilistic DRR

Three-dimensional CT and 2-D x-ray projection imaging are based on roughly the same physical processes.^{5,6} After traversing a homogeneous piece of matter with attenuation coefficient μ and thickness d , an x-ray with initial intensity J_0 is detected with a remaining intensity

$$J_d = J_0 e^{-\mu d}. \quad (1)$$

Conversely, the attenuation coefficient can be determined by measuring the initial and remaining x-ray intensity and computing

$$\mu = \frac{\ln J_0 - \ln J_d}{d}. \quad (2)$$

For a continuous path L in space through a heterogeneously attenuating object, the total attenuation μ_L along L is equal to the integral

$$\mu_L = \int_{-\infty}^{\infty} \mu(L(t)). \quad (3)$$

For paths of finite length, the integral in the above equation can be discretely approximated by a finite sum if the attenuation coefficient is only sampled at N discrete locations \mathbf{x}_i with a uniform spacing δ . In particular, consider a linear path (i.e., ray) originating from a ray source at location \mathbf{s} that reaches an x-ray detector at location \mathbf{x}_{det} . Then the location of the i -th sample along the ray is

$$\mathbf{x}_i = \mathbf{s} + i\mathbf{d} \quad (4)$$

with the normalized direction vector

$$\mathbf{d} = \delta \frac{\mathbf{x}_{\text{det}} - \mathbf{s}}{\|\mathbf{x}_{\text{det}} - \mathbf{s}\|}. \quad (5)$$

The total number of samples along the ray is defined by means of the distance between the location \mathbf{s} of ray source and the location \mathbf{x}_{det} of the detector for this particular ray:

$$N = \frac{\|\mathbf{x}_{\text{det}} - \mathbf{s}\|}{\delta}. \quad (6)$$

Deterministic DRR. The effective attenuation corresponding to the remaining intensity detected at \mathbf{x}_{det} is the DRR value at this location and can, analog to the integral in Eq. (3), be written as discrete sum

$$\text{DRR}(\mathbf{x}_{\text{det}}) = \delta \sum_{i=0}^{N-1} \mu(\mathbf{x}_i) \quad (7)$$

Probabilistic DRR. For the ray associated with the detector position \mathbf{x}_{det} we define the *probabilistic DRR* (pDRR) the distribution P of intensities μ sampled discretely along this ray:

$$\text{pDRR}(\mathbf{x}_{\text{det}}, c) = P[\mu(\mathbf{x}_i) = c | 0 \leq i < N]. \quad (8)$$

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.⁷ The probabilistic DRR can thus be equivalently rewritten as

$$\text{pDRR}(\mathbf{x}_{\text{det}}, c) = P[\mu(\mathbf{x}_i) = c | 0 \leq I_{\text{in}} \leq i \leq I_{\text{out}} < N]. \quad (9)$$

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 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. 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 (see Section 2.3).

2.2. Background detection

Clinical CT 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 DRR computation 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. 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.

Instead of simple thresholding, a finite state machine (FSM)⁸ is implemented to distinguish between air-filled cavities inside the object and surface folds, voxels from which are included in the resulting tissue distribution, and image background, voxels from which are discarded (see Fig. 2). The FSM is initialized in the “**Background**” state with two histograms, “Hray” and “Htemp” initialized as empty distributions. While “Hray” is used to construct the histogram of inside-patient voxel values along the ray, “Htemp” is a temporary histogram that takes up potential background voxels (voxels with values below the threshold T) as long as it has not been decided whether they are inside a cavity or outside the object.

Upon encountering the first voxel with intensity greater than or equal to T , the FSM switches to the “**Foreground**” state, in which it remains as long as all voxel intensities remain above T . When a voxel below the T is encountered in “**Foreground**” state, the FSM switches to the “**Cavity**” state, indicating that this voxel and all subsequent voxels with intensities below T may in fact be inside the patient. This is confirmed if and when another voxel intensity above T is encountered. In this case, the voxels transversed inside the cavity and stored in “Htemp” are collectively added to “Hray”, “Htemp” is reset, and the FSM goes back into the “**Foreground**” state.

If the end point of the ray is encountered while in the “**Cavity**” state (which should usually be the case unless the field-of-view is cropped), then all voxels in “Htemp” have been determined to be outside the object and are discarded.

2.3. Probabilistic mutual information

The mutual information (MI) image similarity measure^{9,10} has been used with great success in the registration of 3-D to 3-D images¹¹ (single or multi modality). Based on our previous experience, we usually apply the normalized mutual information¹² (NMI) image similarity measure that is derived from MI. Both measures are usually computed




Figure 1. Foreground, background and air-filled cavities along a ray. Voxels along the dotted parts of the ray are image background and need to be removed from the ray histogram. Voxels in air-filled cavities inside the patient are useful information and should be entered into the ray histogram (*right*).

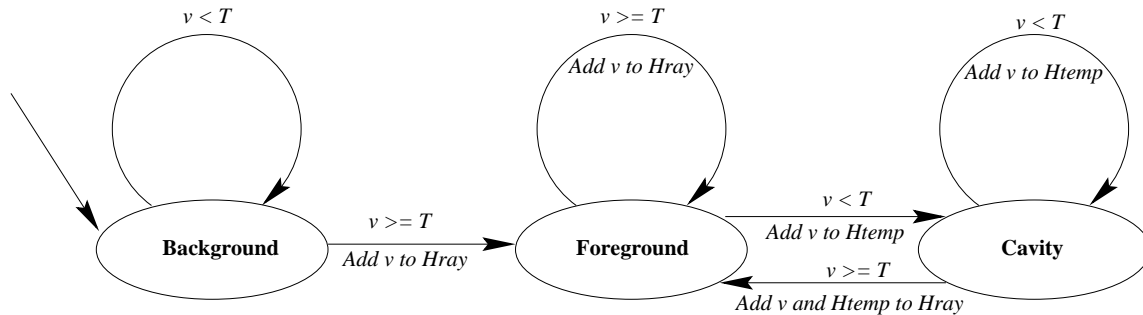


Figure 2. Finite state machine to distinguish image background from air-filled cavities and surface folds. The lower object voxel threshold is denoted by T , the intensity of the next voxel along the ray is denoted by v . The inequalities over each arrow indicate the condition that leads to the respective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

from discrete 2-D histograms. A 2-D histogram is a matrix \mathbf{B} 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 (see Fig. 3). 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 $i = 0, \dots, 7$ 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{B}' as follows:

$$\mathbf{B}'_{r,f} = \mathbf{B}_{r,f} + 1 \text{ where } f = \sum_{i=0}^7 w_i f_i. \quad (10)$$




Figure 3. Registration of CT and projection image using standard DRR and MI. For each projection pixel, the DRR value along the corresponding ray is computed. The bin in the 2-D histogram indexed by this pair of values is incremented by 1.




Figure 4. Registration of CT and projection image using probabilistic DRR and MI. For each projection pixel, the distribution of CT 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.

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 .

Maes et al.¹⁰ suggested a technique called partial volume integration (PVI) 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




Figure 5. Registration of CT and soft-segmented projection image using probabilistic DRR and MI. For each projection pixel, the distribution of CT 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.




Figure 6. Scaling of transformation parameters. A small increment δ is added separately to each parameter (here: for one of the rotational parameters). For each of the eight corners of the 3-D image, the motion vector $\Delta \mathbf{p}$ of its projection in the projection image plane is computed. The step size for the given parameter in the current optimization step is then determined such that the average magnitude of the motion vectors over all eight projected corners (i.e., the effective global motion within the projection plane) is equal to a given value.

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.

$$\mathbf{B}'_{r,f_i} = \mathbf{B}_{r,f_i} + w_i \text{ for } i = 0, \dots, 7. \quad (11)$$

This behavior can be understood as adding to the matrix \mathbf{B} 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{B} while the second vector is the distribution of weights assigned to the columns of \mathbf{B} :

$$\mathbf{B}' = \mathbf{B} + \mathbf{d}_r^T \left(\sum_{i=0}^7 w_i \mathbf{d}_{f_i} \right), \quad (12)$$

Here and in all following equations we assume that the respective vector dimensions match the number of rows and columns of \mathbf{B} , 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. 4).

If both images are probabilistic, i.e., both have probability distributions assigned to their voxels rather than scalar intensities (see Fig. 5), this principle can be generalized as follows:

$$\mathbf{B}' = \mathbf{B} + \left(\sum_{i=0}^7 v_i \mathbf{d}_{r_i} \right)^T \left(\sum_{i=0}^7 w_i \mathbf{d}_{f_i} \right). \quad (13)$$

We refer to the similarity measures MI and NMI computed from the histograms thus generated as probabilistic MI (pMI), and probabilistic NMI (pNMI), respectively. Other entropy-based similarity measures based on 2-D histograms can, of course, be evaluated as well. We have, however, not considered any of those yet.

2.4. Probabilistic DRRs and soft-segmented projection images

The idea of soft segmentation is to assign to each pixel or voxel not a single label but a vector of label probabilities.

Pixels in projection images are associated with a path in space (i.e., a projection ray) rather than a unique location in 3-D space. Therefore, these pixels cannot be assigned a unique tissue label. Instead, one can interpret the vector of label probabilities used in soft segmentation not as a vector of probabilities but as the a-posteriori distribution of tissues along the associated ray. In a related paper⁴ our group describes techniques to reconstruct these distributions from projection images with a known corresponding CT image but unknown mutual pose.

Obviously, the use of the resulting non-scalar images for the purpose of registration adds another level of non-scalar behavior to the method for MI computation described above.

2.5. Scaling of transformation parameters

Due to the projection geometry, not all parameters of the 3-D image transformation have equal effects on the computed projection image. For example, translation of the 3-D image parallel to the projection plane leads to substantially more motion in the projection image than translation orthogonal to the projection plane. Furthermore, the effect of transformation parameter changes also depends on the current set of parameters, i.e., the current pose of the 3-D image relative to the projection geometry.

For each 3-D transformation parameter a scaling factor (parameter weight) is computed such that when adding accordingly weighted increments to the all parameters, the average motion of all projected voxels in the projection plane is always equal to 1. A two-stage model is used to determine these parameter weights: A 3-D weight W_{vol} describes the average 3-D motion induced by changing the given parameter, while a projection component W_P describes its effect after projection onto the projection plane. The total weight for the n -th transformation parameter is the product of both components, i.e.,

$$W^{(n)} = W_{\text{vol}}^{(n)} W_P^{(n)}. \quad (14)$$

3-D component of parameter weights. For translational parameters, the 3-D weight is equal to 1, as each translation by δ causes a global voxel motion equal to δ . Therefore, the 3-D weights for all translational parameters are unity:

$$W_{\text{vol}}^{\text{Trans}} = 1. \quad (15)$$

For rotational parameters, we use the following approximation: All rotations are around an axis parallel to either the x , y , or z coordinate axis. Also, all rotation axes run through the center of the volume. For rotation around the axis parallel to z , the corners of the 3-D image have a distance of $\frac{1}{2}\sqrt{V_x^2 + V_y^2}$ from the rotation axis, where V_x and V_y are the world-coordinate sizes of the image in x and y direction, respectively. When incrementing the rotation angle (in radians) around this axis by δ , the corners will move on a circular trajectory with arc length equal to their distance from the axis multiplied by δ . As coordinates closer to the axis move less (points on the axis do not move at all), we approximate the average motion throughout the whole volume as half this value. In total, the parameter weight for rotation around a parallel to the z axis is

$$W_{\text{vol}}^{\text{RotZ}} = \frac{1}{\sqrt{V_x^2 + V_y^2}}. \quad (16)$$

A similar approximation can be performed for scaling and shearing parameters of a full affine coordinate transformation. For the context of this paper, however, these are irrelevant.

Projection component of parameter weights. In order to determine the projection component W_P of a transformation parameter weight, we analyze the motion of the projected corners of the 3-D image (see Fig. 6). For each of the eight image corners, denoted $\mathbf{v}_{0,0,0}$ through $\mathbf{v}_{1,1,1}$, the motion vector $\Delta\mathbf{p}_{i,j,k}$ is computed when adding δ to the n -th transformation parameter (yielding a 3-D transformation $\mathbf{T}_{\delta,n}$ instead of \mathbf{T}) and projecting the corner $\mathbf{v}_{i,j,k}$ according to the projection Π :

$$\Delta\mathbf{p}_{i,j,k} = \Pi(\mathbf{T}_{\delta,n}(\mathbf{v}_{i,j,k})) - \Pi(\mathbf{T}(\mathbf{v}_{i,j,k})) \quad (17)$$

The magnitude of the resulting motion vectors in the projection plane is averaged over all eight corners

$$\bar{v} = \frac{1}{8} \sum_{i,j,k=0}^1 \|\Delta\mathbf{p}_{i,j,k}\|_2. \quad (18)$$

The projection component of the n -th transformation parameter's weight is then the inverse of the average projected corner motion, normalized by the parameter increment δ :

$$W_P^{(n)} = \frac{\delta}{\bar{v}}. \quad (19)$$




Figure 7. X-ray projection images acquired using amorphous silicon flat panel detectors. The image from patient #1 (*left*) shows part of the lumbar spine; one of the vertebrae has previously been filled with (radio-opaque) bone cement. The image from patient #2 (*right*) shows part of the pelvic bone. Each image contains three implanted BB's, which were used to determine the gold standard transformation between the 2-D projections and the 3-D CT image. The original image resolution was 0.1275 mm pixel size, but the images were resampled to 0.51 mm pixel size for increased SNR and more efficient processing. The field of view is 120×120 mm (240×240 pixels). The distance of the ray source from the center of the projection plane was 329 cm.


3. RESULTS

We assessed our probabilistic image similarity criterion using image data used for planning and motion compensation in a system for robotic radiosurgery.^{13,14} The 2-D x-ray projection geometry was that of two fixed orthogonal amorphous silicon flat panel detectors,¹⁵ which unlike C-arm fluoroscopes are not subject to warping artifacts due to the earth's magnetic field. The projection images therefore do not require any image de-warping and thus do not contain BB's that would otherwise cause image occlusion. Details of the projection geometry are given in the caption of Fig. 7 and illustrated with two example x-ray projection images from two different patients. For each of these patients, a 3-D CT image was also acquired using the following imaging parameters: 512×512 pixels, 0.78 mm pixel size (patient #2: 0.84mm), 180 axial slices, 1.25 mm slice distance.

In order to investigate the convergence properties of the pMI image similarity measure, we have applied it to register probabilistic DRRs of the image data outlined above to deterministic DRRs of the same respective data set (results are plotted in Figs. 8 and 9). This may appear meaningless, as both, the probabilistic and the deterministic DRR, were computed from the very same data set. However, it needs to be noted that we do not compute an intensity-based image similarity between two DRR images. Instead, we compute the similarity of a scalar DRR image with a vector-valued image (the pDRR image). In this context, the registration of DRR and pDRR not only makes sense, but provides for the only mathematically accurate way of comparing projections computed with precisely the same pose and geometry. This is obviously never fully guaranteed with real image data.

The plots of the pMI image similarity measure in Figs. 8 and 9 show that pMI functions as a similarity measure. For all translation and rotation parameters, there is a peak of the similarity function near the correct transformation. For translational parameters, this peak extends beyond 20mm translations, indicating a wide capture range. The peak for translation parallel to the dominant direction of the projection rays, the peak is rather flat, which is due to the almost-parallel projection geometry and its almost negligible perspective scaling. For rotational parameters, the capture range is about 3 degrees on either side of the optimum parameter. The shape of the similarity function is particularly well-behaved with respect to rotation around the y axes (rotation within the coronal image plane). The

Background excluded



Background and soft tissue excluded




Figure 8. Image similarity plots for pMI measure of patient #1 (lumbar spine). *Top row:* Image background excluded (threshold -300 HU); *bottom row:* image background and soft tissue excluded (threshold +200 HU). The horizontal axis of the diagrams is normalized to effective motion in the projection plane (see Section 2.5). The true parameter ranges in 3-D were ± 17 mm for translation in x and z direction (parallel to projection plane), ± 210 mm in y direction (orthogonal to projection plane), ± 4.9 deg for rotations around the x and z axes, and ± 4.3 deg for rotations around the y axis.

removal of soft tissue in addition to image background (see Section 2.2) did not in general enhance the results we found. Only for rotation around the x axis in patient #1 (Fig. 8), the peak of the similarity measure moves closer to the theoretical optimum.

It has to be pointed out that when simultaneously considering a second, orthogonal projection, translation in y direction behaves like translation in x and vice versa. Also, rotation around the x axis behaves like rotation around the y axis. Therefore, the similarity function with respect to the respective transformation components makes up for the rather poor discrimination properties along those parameters in the first projection image.

4. DISCUSSION

This paper has extended the concept of intensity-based image registration algorithms using similarity measures computed from the joint probability distribution function to probabilistic image data. We have also introduced a probabilistic extension to the concept DRRs, which associates with a ray the distribution of tissue- or grey-value frequencies along that ray rather than a scalar ray sum. Our proposed probabilistic similarity measure shows

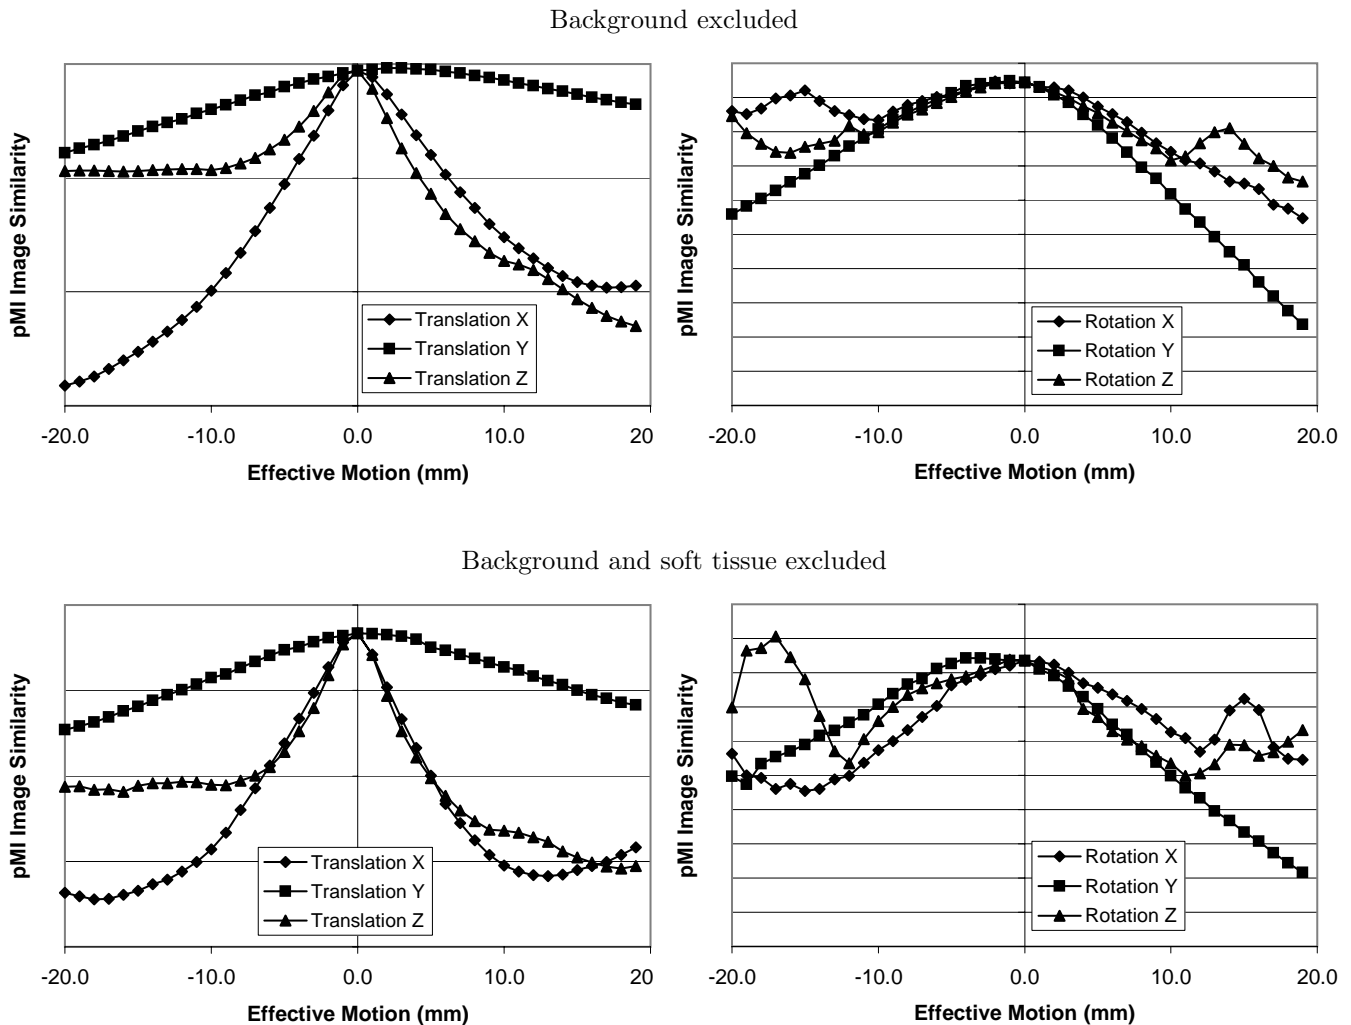


Figure 9. Image similarity plots for pMI measure of patient #2 (pelvis). *Top row:* Image background excluded (threshold -300 HU); *bottom row:* image background and soft tissue excluded (threshold +200 HU). The normalization of the horizontal axes of all diagrams is identical to the one in Fig. 8.

promising convergence properties when applied to registration of pDRR and deterministic DRR images. In addition to the application described in the scope of this paper, the probabilistic extension of histogram-based similarity measures can also be applied to other kinds of probabilistic data, such as probabilistic atlases.

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 projections⁴ is straight forward.

Future work on this project will be mostly focused on carefully comparing the novel probabilistic method for computing image similarity to existing image similarity measures. Also, as the fuzzy segmentation of projection images grows more mature, we will investigate the additional benefits of integrating segmented projection images into the registration process instead of the raw grey-value images.

ACKNOWLEDGMENTS

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

REFERENCES

1. A. L. Boyer, L. Antonuk, A. Fenster, M. Van Herk, H. Meertens, P. Munro, L. E. Reinstein, and J. Wong, "A review of electronic portal imaging devices (EPIDs)," *Med Phys* **19**, pp. 1–16, Jan. 1992.
2. 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.
3. 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.
4. 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.
5. G. T. Herman, *Image Reconstruction from Projections*, Academic Press, 1980.
6. Z.-H. Cho, J. P. Jones, and M. Singh, *Foundations of Medical Imaging*, John Wiley & Sons, New York, NY, 1993.
7. Y.-D. Liang and B. Barsky, "A new concept and method for line clipping," *ACM Transactions on Graphics* **3**, pp. 1–22, Jan. 1984.
8. J. E. Hopcroft and J. D. Ullman, *Introduction to Automata Theory, Languages and Computation*, Addison-Wesley, Reading, MA, 1979.
9. W. M. Wells, P. A. Viola, H. Atsumi, S. Nakajima, and R. Kikinis, "Multi-modal volume registration by maximization of mutual information," *Med Image Anal* **1**(1), pp. 35–51, 1996.
10. 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.
11. 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.
12. 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.
13. 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.
14. J. R. Adler, M. J. Murphy, S. D. Chang, and S. L. Hancock, "Image-guided robotic radiosurgery," *Neurosurgery* **44**, pp. 1299–1307, June 1999.
15. R. L. Weisfield, M. Hartney, R. A. Street, and A. R. B., "New amorphous-silicon image sensor for x-ray diagnostic medical imaging applications," in *Physics of Medical Imaging*, vol. 3336 of *Proceedings of SPIE*, pp. 444–452, SPIE, (Bellingham, WA), 1998.
16. A. Schweikard, G. Glosser, M. Bodduluri, M. Murphy, and J. R. Adler, "Robotic motion compensation for respiratory movement during radiosurgery," *Comput Aided Surg* **5**, pp. 263–277, Oct. 2000.