A Novel Approach for the Registration of 2D Portal and 3D CT Images for Treatment Setup Verification in Radiotherapy

Ravi Bansal[1], Lawrence H. Staib[1], Zhe Chen[2], Anand Rangarajan[1], Jonathan Knisely[2], Ravinder Nath[2], and James S. Duncan[1]

[1] Departments of Electrical Engineering and Diagnostic Radiology,
[2] Department of Therapeutic Radiology, Yale University, New Haven, CT 06520-8042

Abstract. In this paper we present a framework to simultaneously segment portal images and register them to 3D treatment planning CT data sets for the purpose of radiotherapy setup verification. Due to the low resolution and low contrast of the *portal image*, taken with a high energy treatment photon beam, registration to the 3D CT data is a difficult problem. However, if some structure can be segmented in the portal image, it can be used to help registration, and if there is an estimate of the registration parameters, it can help improve the segmention of the portal image. The minimax entropy algorithm proposed in this paper evaluates appropriate entropies in order to segment the portal image and to find the registration parameters iteratively. The proposed algorithm can be used, in general, for registering a high resolution image to a low resolution image. Finally, we show the proposed algorithm's relation to the mutual information [19] metric proposed in the literature for multi–modality image registration.

1 Introduction

Registration is the process that maps pixels from one image, called the *reference image*, to the pixels in another image, called the *test image*. For clinical diagnosis, treatment planning and delivery of a therapy, images from different modalities are often acquired as they provide complementary information about a disease and also can give graphical verification of a delivered therapy. [1]

In radiotherapy for cancer patients, registration of images from different modalities, such as computed tomography (CT) and magnetic resonance imaging (MRI), is often used in treatment planning to delineate the correct spatial extent of the tumor and the surrounding normal structures. Furthermore, registration of images from the same modality is also used for routine radiographical verification of the radiation treatment. In this case, X–ray images are taken at

---

[1] Good reviews of medical image registration, with classification, for multi-modal images can be found in [18], whereas [12] is an excellent review of registration methods in computer integrated surgery.

different energies, one at diagnostic energy (40 – 100 KV) (called the simulation image) when the treatment setup is initially simulated and others using a high energy (4 – 20 MV) treatment beam (called portal images) throughout the treatment course. Due to the high photon energy, portal images are intrinsically of low contrast and poor sharpness. It has remained a difficult task to automate the registration of the simulation and portal images. In fact, such registration is still performed by visual comparison in most radiotherapy centers. Many treatment centers are moving toward offering full 3D conformal treatments that are initially planned from 3D treatment planning CT data sets. Registration of the 3D image set to 2D portal images is therefore necessary to fully quantify the three–dimensional patient setup before treatment [7, 13, 3, 6].

Related Work. To register the portal image to a 3D pre–treatment image, the methods proposed in the literature are either semi–automated or require feature extraction from the portal image. The poor quality of the portal image makes this a difficult problem. However, many registration algorithms require feature extraction as a preprocessing step. Thus, the accuracy of these registration methods is limited by the accuracy of the method to segment features from the portal image. These methods are also called *sparse field* methods.

Instead of using features, some registration methods work directly with the gray scale pixel values of the images [21, 19, 20]. The authors use cross–correlation of pixel intensities has been used as the match metric for aligning portal and simulation images [14]. Woods et al. [21], measure the mis-registration between two images as the dispersion, or variance, of intensities in the 2D histogram. However, this metric is largely heuristic. An information theoretic match metric has been proposed [19] which can be used for multi–modality image registration. These methods which directly manipulate image intensities are called *dense field* methods. The dense field methods, as compared to the sparse field methods, are more robust to noise as these methods do not require extraction of features to be registered.

Our Approach. If the 2D portal and 3D CT images are properly aligned, then the information from the high resolution 3D CT can be used to segment the portal image. On the other hand, if we have an accurate segmentation of the portal image, an accurate registration can be obtained. Thus, this becomes a problem of deciding whether to try to locate structure in the portal film for registration purposes or to perform a crude registration first that might help to look for structure which could be used for better registration. In this paper we propose that there is no clear answer, and the registration and segmentation should be carried out simultaneously, each helping the other.

Thus, we propose an iterative framework in which the segmentation and the registration of two images are estimated in a two stage algorithm using two different entropies, termed *minimax entropy*. In the entropy maximization step, the segmentation of the portal image is estimated, using the current estimates of the registration parameters. In the entropy minimization step, the registration parameters are estimated, based on the current estimates of the segmentation.

The algorithm can start at any step, with some appropriate initialization on the other. We derive this strategy starting from the EM algorithm and show the relationship to the mutual information metric.

2 Problem Definition and Motivation

Patient setup verification, before performing radiotherapy, using a portal and a 3D CT image, is a classical pose estimation problem in computer vision. In the pose estimation problem, the pose of a 3D model is to be determined from a set of 2D images such that the model is aligned with the images. We assume that a rigid transformation of the 3D CT data set will bring it in to alignment with the portal image and hence there are only six transformation parameters to be considered. Since our algorithm works directly on the pixel intensity values, we classify our algorithm as an automated, intrinsic, global, dense–field based algorithm based on the classification in [18].

In the next subsection, we formulate the pose estimation as a maximum likelihood estimation (MLE) problem. This formulation requires the knowledge of the joint density function between the portal and the DRR pixel intensities, denoted as $p(x_i, y_i(T))$ below. As noted earlier in the paper, we want to estimate a segmentation of the portal image in order to help in the registration process. The problem of segmenting the portal image is formulated as a labeling problem in which each pixel is labeled either as bone or background. To incorporate the segmentation information, the joint density function is written in a *mixture density* form where the labels on the portal image pixels are not known. Hence, the segmentation labels are treated as the missing information, and the parameters to be determined are the pose parameters. We formulate our problem as a maximum–likelihood estimation (MLE) problem in section 2.1. An algorithm for computing the MLE from incomplete data is presented in [5] and is called the expectation–maximization (EM) algorithm. Our first thoughts were to select EM algorithm to compute the estimates, are due to its proven monotonic convergence properties, ease of programming and unlike other optimization techniques, it is not necessary to compute Hessians nor it is necessary to worry about setting the step–size. Thus, the EM algorithm may be used to estimate both the pose parameters and the segmentation labels. However, there are several problems with the EM approach that restrict its use in this problem. These are explained in subsection 2.2 and a minimax entropy strategy that addresses these restrictions is described in section 3.

2.1 Problem Definition

Let, $X = \{x(i)\}$, for $i = 1, \ldots, N^2$ be the $N \times N$ portal image. Similarly, let $Y\ (T) = \{y(i, T)\}$ for $i = 1, \ldots, N^2$ be the $N \times N$ projected 3D CT image, at the given transformation parameters T. Let, $G = \{g(i)\}$, for $i = 1, \ldots, N^3$ be the 3D CT image. Note here we index pixels in the images using a single index, even though the images are 2D (or 3D) images. For simplicity in the formulations

below, we have assumed that the portal and the projected 3D CT are of the same dimensions. However, this need not be true. The projected 3D CT image is also called the *digitally reconstructed radiograph* (DRR). Each pixel in the portal image is classified to belonging to one of two classes, bone or background. Thus, we formulate the pose estimation problem as a maximum likelihood estimation problem:

$$\hat{T} = \arg\max_T \ \log p(X, G|T)$$

$$= \arg\max_T \log \left(\frac{p(X, G, T) \ p(G|T)}{p(G, T)} \right)$$

$$= \arg\max_T \ [\log p(X, Y(T)) - \log p(Y(T))]$$

$$= \arg\max_T \sum_i [\log p(x_i, y_i(T)) - \log p(y_i(T))] \qquad (1)$$

where we ignore the term $p(G|T)$, since the 3D CT data set, G, is independent of the transformation parameters, T. In equation (1) we assume that the image pixels are independent. The logarithm of the likelihood function is taken to simplify the mathematical formulation. Note that for notational simplicity, we shall now write $x(i) = x_i$ and $y(i, T) = y_i(T) = y_i$, where the current transformation parameters for which the DRR is obtained should be clear from the context, otherwise it will be made explicit.

Mixture Model and the EM. Letting $A = \{1, 2\}$, be the set of classes (i.e. bone, background), where non–bone pixels are defined to be background and using a mixture density model, we can now write the joint mixture density model, from equation (1), for the portal image and DRR as, at pixel i, $p(x_i, y_i) = \sum_{a \in A} p(m_i = a) \ p_a(x_i, y_i) = \sum_{a \in A} p_i(a) \ p_a(x_i, y_i)$, where m_i is a random variable for the ith pixel taking values in A and $p_a(x_i, y_i)$ is the joint density function of the pixel intensities, given that the portal image pixel is labeled a. Classifying the pixels in the portal image is equivalent to the determination of the probability mass function of m_i. Let Z be the N^2 x 2 *classification matrix* with z_{ai} its ith row and ath column entry satisfying the constraints $\sum_{a=1}^2 z_{ai} = 1$, for $i = 1, \ldots, N^2$. Note that $z_{ai} = 1$ if the ith pixel of the portal image belongs to the class a, i.e., $z_{ai} \in \{0, 1\}$. Thus, z_{ai} are the indicator variables. For each indicator variables, we shall define the expected value as $< z_{ai} > \equiv p_i(a)$.

Using this notation, the EM algorithm [5] for the mixture model can be written as, at the kth iteration,

E–Step:

$$Q(T, T^{(k-1)}) = \sum_i \sum_{a \in A} < z_{ai} >^k \ \log p_a(x_i, y_i) - \sum_i \log p(y_i) \qquad (2)$$

M–Step:

$$T^k = \arg \ \max_T \ Q(T, T^{(k-1)}) \qquad (3)$$

where,

$$< z_{ai} >^k = \left(\frac{< z_{ai} >^{k-1} p_a(x_i, y_i')}{\sum_{b \in \mathbf{A}} < z_{bi} >^{k-1} p_b(x_i, y_i')} \right) \tag{4}$$

where, $y_i' = y(i, T^{(k-1)})$.

2.2   Limitations of EM Approach and Alternative Strategies

For our purposes, the EM algorithm has two key restrictions. First, in the EM algorithm for the mixture model as formulated in equation (2), the function $Q(T, T^{(k-1)})$ is defined only if the component density functions $p_a(x_i, y_i)$, $\forall a$ are specified. The component density functions need to be determined for the specific problem. Thus, for the images used in prostate cancer radiotherapy, the joint density functions need to be determined from the projection model, taking into account the physics of interaction of high energy X–rays with matter. If the joint density functions are not known, then usually a simplified model has to be assumed. In this paper, we propose to estimate the unknown density functions from the given data instead of using simplified models. The second observation comes from Neal and Hinton [15] who provide a view of the general EM algorithm as a coordinate descent algorithm where the coordinates axes are the joint density functions on the missing information and the parameters to be estimated. (A similar view of the EM algorithm for parameter estimation of mixture density models has been given [8] as well). In the E–Step of the EM algorithm, the expected value is calculated, as the missing information is not known. However, this implies that the expected value in equation (2) should be taken with respect to the density function $p(Y)$, instead of $p(Y|X; T^{(k-1)})$. However, Neal and Hinton [15] show that at each iteration of the EM algorithm, the density function $p(Y)$ is estimated to be $p(Y|X; T^{(k-1)})$ and the question still remains as to the relationship between the two density functions. In our approach described in section 3, the *max* step in the *minimax algorithm* formalizes this relation. The *max* step can be viewed as a step for determining the density function on the missing information. The *principle of maximum entropy* [10] (or the *principle of maximal non–committance*) states that in making inferences on the basis of partial information we must use that probability distribution which has maximum entropy subject to whatever is known. This is the only unbiased assignment we can make; to use any other would amount to an arbitrary assumption of information which by hypothesis we do not have.

We overcome the restrictions of the EM algorithm by borrowing the idea of averaging over the estimated density function from *mutual information*. Mutual information as a match metric has been first proposed and successfully applied for multi–modality image registration by [19, 4]. However, the mutual information match metric formulation in the literature assumes that all the pixels in the image are i.i.d., an assumption not true in general. The proposed minimax entropy algorithm described below thus aims at combining the strengths of both the EM algorithm and the mutual information based registration approach to simultaneously segment and register images.

3   Minimax Entropy Approach

The proposed minimax algorithm for solving the basic problem posed by equation (1), in a computational form similar to the EM strategy described by the equations (2), (3) and (4) has two steps, the max step and the min step, which are evaluated iteratively to determine the registration parameters and the probability density function (pdf) of the portal image segmentation. The details of the development of the proposed algorithm are discussed elsewhere [1]. Again, using the notation from section 2, we see the steps mathematically as follows:

    <u>Max Step</u>:

$$p^k(m) = \arg \max_{p(m)} \left[- \int p(m) \log p(m) \; dm \right.$$
$$\left. + \int p(m) \; \log p(m|X,Y;T^{(k-1)}) \; dm \right] \qquad (5)$$

under the constraint $\int p(m) \; dm = 1$, where m is the random variable whose domain is the set of possible segmentations of the portal image, where each pixel can be labeled from the set of labels A. We assume that pixel labels are statistically independent, i.e., $p(m) = \prod_i p(m_i = a) = \prod_i p_i(a)$.

    As formulated above, the max–step simply states that the maximum entropy estimate of the density $p(m)$ is simply $\log p(m|X,Y;T^{(k-1)})$ [2]. This simple formulation of the estimated segmentation density function allows us to systematically put constraints on the density function, as we show below.

    <u>Min Step</u>:

$$T^k = \arg \min_T H(m, X|Y)$$
$$= \arg \min_T \left(\sum_{a \in A} < \overline{z_a} >^k H_a(x, y) - H(y) \right) \qquad (6)$$

where $< \overline{z_a} >^k = (\frac{1}{N^2}) \sum_{i=1}^{N^2} < z_{ai} >^k$. See [1] for the development of the min–step. The component density function for class a, $p_a(x, y)$, is estimated as the weighted sum of Gaussian kernels using the Parzen window method:

$$p_a(w) = \frac{1}{\sum_{w_i \in D} p_i(a)} \sum_{w_i \in D} p_i(a) \; G_\Psi(w - w_i)$$

where, $w = \begin{pmatrix} x \\ y \end{pmatrix}$ and $p_i(a) = p(m_i = a)$ is the probability that the ith pixel in the portal image belongs to class a, estimated in the max step, equation (5), Ψ is 2–by–2 covariance matrix, which is assumed to be diagonal and D is the set of pixels sampled from portal image and the DRR to estimate the conditional density functions. Thus, $H_a(x, y) = - \int \int p_a(x, y) \log p_a(x, y) \; dx \; dy$. In the min step, we have chosen the mutual information (MI) metric currently popular in

the medical image analysis community. We have found it to be more robust than other interesting metrics (e.g. correlation. See [1] where we present the relation between the two metric and show that the MI is more sensitive to mis–registration then correlation). Because MI assumes that pixels are i.i.d., in general this is a problem, which we get around by using a mixture densities. We note that in [17], the authors register images with mutual information as a match metric while incorporating segmentation information on one of the images. However, the image was pre–hand segmented and thus remains fixed throughout the registration. In our proposed algorithm, the portal image segmentation is estimated simultaneously with the transformation parameters.

An annealing schedule [11] is imposed on the estimated portal image pixel segmentation distribution. The modified *max step*, equation (5), can thus be written as:

Max Step:

$$p^k(m) = \arg \max_{p(m)} \left[-\frac{1}{\beta} \int p(m) \log p(m) \, dm \right.$$

$$\left. + \int p(m) \, \log p(m|X,Y;T^{(k-1)}) \, dm \right] \tag{7}$$

under the constraint $\int p(m) \, dm = 1$, where $\beta = \frac{1}{t}$, and t is the *temperature*, which determines the annealing schedule.

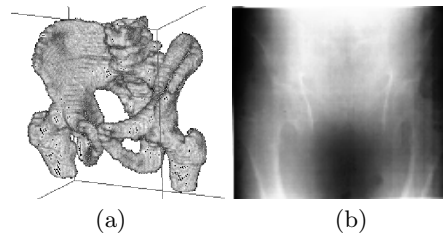## 4   Results

(a)   (b)

Fig. 1. (a) Volume rendered 3D CT phantom for which the simulated portal images were calculated. (b) Real portal image of the phantom obtained by taking the high energy X–ray of the phantom in the treatment room.

The experiments were carried out on simulated and real portal images. To obtain simulated portal images, the 3D CT data set was first rotated or translated in the 3D space, by a known amount, and then a perspective projection was computed. These simulated portal images, rendered under known parameters, are used to study the accuracy of the algorithm. Note that we assume that the perspective parameters are known.

The 3D CT data set, shown in figure 1 (a), is a high resolution diagnostic energy data set obtained by scanning a pelvic phantom. The phantom is made of real human pelvic bone enclosed in plexi–glass with density close to soft tissue. After the CT data set was obtained at diagnostic energy X–rays (40 KeV), its voxel values were mapped to high energy X–rays (6MeV), using attenuation coefficient tables [9], to generate projection images close to the quality of the portal image. Since the diagnostic energy is known, the attenuation coefficient tables are used to estimate the tissue type in each voxel. Once the tissue type is known, the attenuation coefficient for that tissue at high energy can be estimated from the tables. These attenuation coefficients are used to map the phantom voxel values to their high energy values. In figure 1 (a), the 3D CT data set was rendered with the opacity value for soft tissue set close to zero.

Figure 1 (b) shows the portal image obtained by imaging the phantom with X–rays at treatment energy level (6 MeV). The portal image was contrast enhanced by histogram equalization. Experimentally, we found that histogram equalized portal images lead to more accurate results. The images shown in figure 4 are simulated portal images obtained at known transformation parameters to validate the proposed algorithm. Scattering of high energy X–ray photons leads to noise and poor contrast in the portal image. To simulate photon scattering, varying amounts of independent and identically distributed (i.i.d.) Gaussian noise were added to the pixels in the simulated portal image.
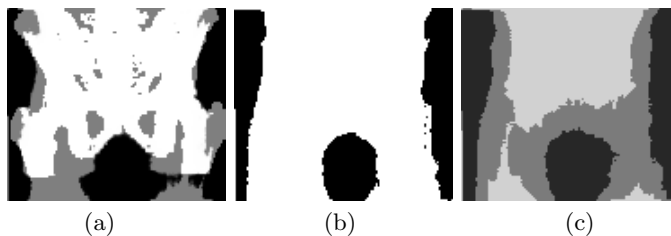

(a)   (b)   (c)

Fig. 2. Estimated segmentation results of the true portal image. (a) Results of using proposed minimax entropy algorithm. (b) Simple threshold. (c) Clustering algorithm in MEDx [16].

Figure 2 shows the segmentation of the real portal image, figure 1 (b), estimated by different segmentation algorithms. Figure 2 (a) shows the segmentation as estimated by the proposed algorithm. The pixel labels are estimated based on the joint density distribution. The bright pixels are labeled bone. The dark pixels are labeled background. Gray pixels are the pixels whose labels cannot be decided based on the joint probability density functions. Figure 2 (b) is the segmentation of portal image under simple, manually selected, threshold. Figure 2 (c) is segmentation using a clustering algorithm, from MEDx. In this clustering algorithm, the image is segmented into three classes based on the nearness of the gray values to the user specified values.

Figure 3 shows the result of registering the real portal image, figure 1 (b), to the 3D CT data set, figure 1 (a). The image in figure 3 (b) is the DRR rendered at the registration parameters estimated by the minimax entropy algorithm. To see the goodness of the estimated parameters, contours were hand drawn on the portal image by an expert, matching closely to key features. These contours were then mapped onto the DRR, using the estimated parameters determined by the proposed algorithm. The contours match closely to the feature points in the DRR, at least showing the accuracy of the estimated transformation parameters in mapping this expert–traced estimate of bone boundaries.

Estimated Parameters					
t_x (vox)	t_y (vox)	t_z (vox)	α (deg)	β (deg)	γ (deg)
2.56	-6.69	2.34	2.51	-1.24	0.73

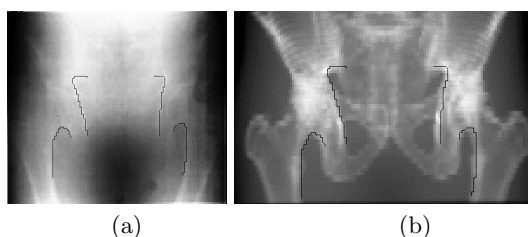(a)   (b)

Fig. 3. (a) Portal image with contours. The reference contours shown are drawn, on the portal image, by an expert, matching anatomical features of interest, to study the registration result. (b) The contours are then mapped on to the DRR rendered using the parameters estimated by the proposed minimax entropy algorithm.

Figure 4 (c) shows the initial results obtained for registration of simulated portal images. Each graph corresponds to the variation of one transformation parameter in the rendering of the simulated portal image. For the graph labeled *X–Trans*, for example, the 3D CT data set was translated along the X–axis by 5 voxels and then a projection was rendered with varying noise. To register the simulated portal image to the 3D CT data set, the transformation parameters were reset to zeros. The results can be summarized as follows. The minimax entropy algorithm is quite robust against noise for in–plane translations along the x–axis and y–axis. The performance of the method deteriorated gracefully for rotations, both in–plane and out–of–plane, as the noise was increased. In the initial steps of the proposed algorithm, when most of the labels have equal probability of belonging to either class, the algorithm is increasing the mutual information between the entire images. Later in the iterations, as the portal image pixels get classified as belonging to one of the two classes, the algorithm is increasing the mutual information of each of the two classes separately.

Overall, for the simulated portal images, where the true registration parameters were known, the estimated parameters converged to the true parameters

after only a few iterations. The later iterations of the algorithm led to better segmentation of the (simulated) portal image, which in turn helped registration to lock into a minimum. The proposed algorithm, as implemented, took about 5 to 10 minutes on SGI Indigo 2 R10K machine.
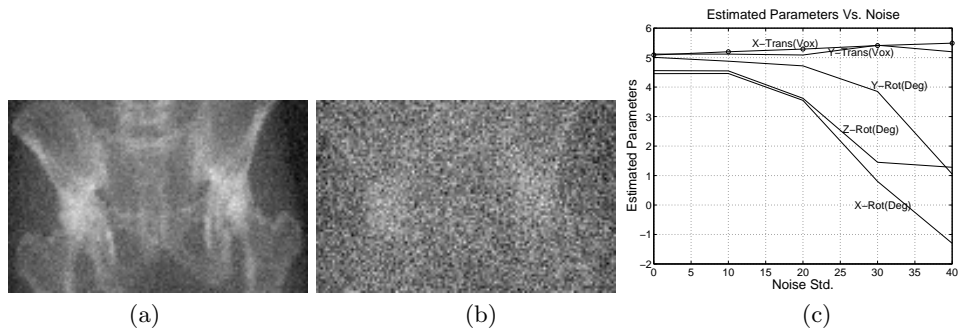

(a)   (b)   (c)

Fig. 4. Digitally rendered simulated portal images from the 3D CT data set in figure 1. The 3D CT data set was rotated by 5^o about the z–axis before taking the perspective projection. Varying amounts of i.i.d. Gaussian noise were added. The maximum pixel value of the simulated portal image is 296. (a) 2D rendered portal image, noise, std (σ) = 10. (b) Std (σ) = 30. (c) This figure shows the graph of estimated parameters against increasing noise added to the simulated portal image.

5   Discussion and the Future work

In this paper we presented an information theoretic framework in which segmentation and registration are carried out together iteratively, with segmentation results helping in the registration and vice–versa. Most registration methods proposed in the literature carry out portal image segmentation as a pre–processing step in the registration process, if at all. Our approach of simultaneously segmenting and registering the images, using a unified framework, leads to a novel and robust algorithm. We need to demonstrate the accuracy of the algorithm in comparison to mutual information based algorithm using raw intensity information only.

The mutual information match metric overcomes the assumption of a linear relationship between the pixel intensities of the images to be registered, an underlying assumption in the correlation match metric. In mutual information based registration, the relationship between the pixel intensities is estimated from the given data itself and thus can register images from different modalities. At an estimated set of transformation parameters, a joint density between the images to be registered can be estimated, from the given data. The mutual information metric assigns a number to each such estimated density, estimated

for different transformation parameters The transformation parameters corresponding to the density having the largest mutual information are chosen as the estimated parameters by the algorithm. The EM algorithm provides an iterative framework to estimate the parameters of a distribution, in the presence of missing data. However, the EM algorithm requires that the parametric form of the distribution to be known.

The proposed minimax entropy algorithm overcomes this restriction of the EM algorithm, by borrowing the idea from the mutual information method of estimating the joint distribution from the given data, and emphasizes the fact that the distribution on the segmentation labels is the maximum entropy distribution, satisfying the given constraints. This fact allows us to invoke other constraints on the distribution systematically, which we used to impose an annealing schedule. In the proposed algorithm, if the joint densities are not known, then instead of assuming some simplified parametric form, we estimate the densities from the data given. Then we use the information theoretic framework to estimate the transformation parameters.

Also, we note the relationship of our algorithm to the mutual information match metric which recently has been shown to be applied successfully to multi–modality image registration.

Our future research includes algorithm speed–up, validation of the accuracy and robustness of the algorithm, especially in comparison to the mutual information based registration and the ridge based algorithm proposed in [6]. Our experiments demonstrated that the estimated registration was not as accurate for out of plane rotations and translations, especially as the noise was increased. Thus, we propose using two orthogonal views, anterior–posterior and lateral views, commonly taken in the treatment room, for more accurate registration results. We will then extend the algorithm to make use of portal images acquired from more that two views which need not be orthogonal.

References

1. R. Bansal, L. Staib, et al. A novel approach for the registration of 2D portal and 3D CT images for treatment setup verification in radiotherapy. Technical Report 1998-1, Dept. of Electrical Engineering and Diagnostic Radiology, Yale University, Jan 1998. (http://noodle.med.yale.edu/~bansal/techreps/regis.ps.gz).
2. R. M. Bevensee. *Maximum Entropy Solutions to Scientific Problems.* P T R Prentice Hall, 1993.
3. J. Bijhold. Three-dimensional verification of patient placement during radiotherapy using portal images. *Med. Phys.*, 20(4):347–, March/April 1993.
4. A. Collignon, F. Maes, et al. Automated multimodality image registration using information theory. *Info. Proc. in Med. Imaging (IPMI)*, pages 263–274, 1995.
5. A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via EM algorithm. *J. Royal Statistical Soc., Ser. B*, 39:1–38, 1977.
6. K. Gilhuijs. *Automated verification of radiation treatment geometry.* PhD thesis, Univ. of Amsterdam, Radiotherapy dept., the Netherlands, 1995.

7. K. G. A. Gilhuijs et al. Interactive three dimensional inspection of patient setup in radiation therapy using digital portal images and computed tomography data. *Int. J. Radiation Oncol. Biol. Phys.*, 34(4):873–885, 1996.

8. R. J. Hathaway. Another interpretation of the EM algorithm for mixture distributions. *Statistics & Probability Letters, North Holland*, 4:53–56, 1986.

9. J. H. Hubble. *Photon Cross Sections, Attenuation Coefficients, and Energy Absorption Coefficients From 10KeV to 100 GeV*. Nat. Stand. Ref. Data. Ser., Nat. Bur. Stand. (U.S.), August 1969.

10. E. T. Jaynes. Information theory and statistical mechanics I. In R. D. Rosenkrantz, editor, *E. T. Jaynes: Papers on probability, statistics and statistical physics*, volume 158, pages 4–16. D. Reidel Publishing Company, Boston. USA, 1983.

11. S. Kirkpatrick, C. Gelatt, and M. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.

12. S. Lavallee. Registration for computer–integrated surgery: Methodology, state of the art. In Russell H. Taylor et al., editors, *Computer–Integrated Surgery: Technology and Clinical Applications*, pages 77–97. The MIT Press, Cambridge, 1996.

13. S. Lavallee, R. Szeliski, and L. Brunie. Anatomy–based registration of three–dimensional medical images, range images, X–ray projections, and three–dimensional models using octree–splines. In Russell H. Taylor et al., editors, *Computer–Integrated Surgery: Technology and Clinical Applications*, pages 115–143. The MIT Press, Cambridge, Massachusetts, 1996.

14. J. Moseley and P. Munro. A semiautomatic method for registration of portal images. *Med. Phys.*, 21(4), April 1994.

15. R. M. Neal and G. E. Hinton. A view of the EM algorithm that justifies incremental, sparse, and other variants. In M. I. Jordan, editor, *Learning in Graphical Models*. Kluwer Academic Press. To Appear. Web address: http://www.cs.toronto.edu/∼radford/papers-online.html.

16. Sensor Systems, Inc. *MEDx^{TM}: Multimodality Radiological Image Processing For Unix Workstations*, 1998. Version 2.1.

17. C. Studholme et al. Incorporating connected region labelling into automated image registration using mutual information. In *Proc. of MMBIA'96*, pages 23–31, 1996.

18. P. A. van den Elsen, E. D. Pol, and M. A. Viergever. Medical image matching– a review with classification. *IEEE Eng. in Med. and Biol.*, 12(1):26–39, 1993.

19. P. Viola and W. M. Wells. Alignment by maximization of mutual information. *Fifth Int. Conf. on Computer Vision*, pages 16–23, 1995.

20. J. Weese et al. 2D/3D registration of pre–operative ct images and intra–operative x–ray projections for image guided surgery. In H.U.Lemke et al., editors, *Comp. Assist. Rad. and Surgery*, pages 833–838. 1997.

21. R. Woods, S. Cherry, and J. Mazziotta. Rapid automated algorithm for aligning and reslicing PET images. *J. Comp. Assist. Tomography*, 16(4):620–633, 1992.