A Multiple-Layer Flexible Mesh Template Matching Method for Non-rigid Registration between a Pelvis Model and CT Images

Jianhua Yao[1], Russell Taylor[2]
1. Diagnostic Radiology Department, Clinical Center, NIH
2. ERC-CISST, Johns Hopkins University

ABSTRACT

A robust non-rigid registration method has been developed to deform a pelvis model to match with anatomical structures in a CT image. A statistical volumetric model is constructed from a collection of training CT datasets. The model is represented as a hierarchical tetrahedral mesh equipped with embedded Bernstein polynomial density functions on the barycentric coordinates of each tetrahedron. The prior information of both shape properties and density properties is incorporated in the model. The non-rigid registration process consists of three stages: affine transformation, global deformation, and local deformation. A multiple-layer flexible mesh template matching method is developed to adjust the location of each vertex on the model to achieve an optimal match with the anatomical structure. The mesh template is retrieved directly from the tetrahedral mesh structure, with multiple-layer structure for different scales of anatomical features and flexible searching sphere for robust template matching. An adaptive deformation focus strategy is adopted to gradually deform each vertex to its matched destination. Several constraints are applied to guarantee the smoothness and continuity. A "leave-one-out" validation showed that the method can achieve about 94% volume overlap and 5.5% density error between the registered model and the ground truth model.
Keywords: Multiple-Layer Flexible mesh template, non-rigid registration

1. INTRODUCTION

An anatomical atlas is a tool to model human anatomy and to represent normal anatomical variations. Non-rigid registration between anatomical atlases and medical images can be regarded as a segmentation tool since the registered model is also the segmentation result. The ability to register an anatomical atlas to individual patient images provides the basis for solving several important problems in medical image interpretation. Once the atlas has been registered to a particular image or set of images, structures of interest can be labeled and extracted for further analysis; surgical plans and physician knowledge defined on the atlas can be transferred to the patient. Non-rigid registration of an atlas also allows population studies to be analyzed in a common frame of reference.

Medical images are usually complex, noisy and possibly incomplete, which makes their interpretation and analysis very difficult without prior knowledge of the anatomy. Statistical models have been used in medical image registration by many researchers. Shen *et al* [1] proposed a statistical model of anatomical surfaces using an affine invariant geometric attribute vector to find the vertex correspondences. Chen *et al.* [2] have built an average brain atlas based on statistical data of voxel intensity values. Cootes *et al.* [3] proposed an Active Appearance Model (AAM) that incorporates both the shape variability and density variability. They implemented their models on 2D MRI brain slices and 2D human face images. Their models have also been extended to 3D surface models by Fleute *et al.* [4]. Among the registration techniques, some are based on anatomical features such as points, curves and surfaces [1, 4]. Others depend on the intensity information in the image for registration [2, 3]. Landmark based methods usually need manual segmentation and hard to automate. Intensity based methods can be automatic, but usually very slow due to vast information involved.

Unlike most other registration methods, our method is based on a statistical *volumetric* atlas model incorporating both shape and density properties of the anatomical structure. The remainder of this paper is organized as follows. Section 2 briefly introduces the statistical bone density atlas model and its construction from a set of training images. Section 3 presents our method for non-rigid registration between the statistical model and a CT image. Finally, Section 4 concludes with some discussion and future directions.

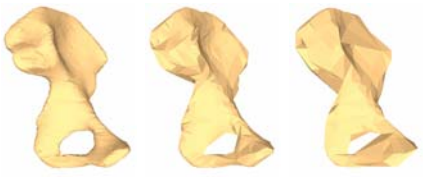

Figure 1. Multiple resolution pelvis model

Figure 2. Density deformation of the model

2. STATISTICAL BONE DENSITY ATLAS MODEL

While most existing anatomical models have only represented the boundary surface or the intensity distribution of anatomical structures, we propose a unique model representation to characterize bone anatomical structures. The atlas is represented as a tetrahedral mesh equipped with embedded Bernstein polynomial density functions on the barycentric coordinates of each tetrahedron. Multiple resolution details of the anatomical structure are described by a hierarchical representation. And the prior information of both shape properties and density properties is incorporated in the atlas.

We propose an efficient and practical "tetrahedral mesh reconstruction from contours" method to construct tetrahedral meshes for bone structures in a CT image. The method can produce tetrahedral meshes with high flexibility and can accommodate any anatomical shape. The meshes align naturally with cortical boundaries within the bones, and minimize residual errors associated with the density representation. We also develop a tetrahedral mesh simplification algorithm based on edge collapsing operations to build a multiple-resolution model representation. We assign an analytical density function to every tetrahedron instead of storing the density value of every voxel within the model. The advantage of such a representation is that it is in an explicit form and is a continuous function in 3D space. Therefore, it is convenient to integrate, to differentiate, to interpolate, and to deform. Currently, the density functions are written as *n*-degree Bernstein polynomials in barycentric coordinates.

We then develop a training strategy to compute a prior model from a collection of training models. In this strategy, we design a model aligning procedure to map all training models into a common mesh structure, and extend the Principal Component Analysis (PCA) method to extract the variation of both shape properties and density properties of the anatomical structure. Using the PCA method, the model can be approximated by a set of model parameters $\{b_i\}$ as:

$$Y = M(\overline{Y}, b) = \overline{Y} + Pb \qquad (1)$$

where Y is a model instance, $M(\overline{Y}, b)$ is the instantiation operation, $\overline{Y}$ is the average model representation, and P is the eigenvector matrix incorporating the prior information.

We have built a statistical density atlas for the hemi-pelvis from eight training images. Figure 1 shows the multiple-resolution pelvis model. Figure 2 shows the density deformation of the model. Details of our model and reconstruction method can be found in [5, 6].

3. NON-RIGID REGISTRATION BETWEEN THE STATISTICAL DENSITY MODEL AND CT IMAGES

We have investigated the non-rigid registration between our statistical bone density atlas and CT images. Unlike most other registration techniques, our method utilizes both the shape properties and the density properties of the model, together with the prior information extracted from a set of training models. The registration procedure can be used to automatically segment the bone structure and generate a patient specific model with the same topological structure as the atlas.

The registration process is divided into three stages: affine transformation, global deformation, and local deformation. The results of the preceding stage are used as the initial value for the next stage. In the affine transformation stage, the translation, rotation, and scale of the atlas model are determined in order to locate the location and orientation of the model. In the global deformation stage, the statistical deformation mode parameters are optimized to align the atlas model with the anatomical structure in the CT images. Due to the limited number of training models in the modeling training stage, the prior information cannot capture all variability inherent in the human anatomy. To compensate for this, a local deformation step is taken to adjust the location of each vertex on the atlas model to achieve a better match.

The registration process is governed by a script file and is highly automatic. Each stage of the registration will be introduced in details in following sections.

3.1 Optimization algorithm and energy function

Both the affine transformation stage and the global deformation stage in the non-rigid registration are essentially multi-dimensional nonlinear optimization problems designed to minimize an energy function between the atlas model and the CT image. Some optimization algorithms [7], such as the gradient descent method, require computing the derivatives of the energy function. In our case, the derivatives of the energy function with respect to both the affine transformation parameters and the global deformation parameters are very difficult to compute. Therefore, we chose Powell's method as our optimization algorithm. To further improve the efficiency and robustness of the algorithm, the process is run in a multiple-resolution framework. This involves first searching for the match in a coarser image using a lower resolution model, and then refining the location in a series of finer resolution images and models. The algorithm is also implemented in a multiple-step-size manner, in which it starts with a large step size and gradually reduces the step size as we get closer to the optimal solution. This leads to a faster algorithm, which is also less likely to fall into a false local minimum.

An energy function is defined as the objective function to evaluate the difference between the atlas model and the CT image, written as:

$$E(mdl,img) = w_s E^{(s)}(mdl,img) + w_d E^{(d)}(mdl,img) \tag{2}$$

$$E^{(s)}(mdl,img) = \sum_{i=1}^{N(v)} \left(\vec{g}^{(mdl)}(v_i) \cdot \vec{g}^{(img)}(v_i) \right) \tag{3}$$

$$E^{(d)}(mdl,img) = \sum_{i=1}^{N(t)} \left(\oint_{\mu} \left(\left(\frac{d^{(mdl)}(t_i,\mu) - d^{(img)}(t_i,\mu)}{d^{(mdl)}(t_i,\mu)} \right)^2 \right) \right) \tag{4}$$

where *mdl* represents the atlas model; *img* represents the CT image; $E^{(s)}$ measures the shape difference; $E^{(d)}$ measures the density difference; $\vec{g}^{(mdl)}(v_i)$ is the surface normal at vertex v_i on the model; $\vec{g}^{(img)}(v_i)$ is the image intensity gradient direction at vertex v_i; $d^{(mdl)}(t_j,\mu)$ is the density value at a voxel within the model; $d^{(img)}(t_j,\mu)$ is the density value at a voxel within the image. $E^{(s)}$ is the sum of the dot product of $\vec{g}^{(mdl)}(v_i)$ and $\vec{g}^{(img)}(v_i)$ over every vertex v_i on the model. $E^{(d)}$ is computed by integrating the density difference over all tetrahedra using barycentric coordinates. w_s and w_d are respective weights for $E^{(s)}$ and $E^{(d)}$ for different applications. For images where edges are more prominent (such as bone images), w_s is supposed to be large. While for images where intensity distributions are more important (such as brain images), $E^{(d)}$ has a larger weight w_d.

3.2 Affine transformation

An affine transformation includes translation $T(t_x, t_y, t_z)$, rotation $R(r_x, r_y, r_z)$, and scale $S(s_x, s_y, s_z)$. After an optimized affine transformation is achieved, the location, orientation, and size of the model should match the anatomical structure. This is a *nine*-dimensional nonlinear optimization problem. It can be written as:

$$\arg\min_{R,S,T} \left(E\left(R \cdot S \cdot mdl + T, img \right) \right) \tag{5}$$

Here $E()$ is the energy function defined in Eq. 2. The affine transformation is divided into three sets of parameters - translation, rotation, and scale; during each pass, only one set of parameters is optimized. We optimize the parameters back and forth several times to get the optimal parameters that minimize the energy function.

3.3 Global deformation

As mentioned in section 2, given a set of model parameters $b=\{b_i\}$, we can generate an instance of the model. The new instance Y of the model is equivalent to a deformed version of the average model. We call the operation of changing the model parameters to obtain a new model instance as the global deformation of the model. The global deformation stage can also be treated as a procedure to minimize an energy function between the atlas model and the CT image. We search throughout the model parameter space in the optimization process. For each parameter set being evaluated, we generate an instance of the atlas. Then we compare this hypothesis with the image, and compute the value of the energy function (Eq. 2). The set of parameters best matching the model with the image is then the parameters that minimize the

Input: A model M before local deformation, and an image I
Output: A model M^* after local deformation

Step 1: Initialize the model;
Step 2: for every vertex v_i in model M
2.1) Construct a multiple-layer flexible mesh template $T(v_i)$ from model M (Section 3.4.1)
2.2) Compute the attribute vector $A^{(m)}(v_i)$ for $T(v_i)$ (Section 3.4.2)
2.3) Find the best matched image location c_i for v_i within searching radius (Section 3.4.3).
end for
Step 3: Adaptively deform $\{v_i\}$. (Section 3.4.4)
Step 4: Repeat Step 2 to Step 3 until the difference between models in two iterations is below a threshold.

Figure 3. Pseudo-code of local deformation procedure

energy function. As in the affine transformation optimization problem, we use Powell's method to optimize the deformation parameter set. The optimization problem can be written as:

$$\arg\min_{b}\left(E\left(M(mdl,b),img\right)\right) \quad (6)$$

Here, $b=\{b_i\}$ is the deformation parameter set of the statistical atlas, $E()$ is the energy function defined in Eq. 2, and $M()$ is the instantiation operation (Eq. 1) to apply the model parameters to the average atlas.

In our experiment on the hemi-pelvis model, we optimize the first five deformation modes. The global deformation provides an initial value for further local deformation.

3.4 Local deformation

Human anatomical structures vary significantly among individuals, yet the statistical atlas can only characterize the variability exhibited in the training set. Given a specific image which is not in the training set, there always exist variations that cannot be reproduced from the statistical model by globally deforming it. Therefore, a local deformation phase is necessary after the global deformation stage to address small discrepancies. The purpose of this local deformation phase is to locally adjust the location of each vertex on the model to match the boundaries of the cortical bones. We propose a multiple-layer flexible mesh template matching method to accomplish the local deformation.

Figure 3 provides the pseudo-code of the local deformation procedure. The details of each step are described in following subsections.

3.4.1 Multiple-layer flexible mesh template

Due to the complex structure of human anatomies, it is not easy to find the correspondence between vertices on the model and voxel locations in the image space. Shen and Davatzikos *et al.* [1] presented an affine invariant attribute vector using the volumes of tetrahedra formed by the vertex neighbors on the surface mesh to find the vertex correspondences. Inspired by their work, we propose a robust method to associate each vertex on the model to its corresponding voxel location.

We utilize the tetrahedral mesh structure of the model to construct a *multiple-layer flexible mesh template* (MLFMT) for each vertex in the model. The mesh template associated with each vertex is retrieved directly from the model, so we do not need extra storage to keep the template. Figure 4 shows a MLFMT in the 2D case, which is a two-layer mesh template with 13 nodes.

The structure of the mesh template of a vertex $v^{(0)}$ can be described as follows: $v^{(0)}$ is the center of the template. The template consists of an array of nodes $\{v_i^{(k)}\}$, where k is the topological distance between $v_i^{(k)}$and $v^{(0)}$ in the mesh. Each node $v_i^{(k)}$ is associated with following attributes: 1) relative position to the center $p_i^{(k)} = v_i^{(k)}-v^{(0)}$; 2) a radius $r_i^{(k)}$; 3) an attribute value; and 4) a weight w, which will be described in Section 3.4.2.

The template associated with a vertex $v^{(0)}$ has multiple layers. $v^{(0)}$ is the center of the template, $\{v_i^{(1)}\}$ are the first layer neighbors of $v^{(0)}$, and $\{v_i^{(2)}\}$ are the second layer neighbors of $v^{(0)}$, and so on. The more layers in the template, the

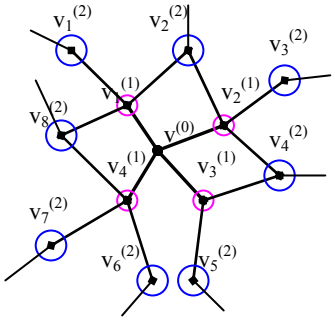

Figure 4 Multiple-layer flexible mesh template

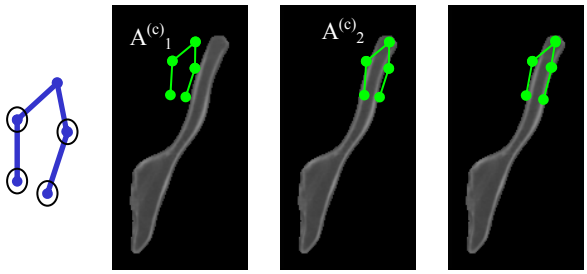

(a) (b) (c) (d)

Figure 5 Flexible template matching result

larger and more global structure can be characterized. Conversely, the fewer layers in the template, the smaller and more local structure it describes. In most cases, a two-layer mesh template is adequate for local deformation.

The template is also flexible. This means instead of just using the locations of nodes on the template, we define a searching sphere $r_i^{(k)}$ for each node. The dashed circles in Figure 4 represent the searching spheres. The searching sphere makes the template deformable, also makes the local deformation stage robust. The center of the template has a searching sphere of zero radius. The nodes further from the template center have larger searching spheres.

The mesh template has the following major advantages over an ordinary fixed-size matrix template: it is retrieved directly from the mesh topological structure; the multiple-layer structure can be used to describe different scale of anatomical features; the flexible searching sphere makes the template match robust.

3.4.2 Template attribute vector and dynamic image attribute vector

Each mesh template has an attribute vector. The attribute vector describes the density and shape properties of the mesh template. Each node on the mesh template is associated with an attribute value $(d;\vec{g})$ and a weight w, where d is the density value of the node and $\vec{g}$ is the gradient direction at the node location. The attribute value of a node $v_i^{(k)}$ can be directly retrieved or computed from the density atlas model. The template attribute vector of the template is basically the concatenation of the attribute value of each node on the template, and can be written as

$$A^{(m)} = (w_0^{(0)} d_0^{(0)}, w_1^{(1)} d_1^{(1)}, w_2^{(1)} d_2^{(1)} \cdots; w_0^{(0)} \vec{g}_0^{(0)}, w_1^{(1)} \vec{g}_1^{(1)}, w_2^{(1)} \vec{g}_2^{(1)} \cdots) \quad (7)$$

In the template matching (details in Section 3.4.3), we move the mesh template of a vertex in the image field and try to find the best matched image location for the vertex. For any voxel location c in the image space, we also define an attribute vector $A^{(c)}$ given the mesh template of vertex v. The image attribute vector has the same format as the template attribute vector $A^{(m)}$ of v, but the attribute values are obtained from corresponding image locations. Because a node $v_i^{(k)}$ on the mesh template is associated with a searching radius $r_i^{(k)}$, when retrieving the attribute value at an image location, we search within a sphere with radius $r_i^{(k)}$ for the most alike attribute value to that of $v_i^{(k)}$.

3.4.3 Flexible mesh template matching

We use a template matching method to find the vertex correspondence because of its simplicity and computational efficiency. Template matching is a simple filtering method to detect a particular shape or object in an image. An object can be detected if its appearance is known accurately in terms of a template. In our case, we want to detect the anatomical structure defined by the mesh structure associated with each vertex on the model. The idea of template matching is straightforward. The template is shifted in the image field, and the difference between the template and the image is calculated at each image position. An optimal match is reported at the position where we get the minimum difference between the template and the image. A shortcoming of the rigid template matching is that it requires that the template be very precise, and it is sensitive to shape deformation. We propose a flexible mesh template to overcome this problem. Each node on the template is associated with a searching radius $r_i^{(k)}$.

The goal of our template matching method is to find the corresponding image location in the image space for each vertex on the atlas model. So that a vertex can be deformed to its corresponding image location. Each vertex has been associated with a mesh template and a template attribute vector. The approach to find the corresponding image coordinate is to shift the mesh template over the image space to find the most similar image attribute vector. Comparing

the difference between the template attribute vector and the dynamic image attribute vector, a best template matching can be found. The following equation is the way to compute the difference between two attribute vectors $A^{(m)}$ and $A^{(c)}$:

$$Diff(A^{(m)}, A^{(c)}) = \frac{1}{len(A^{(m)})} \sum_{i=1}^{len(A^{(m)})} \left(w_i \left| \frac{d_i^{(m)} - d_i^{(c)}}{d_i^{(m)}} \right| \right) + \frac{1}{len(A^{(m)})} \sum_{i=1}^{len(A^{(m)})} \left(\frac{w_i \left(1 - \vec{g}_i^{(m)} \cdot \vec{g}_i^{(c)}\right)}{2} \right) \quad (8)$$

here $A^{(m)}$ is the attribute vector of the mesh template, and $A^{(c)}$ is the image attribute vector of a surrounding image location. It contains two parts, the density vector difference and the gradient vector difference.

In order to demonstrate the flexible mesh template matching, Figure 5 shows a simple example in 2D case. The mesh template in Figure 5a is a simplified 2D two-layer mesh template with five nodes. v_0 is the center of the mesh template; v_1 and v_2 are the first layer neighbors; v_3 and v_4 are the second layer neighbors. Its template attribute vector can be written as:

$$A^{(m)} = (w_0 d_0, w_1 d_1, w_2 d_2, w_3 d_3, w_4 d_4; w_0 \vec{g}_0, w_1 \vec{g}_1, w_2 \vec{g}_2, w_3 \vec{g}_3, w_4 \vec{g}_4) \quad (9)$$

By shifting this mesh template over the image space, we can get image attribute vectors at different image locations, denoted as $A^{(c)}_1$, $A^{(c)}_2$, $A^{(c)}_3$,…. Figure 5b-5d demonstrates some image locations in the template matching process. Figure 5b is a totally unmatched image location. Figure 5c shows a partially matched location by applying a rigid template matching. After deforming the template within the searching sphere of its nodes, a perfect match in Figure 5d is found. The difference between Figure 5c and Figure 5d demonstrates the idea of the flexible template matching. Therefore, the image location of $A^{(c)}_3$ is treated as the corresponding image location of the center of the mesh template, i.e. v_0.

3.4.4 Adaptive deformation and constraints

If any deformability is allowed, one object can always be deformed into any other objects (i.e. a head is morphed into a teapot). Some constraints are necessary for local deformation, especially for meaningful anatomical structure matching. In order to make the local deformation smooth and keep the tetrahedral mesh structure valid, we adopt several strategies, including Gaussian morphing, adaptive deformation focus, and maximum deformation range. Among these techniques, adaptive focusing and Gaussian morphing are inspired by Shen's work [1].

To keep the deformation continuous and smooth, a Gaussian morphing strategy is adopted for the vertex deformation. In this strategy, when one vertex is moved, its neighbors will also be morphed accordingly. The following equation describes the Gaussian morphing:

$$\Delta v_l = \Delta v_0 \cdot e^{-\frac{l^2}{2\sigma^2}} \quad (10)$$

here Δv_0 is the movement of the deformed vertex, Δv_l is the movement of its l^{th} layer neighbors, σ is the morphing parameter. The movement of a neighbor vertex is controlled by the topological distance to the center vertex.

In addition, we apply an adaptive deformation focus strategy, i.e. the vertex with highest matched image attribute vector in a neighborhood will deform first and drag its neighbors to morph. Then, the focus will move to the next highest matched vertex. The adaptive focusing can prevent the deformation from running randomly by restricting the deformation to the prominent matches.

The third strategy to prevent the invalid deformation is to assign a maximum deformation range $MDR(v_i)$ to each vertex v_i, which basically restricts the possible deformation range of a vertex within a sphere. The radius of the maximum deformation range is defined as half of the distance between a vertex and its closest neighbor. This range is adaptively updated at the end of each iteration.

Under these constraints, the local deformation will not go "wild" and change the appearance shape of the model. Instead, it will gradually morph the model to make it better reflect the anatomical structures.

3.5 Results

We have tested the non-rigid registration algorithm using the statistical hemi-pelvis density atlas. The statistical model is computed from eight training CT images (Section 2). We then took an additional pelvis CT image, and applied the non-rigid registration between the atlas model and the CT image. Figure 6 shows some visual results of the non-rigid registration at difference stages. The first row is a 3D visualization of the process, and the second row shows one cross section of the data set. The atlas model is superimposed on the CT image to show the matching between the model and the anatomical structure. Figure 6a is the initial state of the model. Figure 6b is the result after the affine transformation. Figure 6c is the result after the global deformation by optimizing five global deformation parameters. After this stage, the shape of the model almost matches the anatomy, but still some small differences exist due to limited variability in the

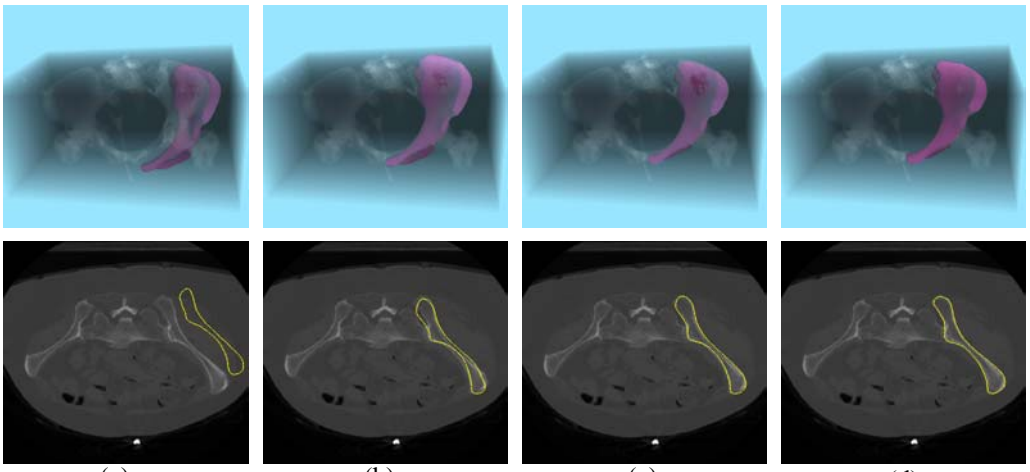

(a) (b) (c) (d)

Figure 6. Visual results of non-rigid registration between the statistical model and CT images

statistical model. Finally, Figure 6d is the result after the local deformation, which reveals a very close match between the image and the deformed model.

To quantitatively validate our method, we have implemented volume and density-based error metrics. The error metric measures the difference of density values of voxels within two volumetric models. Assume model 1 is the model produced by the registration procedure, and model 2 is the ground truth model. The voxels within the two models are obtained by a volume scanning algorithm. Two scans of voxels are compared in the following ways: 1) the percentage of the number of overlapping voxels to the total number of voxels, written as

$$Overlap = \|V_1 \cap V_2\| / \|V_1\| \times 100\% \quad (11)$$

where V_1 is the set of voxels in model 1, V_2 is the set of voxels in model 2, and $\|\cdot\|$ represents the size of a set; and 2) the statistics of density differences between voxels.

Table 1. Leave-one-out validation of atlas/CT registration method for the half pelvis atlas

	Initial stage	Affine transformation		Global deformation		Local deformation (final result)			
Data set	Overlap	Avg	Overlap	Avg	Overlap	Avg	Stdev	Max	Overlap
1	0.0%	7.4	84.5%	6.7	91.3%	6.0	5.1	39.2	96.4%
2	5.4%	6.7	82.4%	6.2	89.1%	5.3	4.7	40.9	94.3%
3	0.0%	6.9	82.4%	6.1	88.6%	5.3	4.7	42.7	94.1%
4	0.0%	6.2	78.2%	5.4	89.1%	5.1	4.6	42.5	93.2%
5	0.0%	6.5	81.1%	6.8	89.3%	5.6	5.4	36.1	94.1%
6	1.3%	6.8	84.6%	6.1	91.6%	5.9	5.0	40.6	95.6%
7	3.8%	6.7	90.2%	5.9	94.4%	5.4	4.8	36.4	97.4%
8	0.0%	6.7	79.2%	6.3	88.4%	5.9	5.2	42.4	92.2%
Avg	1.31%	6.7	82.81%	6.2	90.21%	5.5	4.9	40.1	94.67%
Std Dev	2.1%	0.3	3.7%	0.4	2.1%	0.3	0.3	2.6	1.7%

Currently, we have eight CT images in the pelvis training set. We conducted a "leave-one-out" validation on the training data set. That is, we selected one data set in the training set as the testing data set, and used the other seven data sets to build a statistical atlas. We then registered the statistical atlas with the testing data set. We used the original test model as the ground truth model, and compared the registration result with the ground truth model to validate the registration. Table 1 lists the results of this experiment. We evaluated the results at the beginning of registration and at the end of each stage. The percentage of overlapping voxels versus all voxels is reported. We also report the statistics of

the density difference at corresponding voxel locations between the transformed atlas model and the ground truth model. In the table, "Avg" is the average voxel density difference, and "Stdev" is the standard deviation of density differences. The density values of the data set have been scaled to the range [0-255]. "Overlap" is the percentage of overlapping voxels to all voxels within the model. From the result, the average overlap between the ground truth model and the registered atlas is 94.7%.

4. DISCUSSIONS

We have presented a new automatic method for non-rigid medical image registration and model assisted segmentation using a statistical density atlas model. The model is represented as a tetrahedral mesh and contains both shape and density properties and their variations.

The non-rigid registration procedure is divided into three stages. First, the affine transformation is determined to set the pose and size of the model. Then, global deformation is computed by optimizing statistical deformation parameters. To compensate for the limited deformation allowed in the global deformation, we allow further local deformation driven by a multiple-layer flexible mesh template associated with each model vertex. The template itself is flexible and non-rigid. We find out that allowing local deformation can lead to more accurate matching, and that it can be more effective than simply adding more global modes. The registration process is governed by a script file, and the control parameters are provided by a user. A more intelligent method needs to be developed to automate the process. Atlas models built from a larger training set will also improve the accuracy.

In future works, we plan to build a more computerized bone density atlas based on the statistical bone density model presented in this paper. In this atlas, anatomical landmarks are marked and certain surgical procedures are defined. We also plan to further extend this model to other anatomies such as knees and vertebras. In our other work [6], we have demonstrated 2D/3D non-rigid registration between the atlas model and a set of X-ray images. In general, this technique can be used to perform 3D patient specific modeling and analysis without the patient specific CT image. Clinical tasks such as planning from X-rays, intra-operative guidance, post-operative analysis and retrospective studies are also potential applications of our atlas model.

ACKNOWLEDGEMENTS

This work was partially funded by NSF Engineering Research Center grant EEC9731478. We thank Shadyside Hospital for providing the pelvis data. We specially thank Dinggang Shen for his suggestion on deformable registration algorithm. We also thank Jerry Prince, Christos Davatzikos, Stephen Pizer, Sandy Wells, and Chengyang Xu for their useful discussions.

REFERENCE

1. Shen, D. and C. Davatzikos. *Adaptive-Focus Statistical Shape Model for Segmentation of 3D MR Structures*. in *MICCAI 2000*. 2000. Pittsburgh, PA.
2. Chen, M., *et al.*, *3-D Deformable Registration of Medical Images Using a Statistical Atlas*. PhD Thesis, 1998, Carnegie Mellon University: Pittsburgh, PA.
3. Cootes, T.F. and C.J. Taylor, *Statistical Models of Appearance for Computer Vision*, 2000.
4. Fleute, M. and S. Lavallee. *Nonrigid 3-D/2-D Registration of Images Using Statistical Models*. in *MICCAI 1999*. 1999. Cambridge, UK.
5. Yao, J. and R. Taylor. *Construction and Simplification of Bone Density Models*. in *SPIE Medical Imaging 2001*. 2001. San Diego, CA.
6. Yao, J., *A statistical bone density atlas and deformable medical image registration*, PhD Thesis, 2001, the Johns Hopkins University: Baltimore, MD.
7. Press, W.H., *et al.*, *Numerical Recipes in C*. Second ed. 1992: Cambridge University Press.