Integrating Statistical Models of Bone Density into Shape Based 2D-3D Registration Framework

Gouthami Chintalapani[1], Ofri Sadowsky[1], Lotta M. Ellingsen[2], Jerry L. Prince[2], and Russell H. Taylor[1]

[1] Department of Computer Science
[2] Department of Electrical and Computer Engineering
The Johns Hopkins University, Baltimore, MD, USA
{gouthami,rht}@jhu.edu

Abstract. We present a framework to deformably register simulated X-ray images to a combined statistical model of pelvis anatomical structure, created from a population of CT scans. The primary contributions are: 1) a framework to create and analyze bone density variations, separate from shape variations and 2) an augmented 2D/3D registration framework that couples shape and density priors to create accurate patient specific models. Our statistical model representation consists of a tetrahedral mesh for approximating bone shape and Bernstein polynomials defined within each tetrahedron for bone density. All datasets in the given population are registered deformably to a template CT dataset. The shape and density statistics are extracted using principal component analysis on the corresponding mesh vertices and voxels of the shape-free deformed subjects respectively. In the registration framework, we register the 2D input images to the 3D shape prior and estimate the bone density parameters in a least-squares like setup by projecting the density modes on to the input image space. This approach was tested using leave-n-out experiments, with n = 8, datasets using an atlas of 63 full pelvis CT datasets.

1 Introduction

Statistical modeling and analysis of anatomical shapes is a promising research area in medical imaging with a variety of applications such as segmentation, 2D/3D registration and reconstruction of anatomical structures, pre-operative planning, analyzing population variations [1][2]. Statistical models are primarily built from principal component analysis (PCA) on point distribution models representing anatomical shapes [3] or from deformation fields [4]. More often than not, point distribution models characterize shape variations only. Hence, when used in applications such as 2D/3D registration, the resulting models would match the input images in shape but not in bone density properties. Cootes *et al.* have proposed a method to create active appearance models (AAM) by concatenating shape and intensity vectors into a single vector and computing PCA

on the resulting matrix [3]. A weight factor W_s is introduced to scale the shape vector so that the units of the shape and the intensity vectors are normalized. Following this work, 3D shape-intensity atlases were created by combining shape and intensity vectors, with $W_s = 1$, as seen in [5], [6], [7]. After computing PCA on the combined vector, the shape and intensity modes were separated and re-normalized. The correlation between the bone shape and the density sub-spaces is not clearly explained in these works and needs to be explored if the shape vectors were to be combined with the intensity vectors.

The application of such shape-intensity atlas in a 2D/3D registration framework can be seen in [5], [8], [9]. Hurvitz *et al.* have proposed a combined atlas with three components consisting of a shape model, a CT-like reference image and a reference surface of the bone and an inverse warping method to estimate the intensity properties [9]. Their model consists of only mean intensity and does not incorporate any intensity variations present in the population. Yao *et al.* [5] and Steininger *et al.* [8] use shape and intensity variations in the registration framework. The proposed approach in these works is to project the prior model and estimate model parameters from the shape-intensity priors that maximizes a given similarity metric. This method requires model instantiation and projection in each step of the optimization algorithm which could be computationally expensive.

In this paper, we extend an existing 2D/3D deformable registration method to incorporate bone density variations. The primary contributions are: 1) a framework to create and analyze bone density variations, separate from shape variations and 2) an augmented 2D/3D registration framework that couples shape and density priors to create accurate patient specific models. We create two models, one for shape and one for CT intensities and integrate them into a 2D/3D registration method. We present alternate methods of computing density statistics in both the CT voxel intensity space and the polynomial approximation. In the registration framework, we register 2D patient images to a shape statistical model first, and successively estimate the density parameters in a least-squares setup. This registration method enables us to create CT-like patient specific models from the 2D X-ray images that match the shape and density properties of the patient. The primary difference between the above cited works and our proposed approach is that we recover the density parameters in a single step by linearly projecting the density model on to the 2D input image space. After the rigid and shape registration of the model, projection images are created for each density mode. The density parameters are computed such that the similarity between the linear combination of these projection images and the input images is maximized. The rest of the paper is organized as follows. We present our density atlas construction and registration method in Section 2, experiments and results in Section 3, followed by discussion in Section 4.

2 Method

In this section, we briefly outline our earlier work on shape statistical models, discuss our new framework to create and integrate density models into a 2D/3D registration method.

2.1 Model Representation

Following Yao's work [5], we use a tetrahedral mesh to represent bone shape and Bernstein polynomials to approximate bone density or CT numbers within each tetrahedron. A tetrahedral mesh is defined by a set of vertices P and a list of tetrahedra T. Each vertex P_i is a point in 3D. The CT densities in each tetrahedron are approximated using Bernstein polynomials. Given a tetrahedron $T_j = (v_0, ..., v_3)$, where v_i is a vertex, they can be arranged in a homogeneous matrix of the form [2]:

$$M_T = \begin{bmatrix} v_0 & v_1 & v_2 & v_3 \\ 1 & 1 & 1 & 1 \end{bmatrix} \tag{1}$$

The density of each cell is approximated using a local Bernstein polynomial function inside the volume of the cell:

$$f_u^j = \sum_{\|k\|=d} \beta_k^j B_k^d(u) = \sum_{\|k\|=d} \beta_k^j \binom{d}{k} u_0^{k_0} u_1^{k_1} u_2^{k_2} u_3^{k_3} \tag{2}$$

Here, $u = (u_0, u_1, u_2, u_3) = M_T^{-1} x$ are the *barycentric coordinates* of a point x in homogeneous coordinates; $k = (k_0, k_1, k_2, k_3)$ is the power index of the basis function B_k^d, with d being the degree of the polynomial; $\binom{d}{k} = \frac{d!}{k_0!k_1!k_2!k_3!}$ is a multinomial factor; and β is a free coefficient. The advantage of this polynomial representation is that it enables us to compute line integrals in a closed form to create digitally reconstructed radiographic (DRR) images in the 2D/3D registration algorithm [2]. Polynomial representation can also be viewed as a dimensionality reduction step.

2.2 Shape Statistical Model

From a population of CT datasets, $V_i, i = 1, 2...n$, we manually select a template dataset, V_T, segment and mesh the anatomy of interest resulting in a template mesh, M_T. We register each CT dataset, V_i, deformably to the template dataset using a 3D grayscale deformable registration method proposed in [10]. This results in a warped volume, V_i^T and a deformation field, D_T^i. We interpolate this deformation field at each vertex location of the template mesh to create mesh instances of the subjects. The vertices of these mesh instances define our point distribution model. We compute PCA on the mesh vertices to create a shape model consisting of mean shape and dominant shape modes. This shape model is further refined by using an iterative bootstrapping technique proposed in [11]. With this method, we update the atlas as datasets are being added and stabilize

the modes by registering the datasets to the atlas iteratively. This method yields consistent shape models. Given the linear parameterization of the shape prior, any new shape instance can be approximated as follows:

$$S = S^0 + \Sigma_{j=1}^{m} \lambda_j S^j \tag{3}$$

where $S^0$is the mean shape, S^j are orthonormal shape mode vectors, and λ_j are mode weights/parameters.

2.3 Density Statistical Model

For each dataset V_i, we have a deformation field, D_T^i, defining the vector flow from the template to the subject space, and a warped volume,V_i^T, obtained by deforming the subject into the template. These warped volumes are shape-independent, i.e., they resemble the template shape, but the intensities are derived from the individual subject datasets. The voxels of these volumes define PDM for analyzing bone density properties. To extract intensity variations, we stack these shape-free volumes into a big matrix and perform PCA on the data matrix resulting in a mean density volume and dominant density modal volumes (see Figure 1(a)).

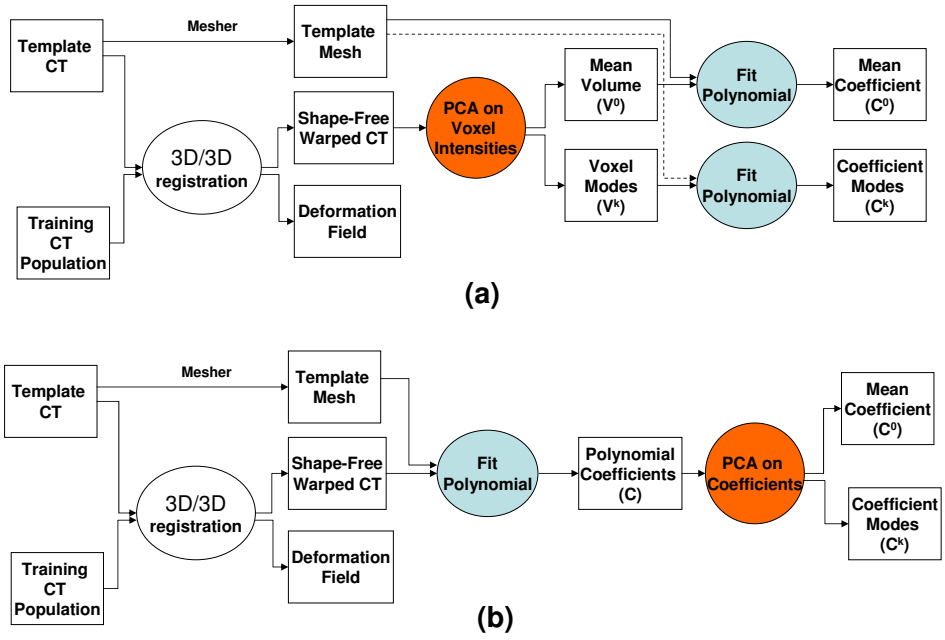

Fig. 1. Density Statistical Model Construction Pipeline: (a) Voxel based density statistical model; (b) Polynomial representation based denstity statistical model;

The parametric linear model of volumetric density model is given as follows.

$$V = V^0 + \Sigma_{k=1}^{n} \mu_k V^k \tag{4}$$

To map voxel-based density modes to polynomial space, we fit 3^{rd} degree Bernstein polynomials to the template mesh and modal density volumes. This is achieved by solving for the unknown coefficients β in equation (2), where f_u^j are derived from the density modes and u from the template mesh.

Alternatively, we fit polynomials to the master mesh M_T and the deformed subjects V_i^T to create polynomial approximation of the intensities of each subject. We then perform PCA on the polynomial coefficients resulting in polynomial density modes (See Figure 1(b)). The linearized parametric form is given as follows:

$$C = C^0 + \Sigma_{k=1}^{n} \mu_k C^k \tag{5}$$

where C^0 is the mean intensity polynomial, C^k are orthonormal intensity mode vectors, and μ_k are mode weights/parameters. To map the polynomial space to voxel space, we transform voxel coordinates to barycentric coordinates, relative to the tetrahedron in which the voxel is contained and then evaluate the polynomial at these points, as shown in equation (2). Here the goal is to compute f_u^j matrix, given the coefficients β and the 3D location, u.

2.4 Registration Framework

We present a new registration framework that combines the shape atlas with a density atlas to create accurate patient specific models (See Figure 2). We use a mutual information based rigid 2D/3D registration method presented in [12], although any such 2D/3D deformable registration method would satisfy our requirements in this work. A detailed analysis of the shape based 2D-3D registration method used in this paper is given in [12]. In this method, projection images or DRRs are created by computing line integrals of the density polynomials along the lines of sight through the space of the mesh cells. This deformable registration method estimates the approximate pose and shape of the patient anatomy by matching the atlas to a set of 2D input images. After estimating the shape, we create projection images of the atlas with various density modes from the intensity atlas. We then solve for the density parameters in a least-square like method that maximizes the similarity between the linear combination of the density projections and the input images. The density parameters are estimated in a single step with this method.

The pseudocode for the algorithm is as follows:

1. *Input:* X-ray images/DRRs I_i, where i is the image number/view angle, camera pose parameters (intrinsic and extrinsic) P_i, mean shape S^0 and shape modes $S^k, k = 1, 2, ..n$, mean density C^0 and density modes $C^l, l = 1, 2, ..m$
2. Register the input images to the shape model, rigidly and deformably to estimate the pose and shape parameters respectively. The output consists of a pose $F(R, T, S)$ and mode weight parameters λ such that

$$(F, \lambda) = \underset{(F,\lambda)}{\operatorname{argmax}} \sum_i MI\left(I_i, DRR\left(F.\left(S^0 + \Sigma_{j=1}^{n} \lambda_j S^j\right)\right)\right) \tag{6}$$

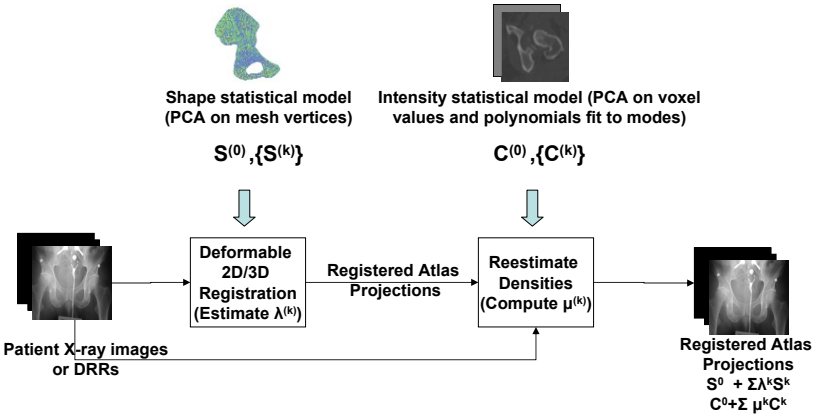

Fig. 2. Registration framework flowchart

3. Transform the shape model and create DRRs of the registered shape model with mean density and density modes.

$$S^{est} = F * \left(S^0 + \Sigma_{j=1}^{m} \lambda_j S^j\right) \tag{7}$$

$$d_i^{mean} = DRR\left(S^{est}, P_i, C_0\right) \tag{8}$$

$$d_i^k = DRR\left(S^{est}, P_i, C_k\right) \tag{9}$$

where d_i^k is the k^{th} density mode projected in i^{th} view direction

4. Formulate a least-squares problem to solve for density mode parameters μ such that

$$\mu = \underset{\mu}{\operatorname{argmin}} \sum_i \left(I_i - \left(d_i^{mean} + \Sigma_{k=1}^{m} \mu_k d_i^k\right)\right)^2 \tag{10}$$

5. Generate the new patient specific model by sampling the estimated shape model on a voxel grid using the estimated coefficients

$$C^{est} = C^0 + \Sigma_{k=1}^{m} \mu_k C^k \tag{11}$$

$$CT_{patient}^{est} = Voxelize\left(S^{est}, C^{est}\right) \tag{12}$$

3 Experiments and Results

3.1 Atlas Experiments

A statistical model of full pelvis anatomy is created from CT scans of 63 healthy individuals. The template mesh model consists of 26875 vertices and 105767 tetrahedra. We have used a $3rd$ order Bernstein polynomials to approximate the CT intensities. The original CT datasets are downsampled from 512x512x82 voxels with voxel size of $0.9375mm^3$ to 256x256x128 voxels and $1.875mm^3$. The

original CT intensities are specified in Hounsfield units (HU) as signed integers. The dynamic range is adjusted such that all the voxels have positive intensity numbers by adding +1000 (HU for air is -1000) to the original units. For volumetric density model, we cropped the volumes to contain 165x95x128 voxels. The projection images of the shape atlas and the density atlas are shown in Figures 3 and 4 respectively.

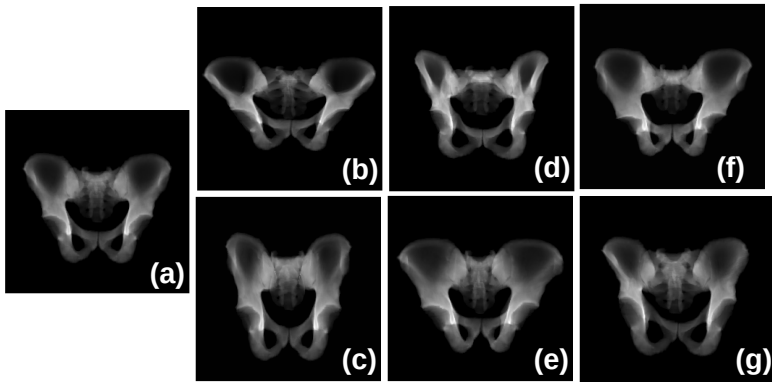

Fig. 3. DRRs of mean shape and first three principal modes. (a) mean shape (S^0); (b)-(c) Mode 1($S^0 \pm 3\lambda_1 S^1$); (d)-(e) Mode 2 ($S^0 \pm 3\lambda_2 S^2$; (f)-(g) Mode 3 ($S^0 \pm 3\lambda_3 S^3$)

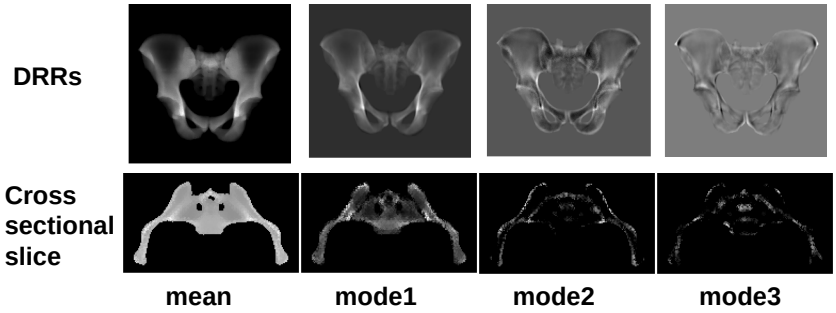

Fig. 4. Top row: DRRs of mean density and first three principal modes (left to right) from polynomial coefficient based density model. Bottom row: An example cross sectional slice of volumes representing mean density and the first three principal modes (left to right) from the volumetric density model

To assess these models, we have performed leave-one-out experiments in which an atlas is created from all but one dataset and the left-out dataset is

reconstructed from the model. For shape atlas, we have measured surface distance between the left out subject and the reconstructed subject from the model. Our analysis shows that we can estimate any given instance with an average accuracy of 1.5 mm using the first 15 dominant shape modes. For density atlas, we have measured the RMS error between the estimated volume and the left out volume in Hounsfield Units.

$$E(V^{true}, V^{est}) = \sqrt{\frac{1}{n_{vox}} \sum_{i=1}^{n_{vox}} \{V^{true} - V^{est}\}^2} \tag{13}$$

Since our registration method requires the density modes in polynomial representation, we have computed the density models in both the volumetric and polynomial space. The leave-out analysis for both volumetric and polynomial based density models is shown in Figure 5. The results show that the density modes from these two representations (voxel and polynomial), are identical and the residual errors from the intensity atlases are comparable. By identical modes, we mean that the intensity mode contributions are similar, i.e., similar anatomical regions are bright in both the images.

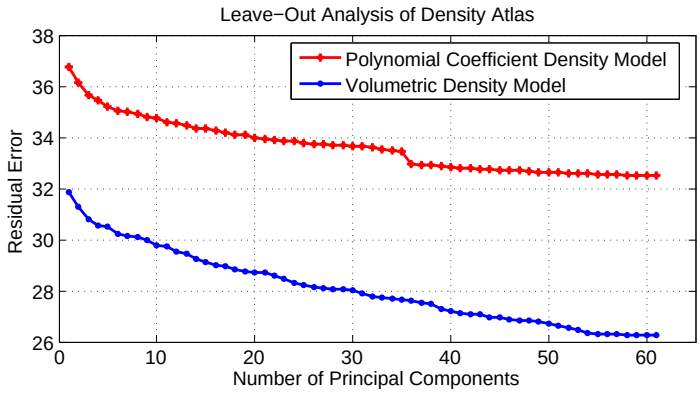

Fig. 5. Root Mean Squared residual errors from leave-out validation using density models from polynomial coefficients and shape free warped volumes

3.2 Registration Experiments

The proposed method is validated using a leave-n-out validation method, with $n = 8$. Eight datasets were randomly selected from a population of 63 datasets. The remaining 55 datasets were used to build the shape and density statistical models. Projection images (DRRs) were created for these 8 datasets with the image size of 512x512 pixels, at views $0°, 90°, 135°$. Based on our analysis presented

above and in [11], we have retained 15 shape modes and 12 density modes along with the mean shape and mean density in the atlas. For each dataset, all three projection images are registered to the shape atlas simultaneously to estimate the shape parameters.

Prior to estimating the density parameters, we render the estimated shape with mean density (1) and density modes (12 modes) for each view angle (3 views) resulting in (12+1) x 3 = 39 DRRs for each left-out subject. After solving for the density parameters in a least-squares setup given in equation (10), we estimate the new polynomial coefficient using equation (5). A CT-like volume is created using the estimated shape and the polynomial. The root mean squared residual errors are computed between the approximated CT and the original CT. Since the goal of the proposed method is to create patient specific models that match both in shape and intensity properties, we chose to use the root mean squared error metric, (See equation 13), between voxel intensities to validate our approach. As we are using DRRs created from the original CT datasets, we could say that the spectrum of CT intensities for both the volumes and the polynomial representation as well as the input DRRs is the same and hence this error metric would be valid. When considering registration with real X-ray images, an intensity calibration step is needed to map the intensities of the input X-ray images to the underlying CT intensities in the density model. The results from these leave-out experiments are shown in Table 1.

(1)	(2)	(3)	(4)	(5)
#	$\mathbf{S^{true}}$– $\mathbf{S^{est}}$	**RMS (** $\mathbf{V^{true}}$, $\mathbf{V^{est}_{mean}}$**)**	**RMS(** $\mathbf{V^{true}}$, $\mathbf{V^{est}_{modes}}$ **)**	Δ ((3)-(4))/(3)
	(mm)	(HU)	(HU)	%
1	1.94	109.92	58.88	**46.43**
2	1.62	128.32	96.0	**25.19**
3	1.90	98.4	77.12	**21.63**
4	2.60	51.68	41.6	**19.50**
5	2.48	109.44	84.8	**22.51**
6	1.95	73.44	50.56	**31.15**
7	2.30	72.96	47.52	**34.84**
8	2.93	101.28	85.76	**15.32**
avg	**2.21**	**93.18**	**67.78**	**27.07**

Table 1. Residual errors from leave-out-validation tests of the augmented registration algorithm with the combined atlas of shape and density variations. Columns 3 shows residual errors when using mean density only and column 4 shows residual errors with mean density and density modes. The % reduction in RMS error between columns 3 and 4 is given in Column 5.

The second column shows the accuracy of the shape registration in terms of the surface distance between the registered atlas shape, S^{est} and the original shape S^{true}. The third column presents the root mean squared error between

the original CT volume, V^{true}, and the CT-like approximation of the registered atlas with mean density only, V^{est}_{mean}. The fourth column presents the residual error between the original CT, V^{true}, and the CT-like approximation of the registered atlas with mean density and density modes, V^{est}_{modes}. Column 5 shows the % reduction in the residual errors from using mean density (column 3) to the density modes (column 4). The error values are specified in Hounsfield units.

4 Discussion and Conclusion

We have proposed a framework to analyze bone density variations in a given population. We presented alternate ways of computing density models in either the voxel space or the polynomial space. Our analysis shows that the first 12 eigen modes explain 80% of the variation. The augmented registration framework shows an average reduction of 27% in terms of root mean squared error of the voxel intensities in Hounsfield units, measured after compensating for the CT intensities with the density model. The best case error reduction was 46%. These results show that the bone density properties can be recovered in a single step by linearly projecting the 3D density modes onto the 2D image space. And the resulting 3D registered models exhibit similar bone density properties as the input X-ray images. Although, we have demonstrated the method using a specific 2D/3D registration algorithm, the extension presented here can be applied to any volumetric or mesh-based 2D/3D registration algorithm.

The accuracy of the density atlas depends on the deformable registration method as well as the number of datasets in the population. The sample size is small compared to the dimension of the input vector. We believe that an increase in the population size would result in more compact model generation. We are currently working towards increasing the atlas population to 150 datasets. With a more accurate density model, the registration accuracies can be further improved. Although the accuracy of the density registration depends on the accuracy of shape registration, our results indicate that with a reasonable shape registration, we are still able to achieve a significant reduction in RMS error values with the density model. We are exploring ways to refine the shape registration by estimating the shape parameters and the density parameters in an iterative approach. Other experiments such as using real X-rays instead of DRRs, truncation, partial anatomy etc. to understand the behavior of the density properties in relation to the shape properties need to be performed. However, we believe that the results presented here are promising and validate the hypothesis that the combined shape and intensity atlas would yield better patient specific models.

Acknowledgments

This work was funded in part by NIH grants 1-R01-EB006839-01A1 and 1-R21-EB003616-01, in part by NSF ERC cooperative agreement EEC9731478, and by

Johns Hopkins University internal funds. We would also like to thank Dr. Ted DeWeese and Dr. Lee Myers for providing us with the CT datasets.

References

1. Tang, T., Ellis, R.: 2d/3d deformable registration using a hybrid atlas. In: MICCAI. Volume LNCS 3750, Part II. (2006) 223–230
2. Sadowsky, O.: Image Registration and Hybrid Volume Reconstruction of Bone Anatomy Using a Statistical Shape Atlas. PhD thesis, The Johns Hopkins University (2008)
3. Cootes, T., Beeston, C., Edwards, G., Taylor, C.: A unified framework for atlas matching using active appearance models. In: IPMI. (1999) 322–333
4. Rueckert, D., Frangi, A., Schnabel, J.: Automatic construction of 3d statistical deformation models using non-rigid registration. In: Medical Image Computing and Computer Assisted Intervention. (2001) 77–84
5. Yao, J.: A statistical bone density atlas and deformable medical image registration. Ph. D. Thesis, The Johns Hopkins Univeristy, 2002.
6. Fritscher, K.D., Grunerbl, A., Schubert, R.: 3d image segmentation using combined shape-intensity prior models. International Journal of Computer Assisted Radiology and Surgery **1** (Februaury 2007) 341–350
7. Querol, L., Buchler, P., Rueckert, D., Nolte, L.P., Ballester, M.: Statistical finite element model for bone shape and biomechanical properties. In: R. Larsen, M. Nielsen, and J. Sporning (Eds) MICCAI 2006. Volume 4190. (2006) 405–411
8. Steininger, P., Fritscher, K.D., Kofler, G., Schuler, B., Hanni, M., Schweiger, K., Schubert, R.: Comparison of different metrics for appearance-model-based 2d/3d registration with x-ray images. In: Proc. BMV, Berlin. (2008) 122–127
9. Hurvitz, A., Joskowicz, L.: Registration of a ct-like atlas to fluoroscopic x-ray images using intensity correspondences. International Journal of Computer Assisted Radiology and Surgery **3** (October 2008) 493–504
10. Ellingsen, L., Prince, J.: Deformable registration of ct pelvis images using mjolnir. In: IEEE 7th Nordic Signal Processing Symposium (NORSIG). (2006)
11. Chintalapani, G., Ellingsen, L., Sadowsky, O., Prince, J., Taylor, R.: Statistical atlases of bone anatomy: Construction, iterative improvement and validation. In: MICCAI. Volume 4791/2007. (2007) 499–506
12. Sadowsky, O., Chintalapani, G., Taylor, R.H.: Deformable 2d-3d registration of the pelvis with a limited field of view, using shape statistics. In Ayache, N., Ourselin, S., Maeder, A., eds.: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2007. Volume 4792 of LNCS., Springer (2007) 519–526