Fully Automatic X-Ray Image Segmentation via Joint Estimation of Image Displacements

Cheng Chen[1], Weiguo Xie[1], Jochen Franke[2], Paul A. Grützner[2], Lutz-P. Nolte[1], and Guoyan Zheng[1]

[1] Institute for Surgical Technologies and Biomechanics, Universität Bern, Switzerland
[2] BG Clinic Ludwigshafen, Ludwigshafen, Germany

Abstract. We propose a new method for fully-automatic landmark detection and shape segmentation in X-ray images. Our algorithm works by estimating the displacements from image patches to the (unknown) landmark positions and then integrating them via voting. The fundamental contribution is that, we jointly estimate the displacements from all patches to multiple landmarks together, by considering not only the training data but also geometric constraints on the test image. The various constraints constitute a convex objective function that can be solved efficiently. Validated on three challenging datasets, our method achieves high accuracy in landmark detection, and, combined with statistical shape model, gives a better performance in shape segmentation compared to the state-of-the-art methods.

1 Introduction

Segmenting anatomical regions such as the femur and the pelvis in the clinical X-ray images provides invaluable information for computer aided diagnosis [1,2], surgery planning and image-guided intervention [3], and three-dimensional (3D) model reconstruction [4,5]. Manual landmarking and segmentation are both time-consuming and error-prone. Therefore, automatic landmark detection and shape segmentation techniques have been an active research topic in medical image analysis [5]-[10].

Landmarking and segmentation X-Rays has to deal with many challenges such as poor image illumination, unknown image projection and unexpected appearance caused by trauma or implants. Many algorithms have been proposed, such as methods based on local image features [1,3] and model-based methods [5,6]. A large body of methods rely on machine learning techniques, such as the shape regression machine proposed in [7] and the marginal space learning method in [8]. Perhaps the most popular method up to now is the voting scheme based on random forest (RF) regression [9,10], which estimates the displacements from a set of sampled image patches to the landmark by random forest (RF) regression, and then the landmark position is estimated by aggregating the votes made by all the patches. In [2], Lindner et al. combine this method with a constrained local model (CLM) for automatic segmentation of proximal femur.

K. Mori et al. (Eds.): MICCAI 2013, Part III, LNCS 8151, pp. 227–234, 2013.

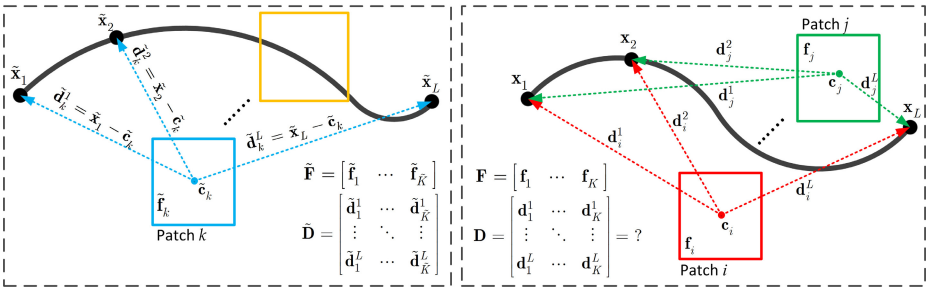

Fig. 1. Schematic illustration of the semi-supervised joint estimation of image displacements. Left: training data. Right: test data.

There are two key components behind the success of RF voting based method. The first is the strategy of positioning landmarks by estimating its relative displacements with regard to other image parts (e.g. patches). This is reasonable since medical image is highly structured. The second is the discriminative power of the RF model. In this paper, we focus on improving the estimation the displacements from image patches to landmarks. In previous methods based on RF regression, the displacement from each patch to each landmark is estimated independently by the pre-trained RF model, i.e. the displacements are determined solely by the training data. Our method is fundamentally different, as we estimate the displacements from all patches to multiple landmarks all together. This joint estimation scheme allows us to exploit the mutual interactions among the displacements that are being estimated by considering the geometric constraints. In this way, our joint displacement estimation method achieves better accuracy. We tested our method on three large and challenging datasets: CompleteFemur, ProximalFemur and Pelvis. The experimental results show that our method achieves better performance compared to the state-of-the-art methods.

2 Landmark Detection by Joint Displacement Estimation

2.1 Problem Formulation

Training Data. Assume that we are interested in L landmarks, and the ground-truth position of these landmarks are known in a set of training images. As shown in Fig. 1 (left), $\tilde{\mathbf{x}}_l \in \mathbb{R}^2$ is the position of the lth landmark. Furthermore, we randomly sample a number of square patches around all the landmarks. For the kth patch, we denote $\tilde{\mathbf{c}}_k \in \mathbb{R}^2$ as its center position, $\tilde{\mathbf{f}}_k \in \mathbb{R}^{d_f}$ as its visual feature, and $\tilde{\mathbf{d}}_k^l = \tilde{\mathbf{x}}_l - \tilde{\mathbf{c}}_k \in \mathbb{R}^2$ is the displacement from the kth patch center to the lth landmark. We totally sample $\tilde{K}$ patches over all the training images, and we denote $\tilde{\mathbf{F}} = [\tilde{\mathbf{f}}_1, ..., \tilde{\mathbf{f}}_{\tilde{K}}] \in \mathbb{R}^{d_f \times \tilde{K}}$ as the matrix of features of all training patches, and $\tilde{\mathbf{D}} \in \mathbb{R}^{2L \times \tilde{K}}$, whose element $\tilde{\mathbf{D}}_{ij} = \tilde{\mathbf{d}}_j^i$, as the matrix of displacements.

Test Data. During test, we have a new image, on which we want to estimate the positions of the L landmarks, as shown in Fig. 1 (right). We randomly sample

K patches, where $\mathbf{c}_k \in \mathbb{R}^2$ and $\mathbf{f}_k \in \mathbb{R}^{d_f}$ are the center position and the visual feature of the kth patch that we sampled. We denote $\mathbf{F} = [\mathbf{f}_1, ..., \mathbf{f}_k] \in \mathbb{R}^{d_f \times K}$ as the matrix of features of all test patches.

Strategy. To estimate the position of the L landmarks on the test image, we first want to estimate $\{\mathbf{d}_k^l\}_{k=1...K, l=1...L}$, which is the displacement from each patch to each landmark. Then, $\{\mathbf{c}_k + \mathbf{d}_k^l\}_{k=1...K}$ will be the set of votes of the lth landmark's position from all the test patches, from which we can compute the response image by a voting scheme (details on response image will be presented in Section 2.4). Therefore, if we denote $\mathbf{D} \in \mathbb{R}^{2L \times K}$, whose element $\mathbf{D}_{ij} = \mathbf{d}_j^i$, as the matrix of displacements in the test image, our goal is to estimate $\mathbf{D}$.

2.2 Objective Function

First, we construct a compound displacement matrix:

$$\hat{\mathbf{D}} = \left[\, \tilde{\mathbf{D}}\ \mathbf{D} \,\right] = \begin{bmatrix} \tilde{\mathbf{d}}_1^1 & \cdots & \tilde{\mathbf{d}}_{\tilde{K}}^1 & \mathbf{d}_1^1 & \cdots & \mathbf{d}_K^1 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \tilde{\mathbf{d}}_1^L & \cdots & \tilde{\mathbf{d}}_{\tilde{K}}^L & \mathbf{d}_1^L & \cdots & \mathbf{d}_K^L \end{bmatrix} \in \mathbb{R}^{2L \times (\tilde{K}+K)} \tag{1}$$

The left part (the first $\tilde{K}$ columns) of $\hat{\mathbf{D}}$ contains the displacements in the training images, and the right part (the last K columns) is the displacements in the test image. Note that we have $\mathbf{D} = \hat{\mathbf{D}}\mathbf{Q}$ if we define $\mathbf{Q} = \begin{bmatrix} \mathbf{0}_{\tilde{K} \times K} \\ \mathbf{I}_K \end{bmatrix}$, where $\mathbf{I}_n$ is the $n \times n$ identity matrix, and $\mathbf{0}_{m \times n}$ is the $m \times n$ zero matrix.

Treating $\hat{\mathbf{D}}$ as a variable, our problem can be converted to finding the optimal $\hat{\mathbf{D}}$. Then, the optimal $\mathbf{D}$ can be computed by $\mathbf{D} = \hat{\mathbf{D}}\mathbf{Q}$. To this end, we design an objective with regard to $\hat{\mathbf{D}}$:

$$E(\hat{\mathbf{D}}) = E_g(\hat{\mathbf{D}}) + \alpha E_s(\hat{\mathbf{D}}) + \beta E_p(\hat{\mathbf{D}}) + \gamma E_l(\hat{\mathbf{D}}) \tag{2}$$

Ground-Truth Penalty $E_g(\hat{\mathbf{D}})$. The left part of $\hat{\mathbf{D}}$ should be close to the ground-truth displacements in the training data, which is encoded in $\tilde{\mathbf{D}}$. Therefore, we want to minimize the Ground-truth Penalty:

$$E_g(\hat{\mathbf{D}}) = \frac{1}{2L\tilde{K}} \left\| \hat{\mathbf{D}} \begin{bmatrix} \mathbf{I}_{\tilde{K}} \\ \mathbf{0}_{K \times \tilde{K}} \end{bmatrix} - \tilde{\mathbf{D}} \right\|_F^2 = \frac{1}{2L\tilde{K}} \left\| \hat{\mathbf{D}}\mathbf{P} - \tilde{\mathbf{D}} \right\|_F^2 \tag{3}$$

where $\|.\|_F$ is the Frobenius norm, and $\mathbf{P} = \begin{bmatrix} \mathbf{I}_{\tilde{K}} \\ \mathbf{0}_{K \times \tilde{K}} \end{bmatrix}$.

Smooth Mapping Penalty $E_s(\hat{\mathbf{D}})$. First, we construct a compound feature matrix $\hat{\mathbf{F}} = [\,\tilde{\mathbf{f}}\ \mathbf{f}\,] \in \mathbb{R}^{d_f \times (\tilde{K}+K)}$. Now, each column of $\hat{\mathbf{F}}$ is the feature of a patch, and each column of $\hat{\mathbf{D}}$ is the displacement vector (to all landmarks) of a patch. Then, $\|\mathrm{col}_i(\hat{\mathbf{F}}) - \mathrm{col}_j(\hat{\mathbf{F}})\|_{L2}$ is the $L2$ feature distance of a pair of patches

(i,j), where $\text{col}_i(\hat{\mathbf{F}})$ denotes the ith column of $\hat{\mathbf{F}}$. From all pairwise distances, we construct a binary affinity matrix $\mathbf{S} \in \{0,1\}^{(\tilde{K}+K)(\tilde{K}+K)}$, where $\mathbf{S}_{ij} = 1$ if and only if the ith and the jth patches are mutually $k(k = 20)$ nearest neighbors in the feature space. Note that the edges in the affinity matrix might link two training pathes, two test patches, or a training patch and a test patch.

The mapping from the feature space to the displacement space should be smooth. That is, for every pair of patches (i,j), if they are similar in the feature space, their displacements to landmarks should also be similar. We define the Smooth Mapping Penalty $E_s(\hat{\mathbf{D}})$ as the violation from this assumption:

$$E_s(\hat{\mathbf{D}}) = \frac{1}{2L\sum_{i\neq j}} \sum_{i\neq j} \mathbf{S}_{ij} \left\| \text{col}_i(\hat{\mathbf{D}}) - \text{col}_j(\hat{\mathbf{D}}) \right\|_{L2}^2 \tag{4}$$

For each pair of patches, E_s introduces a high penalty if the two patches are similar in the feature space (i.e. $\mathbf{S}_{ij} = 1$) but their displacements are very different (i.e. $\left\| \text{col}_i(\hat{\mathbf{D}}) - \text{col}_j(\hat{\mathbf{D}}) \right\|_{L2}$ is large). If we construct $\mathbf{M}$ as the (trace normalized) laplacian matrix [11] of $\mathbf{S}$, E_s can be compactly written as:

$$E_s(\hat{\mathbf{D}}) = \frac{1}{L} \text{Tr}\left(\hat{\mathbf{D}} \mathbf{M} \hat{\mathbf{D}}^\top \right) \tag{5}$$

Patch Offset Penalty $E_p(\hat{\mathbf{D}})$. Each column of $\mathbf{D}$ is the displacements from a single patch in the test image to all the landmarks. If we take the subtraction of two columns $\text{col}_{i-j}(\mathbf{D}) = \text{col}_i(\mathbf{D}) - \text{col}_j(\mathbf{D})$, it can be written as: $\text{col}_{i-j}\mathbf{D} = \mathbf{D}(\mathbf{e}_i^K - \mathbf{e}_j^K) = \begin{bmatrix} \mathbf{d}_i^1 - \mathbf{d}_j^1 \\ \cdots \\ \mathbf{d}_i^L - \mathbf{d}_j^L \end{bmatrix}$, where $\mathbf{e}_i^K$ is a K dimensional column vector whose ith element is 1 and all other elements are 0s. From Fig. 1 (right), we can see that $\mathbf{d}_i^1 - \mathbf{d}_j^1 = ... = \mathbf{d}_i^L - \mathbf{d}_j^L = \mathbf{c}_j - \mathbf{c}_i$, because $(\mathbf{d}_i^1, \mathbf{d}_j^1), ..., (\mathbf{d}_i^L, \mathbf{d}_j^L)$ form triangles with the same edge $\mathbf{c}_j - \mathbf{c}_i$. Therefore, we impose a penalty $E_p^{i-j}(\mathbf{D}) = \left\| \mathbf{D}(\mathbf{e}_i^K - \mathbf{e}_j^K) - \bar{\mathbf{c}}_{j-i} \right\|_F^2$, where $\bar{\mathbf{c}}_{j-i}$ is the L times vertical replicate of $\mathbf{c}_j - \mathbf{c}_i$. We can include a penalty for each pair (i,j) of columns. For efficiency reasons, we eliminate redundancies and use $K-1$ pairs to define the Patch Offset Penalty:

$$E_p(\hat{\mathbf{D}}) = \frac{1}{2LK} \sum_{i=1}^{L-1} E_p^{i-(i+1)}(\mathbf{D}) = \frac{1}{2LK} \left\| \hat{\mathbf{D}}\mathbf{Q}\mathbf{U} - \bar{\mathbf{C}} \right\|_F^2 \tag{6}$$

where $\mathbf{U} = \left[\mathbf{e}_1^K - \mathbf{e}_2^K, ..., \mathbf{e}_{K-1}^K - \mathbf{e}_K^K \right]$ and $\bar{\mathbf{C}} = \left[\bar{\mathbf{c}}_{2-1} \ ... \ \bar{\mathbf{c}}_{K-(K-1)} \right]$.

Landmark Offset Penalty $E_l(\hat{\mathbf{D}})$. Now we investigate the subtraction of rows in $\mathbf{D}$. For each pair of rows (i,j), we can write $\text{row}_{i-j}(\mathbf{D}) = \text{row}_i(\mathbf{D}) - \text{row}_j(\mathbf{D}) = (\mathbf{e}_i^{2L} - \mathbf{e}_j^{2L})^\top \mathbf{D}$. Note that $\mathbf{D}$ has $2L$ rows with interleaving coordinates in two axis as each $\mathbf{d}_j^i$ is two dimensional. For any even number i between 1 and $2L$, we have $\begin{bmatrix} \text{row}_{(i-1)-(i+1)}(\mathbf{D}) \\ \text{row}_{(i)-(i+2)}(\mathbf{D}) \end{bmatrix} = [\mathbf{d}_1^{i/2} - \mathbf{d}_1^{i/2+1}, ..., \mathbf{d}_K^{i/2} - \mathbf{d}_K^{i/2+1}]$. Note that,

$\left(\mathbf{d}_1^{i/2}, \mathbf{d}_1^{i/2+1}\right), ..., \left(\mathbf{d}_K^{i/2}, \mathbf{d}_K^{i/2+1}\right)$ all form triangles with edge $\mathbf{x}_{i/2+1} - \mathbf{x}_{i/2}$ (Fig. 1 (right)). Therefore, although the value $\mathbf{x}_{i/2+1} - \mathbf{x}_{i/2}$ is not known, we however know that elements of $\text{row}_{i-(i+2)}(\mathbf{D})$ should be identical, i.e. with zero variance:

$$\text{Var}\left(\text{row}_{i-(i+2)}(\mathbf{D})\right) = \left(\mathbf{e}_i^{2L} - \mathbf{e}_{i+2}^{2L}\right)^\top \mathbf{D}\mathbf{H}_k\mathbf{H}_k^\top\mathbf{D}^\top\left(\mathbf{e}_i^{2L} - \mathbf{e}_{i+2}^{2L}\right) = 0 \quad (7)$$

where $\mathbf{H}_k$ is the $K \times K$ centering matrix $\mathbf{H}_k = \mathbf{I}_k - \frac{1}{K}\mathbf{1}_{K\times K}$, where $\mathbf{1}_{K\times K}$ is the $K \times K$ matrix whose elements are all 1s. Eq. (7) holds for every pair, and since variance is always non-negative, this is equivalent to say that the summation of all variances is also zero, and we define Landmark Offset Penalty:

$$E_l(\hat{\mathbf{D}}) = \frac{1}{2LK}\sum_{i=1}^{2L-2}\text{Var}\left(\text{row}_{i-(i+2)}(\mathbf{D})\right) = \frac{1}{2LK}\text{Tr}\left(\mathbf{V}\hat{\mathbf{D}}\mathbf{Q}\mathbf{H}_k\mathbf{H}_k^\top\mathbf{Q}^\top\hat{\mathbf{D}}^\top\mathbf{V}^\top\right) \quad (8)$$

where $\mathbf{V} = [\mathbf{e}_1^{2L} - \mathbf{e}_3^{2L}, \mathbf{e}_2^{2L} - \mathbf{e}_4^{2L}, ..., \mathbf{e}_{2L-2}^{2L} - \mathbf{e}_{2L}^{2L}]^\top$.

2.3 Optimization

Substituting Eqs. (3),(5),(6) and (8) into Eq. (2), we get the final objective function. We can prove that Eq. (2) is convex, and the derivative is:

$$\partial E(\hat{\mathbf{D}}) \Big/ \partial\hat{\mathbf{D}} = \hat{\mathbf{D}}\mathcal{A} + \mathcal{B}\hat{\mathbf{D}}\mathcal{C} + \mathcal{G} \quad (9)$$

where $\mathcal{A} = \frac{1}{L\tilde{K}}\mathbf{P}\mathbf{P}^\top + \frac{2\alpha}{L}\mathbf{M} + \frac{\beta}{LK}\mathbf{Q}\mathbf{U}\mathbf{U}^\top\mathbf{Q}^\top$, $\mathcal{B} = \frac{\gamma}{LK}\mathbf{V}^\top\mathbf{V}$, $\mathcal{C} = \mathbf{Q}\mathbf{H}_k\mathbf{H}_k^\top\mathbf{Q}^\top$, $\mathcal{G} = -\frac{\tilde{\mathbf{D}}\mathbf{P}^\top}{L\tilde{K}} - \frac{\beta\bar{\mathbf{C}}\mathbf{U}^\top\mathbf{Q}^\top}{LK}$. We use gradient descend for the global optimum $\hat{\mathbf{D}}$.

Note that in our optimization process we jointly optimize the training data $\tilde{\mathbf{D}}$ and the test data $\mathbf{D}$. We could also take another way where the we optimize only $\mathbf{D}$. However, the advantage of our approach is that we can exploit the complex and relations between each training and test patch which is reflected in the matrix $\mathbf{S}$ in the term E_s.

2.4 Constructing Response Image

After we find the optimum $\hat{\mathbf{D}}$, we have $\mathbf{D} = \hat{\mathbf{D}}\mathbf{Q}$, and $\{\mathbf{c}_k + \mathbf{d}_k^l\}_{k=1...K}$ will be the set of votes for the position of the lth landmark. From these votes we perform kernel density estimation, which gives us the probability of landmark at each pixel location of the image (this is called the *response image* of the landmark). The response images of all landmarks will be used later for shape segmentation.

3 Shape Segmentation via Statistical Shape Model

Our shape segmentation algorithm works by combining the landmark detection with a statistical shape model. First, we detect a small set of global landmarks by exhaustively search in the image in different scales and rotations.

According to the position of these global landmarks, the image is scaled, translated and rotated to compensate for the global transformation. Then, the landmarks defining the shape are detected. During this step, we divide the whole shape into multiple subshapes consisting of several nearby successive landmarks. The method introduced in Section 2 is performed on each subshape. Then, shape segmentation is performed using response images of landmarks along with a statistical shape model. Instead of using the classical Active Shape Model [13] via PCA, we adopt the shape model based on sparse representation technique introduced in [14].

4 Experiments

Data: We tested our method on three tasks: segmentation of complete femur, proximal femur and pelvis. The X-ray data come from our clinical partner:

- CompleteFemur: 80 training images, 109 test images.
- ProximalFemur: 100 training images, 188 test images.
- Pelvis: 100 training images, 163 test images.

Note that a considerable part of the images are post-operative x-ray radiographs after trauma or joint replacement surgery, which significantly increases the challenge due to large variation of femur/pelvis appearance and the presence of implants (as can be seen in Fig. 2). As a indication, we made a manual counting, which shows that 32% of the test images contain implants.

Implementation Details: Each shape is divided into subshapes of 4 successive landmarks (i.e. L=4 in Section 2). The landmark detection algorithm in Section 2 is performed on each subshape, and then the segmentation of the whole shape is derived as in Section 3. In Section 2, for the patch visual feature, we use multi-level HoG (Histogram of Oriented Gradient [12]) feature with block sizes 1×1 and 2×2. Each block is divided into 2×2 cells and for each cell an 18 dimensional HoG feature is extracted by histogramming the gradient direction of each pixel. Therefore, feature dimension $d_f = 360$. For each subshape, we sample $\tilde{K} = 2000$ training patches and $K = 2000$ test patches. For the objective function, we use $\alpha = 0.1, \beta = 0.1, \gamma = 0.1$.

4.1 Results

Fig. 2 shows the qualitative result for all the three segmentation tasks. We can see that despite the challenges such as significant variation of appearance, poor image contrast and implants, our method achieves good results. For quantitative evaluation, we manually annotate the contour on all the test images, and the segmentation error is then calculated by the average point-to-curve distance between the landmarks of the segmented shape and the annotated contour. The error is then converted from pixel unit to physical unit (mm) as the pixel spacing of our radiograph is known. Table 1 shows the result. Note that to calculate the success rate, the failure cases are manually identified.

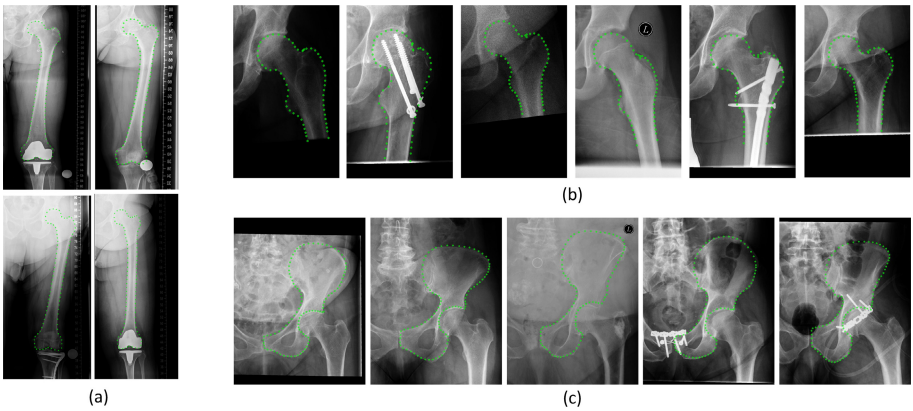

Fig. 2. Segmentation result on (a): CompleteFemur; (b): ProximalFemur; (c): Pelvis

Table 1. Quantitative evaluation. Numbers are in unit mm.

Anatomy	Success rate	Median	Minimum	Maximum	Mean	Standard deviation.
CompleteFemur	100%	1.2	0.7	3.4	1.3	0.4
ProximalFemur	98.4%	1.3	0.6	3.8	1.4	0.6
Pelvis	98.8%	1.9	1.0	4.3	2.2	0.8

Table 2. Comparison of our method with RF method. Numbers are in unit mm.

	CompleteFemur		ProximalFemur		Pelvis	
	Mean	Std	Mean	Std	Mean	Std
Our method	1.3	0.42	1.4	0.62	2.2	0.80
RF method	1.8	0.55	1.3	0.54	2.5	0.78
p-value	0.00005		0.31		0.05	

Our method processes one image of resolution 3000×3000 in about 2 minutes, with an unoptimized Matlab implementation on a computer with 3.0GHz CPU.

We compare our method with the random forest (RF) regression method. We use the same parameter for both methods when applicable (e.g. the same patch feature, the same number of training/test patches). For RF, we use 5 trees per forest[1]. The results are shown in Table 2, which shows the average error and standard deviation (in mm) of the two methods as well as the p-value. From the table we see that our method have comparable performance in proximal femur segmentation, and outperforms RF in complete femur and pelvis segmentation.

5 Conclusions

We have proposed a new method for fully-automatic landmark detection and shape segmentation in X-ray images. Our method works by jointly estimating

[1] We also tried more than 5 trees and no notable improvement was found.

the displacements from test patches to landmarks by considering both training data and geometric constraints. Experiments show that our method improves the landmark detection accuracy, and, combined with statistical shape model, can accurately segment shapes on the challenging datasets. In the future, we would like to integrate the shape model into our joint regression framework.

Acknowledgement. This work is supported by SNSF Project 51NF40-144610.

References

1. Chen, Y., Ee, X., Leow, W.-K., Howe, T.S.: Automatic extraction of femur contours from hip X-ray images. In: Liu, Y., Jiang, T.-Z., Zhang, C. (eds.) CVBIA 2005. LNCS, vol. 3765, pp. 200–209. Springer, Heidelberg (2005)
2. Lindner, C., Thiagarajah, S., Wilkinson, J.M., Wallis, G.A., Cootes, T.F.: Accurate fully automatic femur segmentation in pelvic radiographs using regression voting. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI 2012, Part III. LNCS, vol. 7512, pp. 353–360. Springer, Heidelberg (2012)
3. Gottschling, H., Roth, M., Schweikard, A., Burgkart, R.: Intraoperative, fluoroscopy-based planning for complex osteotomies of the proximal femur. Int. J. Med. Robot. 1(3), 67–73 (2005)
4. Baka, N., Kaptein, B.L., Bruijne, M., van Walsum, T., Giphart, J.E., Niessen, W.J., Lelieveldt, B.P.: 2D-3D shape reconstruction of the distal femur from stereo X-ray imaging using statistical shape model. Med. Image Anal. 15(6), 840–850 (2001)
5. Dong, X., Zheng, G.: Automatic extraction of proximal femur contours from calibrated X-ray images using 3D statistical models: an in vitro study. Int. J. Med. Robot. 5(2), 213–222 (2009)
6. Cristinacce, D., Cootes, T.: Automatic feature localization with constrained local models. Pattern Recognition 41(19), 3054–3067 (2008)
7. Zhou, S.K., Comaniciu, D.: Shape regression machine. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 13–25. Springer, Heidelberg (2007)
8. Zheng, Y., Barbu, A., Georgescu, B., Scheuering, M., Comaniciu, D.: Four-chamber heart modeling and automatic segmentation of 3-D cardiac CT volumes using marginal space learning and steerable features. IEEE T. Med. Imaging 27(11), 1668–1681 (2008)
9. Pauly, O., Glocker, B., Criminisi, A., Mateus, D., Möller, A.M., Nekolla, S., Navab, N.: Fast multiple organ detection and localization in whole-body MR Dixon sequences. In: Fichtinger, G., Martel, A., Peters, T. (eds.) MICCAI 2011, Part III. LNCS, vol. 6893, pp. 239–247. Springer, Heidelberg (2011)
10. Criminisi, A., Shotton, J., Robertson, D., Konukoglu, E.: Regression forests for efficient anatomy detection and localization in CT studies. In: Menze, B., Langs, G., Tu, Z., Criminisi, A. (eds.) MICCAI 2010. LNCS, vol. 6533, pp. 106–117. Springer, Heidelberg (2011)
11. Kokiopoulou, E., Chen, J., Saad, Y.: Trace optimization and eigenproblems in dimension reduction methods. Numerical Linear Algebra with Applications 18(3), 565–602 (2011)
12. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: CVPR (2005)
13. Cootes, T.F., Taylor, C.J.: Active shape models-'smart snakes'. In: BMVC (1992)
14. Zhang, S., Zhan, Y., Dewan, M., Huang, J., Metaxas, D.N., Zhou, X.S.: Sparse shape composition: a new framework for shape prior modeling. In: CVPR (2011)