

FULLY AUTOMATIC SEGMENTATION OF WRIST BONES FOR ARTHRITIS PATIENTS

Martin Koch^{1*}, Alexander G. Schwing^{2*}, Dorin Comaniciu³, Marc Pollefeys²

¹ Pattern Recognition Lab, Friedrich-Alexander-University Erlangen-Nuremberg, Erlangen, Germany

² Institute for Visual Computing, ETH Zürich, Zürich, Switzerland

³ Siemens Corporate Research, Princeton, NJ, USA

ABSTRACT

A reliable method to evaluate and follow the course of arthritis is given by examination of the carpal bones within the wrist joint. Humans typically have eight such small angular bones arranged in two rows. The small size as well as the number make manual segmentation for an analysis of the disease progression a tedious process. Further, fully automatic approaches are still not very reliable. To support medical treatment we present a fully automatic machine learning approach which (i) finds a bounding box around every bone and (ii) outlines the contour and computes a 3-D model of every carpal. The proposed approach has been successfully evaluated on 110 clinical wrist data sets of arthritis patients. The data consists of 59 T1 and 51 T2 weighted MRI images. With the point-to-mesh error deviating from ground truth an average of 0.48 ± 0.45 mm / 0.59 ± 0.49 mm on T1 / T2 modality, accurate segmentation results have been achieved.

Index Terms— arthritis, segmentation, wrist bones, 3-D model

1. INTRODUCTION

Approximately 22% adults in the United States report having doctor-diagnosed arthritis and by 2030, an estimated 67 million Americans ages 18 years or older are projected to have doctor-diagnosed arthritis [1]. Besides significant limitations in vital activities, work related implications are reported [2]. But most importantly arthritis and other rheumatic conditions have been and remain the most common cause of disability in the United States [3].

It is well known that arthritis erodes the carpal bones. Thus one possibility to evaluate and follow the course of disease is given by examining the wrist. It is however a tedious process to manually segment the eight carpal bones within both wrists accurately. Considering the amount of patients reported in above quoted statistics, this is certainly an intractable task. Therefore we present a reasonably fast, yet accurate framework that enables automatic segmentation of the eight carpal bones for both hands.

*This work was partially conducted while Martin Koch and Alexander G. Schwing were with Siemens Corporate Research, Princeton, NJ, USA.




Fig. 1. (a), (c) show slices of the T1 and T2 volumes respectively. Our resulting 3-D model is illustrated in (b) and (d).

Taouli *et al.* [4] showed that magnetic resonance imaging (MRI) is a good technique to visualize bone erosion. They did not observe a significant difference between high-field 1.5-T(esla) MRI and low-field 0.2-T MRI. We aim at processing either T1- or T2-weighted sequences of the wrist in any orientation as exemplarily shown in Figure 1(a) and Figure 1(c) respectively. The final 3-D(imensional) model obtained with our approach described in Section 3 is illustrated in Figure 1(b) and Figure 1(d).

Before diving into the details we review related work in Section 2. We present and discuss our results on a heterogeneous 3-D data set consisting of 59 T1- and 51 T2-weighted volumes of patients from three different clinics using different scanners in Section 4 and conclude in Section 5.

2. RELATED WORK

Despite the previously described necessity, there is relatively few work on segmentation of carpal bones in 3-D MRI se-

quences.

Sebastian *et al.* describe an approach to segment carpal bones from computed tomography (CT) sequences using skeletally coupled deformable models [5]. They emphasize that the non-uniformity of bone tissue, the irregular shape of closely packed carpal bones and small inter-bone spaces compared to the resolution render the segmentation of carpal bones a challenging task. Additionally we need to cope with the lower Signal-to-Noise ratio of MRI compared to CT sequences. The authors use a curve evolution implementation of region growing from manually initialized seeds, where growth is modulated by a competition between neighboring regions.

Similarly Duryea *et al.* describe their semi-automated approach using CT data [6]. However, their tool was developed to assist in a study of rheumatoid arthritis progression that used CT as well as MRI. From a technological perspective a reader had to segment a baseline volume semi-automatically, *i.e.* supported by slice based edge tracking. A 3-D image registration algorithm was applied to follow up scans. They mention that MRI has great potential as an imaging modality for arthritis which causes inflammation of the soft tissue which in turn is visible with MRI. According to our evaluation presented in Section 4, MRI data can be used to segment the bones. Further the segmentation can be used to tremendously reduce the search space for possible inflammations.

Aifeng *et al.* present a method for fully automatic carpal bone segmentation and feature analysis in hand X-ray radiographs [7]. The purpose of their work is to automatically segment the carpal bones by anisotropic diffusion and Canny edge detection techniques. To achieve segmentation they utilize a four stage procedure which first finds the region of interest within the acquired image. Secondly they apply an anisotropic diffusion filter to remove artifacts. Their third step is Canny edge detection followed by morphological operations to clean up the result in the fourth stage.

All of the above mentioned achievements use traditional non-learning based image processing in one or another form. To the best of our knowledge no framework is yet described using recent advances in machine learning, *i.e.* a combination of both discriminative and generative methods, to segment the eight carpal bones as described in the following section.

3. OUR APPROACH

Suppose given a volumetric image either T1 or T2 consisting of slices exemplarily illustrated in Figure 1(a) and Figure 1(c). From a high level perspective our approach consists of two steps. We first find the bounding boxes of the different carpal bones using discriminative classifiers. We provide a brief overview in Section 3.1. Within each bounding box we find in a second step the maximum a-posteriori (MAP) labeling, *i.e.* a mapping of each voxel to the class labels $\mathcal{L} = \{0, 1\}$ for no-bone and bone respectively. We provide the details in Section 3.2. Due to several reasons we split the original multi-

label task of segmenting the eight carpal bones into several binary classification problems. First, not all classifiers can be easily extended for the multi-label case. Second, there are not yet any multi-label (9 in our case) MAP algorithms that are computationally efficient when being applied to a reasonably sized 3-D volume.

3.1. Finding the bounding boxes

To find the bounding boxes we apply marginal space learning (MSL) proposed by Zheng *et al.* [8]. Contrasting the aforementioned work we however use a Random Forest (RF) classifier [9]. The fact that we target multiple modalities is the main reason. An RF classifier is particularly well suited due to its structure. It is formed out of multiple decision trees which in turn are composed by nodes performing binary splits depending on the features of a sample. The immense differences between T1 and T2 data are reflected in entirely different feature values. Thus T1 and T2 can easily be told apart, which in turn makes a separate treatment of T1 and T2 data more complex than necessary. Due to the page limit and as the procedure of finding bounding boxes as described in detail in [8] is not the focus of this work, we omit the subtleties and turn our attention in the following section to segmentation of the carpal bones.

3.2. Segmenting the carpal bones

Given the bounding box around a carpal, computed with the method and the deviations referenced and described in the previous section, we now want to label the N interior voxels according to their probability of belonging to class 1, *i.e.* how likely the particular voxel adheres to a bone structure. We denote by $\mathbf{y} \in \mathcal{L}^N$ and $\mathbf{x} = [x_1, \dots, x_N]^T \in \mathbb{R}^N$ the final binary labeling and the intensity values respectively. To obtain the most likely labeling we maximize the joint probability

$$p(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}) = \exp(\langle \mathbf{s}(\mathbf{y}), \boldsymbol{\theta}(\mathbf{x}) \rangle - A(\boldsymbol{\theta}(\mathbf{x}))) \quad (1)$$

modeled as a member of the exponential family [10]. The vector $\mathbf{s}(\mathbf{y})$ denotes the sufficient statistics (*i.e.* that combination of the values of the labeling that is necessary to compute the distribution). The parameters of the distribution (*e.g.* mean and variance for a Gaussian) are summarized in the vector $\boldsymbol{\theta}(\mathbf{x})$. The sufficient statistics depend on the chosen labeling \mathbf{y} . The dependence of the parameters $\boldsymbol{\theta}$ on the intensities \mathbf{x} is highlighted below. The log-partition function normalizing the distribution is denoted by $A(\boldsymbol{\theta}(\mathbf{x}))$. We model the distribution as a pairwise, binary Markov random field, *i.e.* it factors according to the structure of an undirected graph $G(\mathcal{V}, \mathcal{E})$ with node set \mathcal{V} and edge set $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ and thus a maximum clique size of two. To show that we indeed compute the MAP labeling we note that

$$\operatorname{argmax}_{\mathbf{y}} p(\mathbf{y} | \mathbf{x}; \boldsymbol{\theta}) = \operatorname{argmax}_{\mathbf{y}} p(\mathbf{x} | \mathbf{y}; \boldsymbol{\theta}) p(\mathbf{y}; \boldsymbol{\theta}), \quad (2)$$

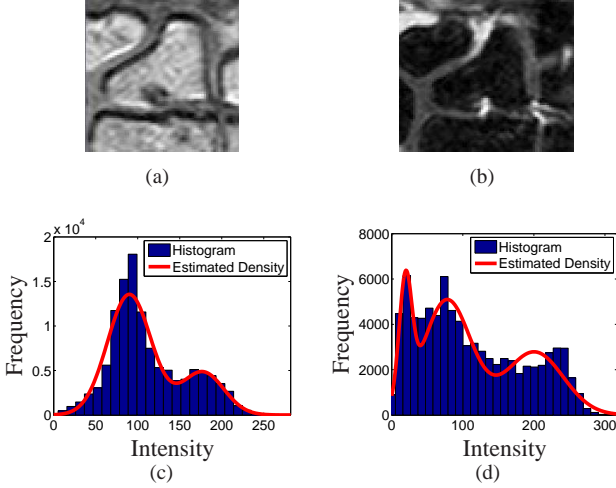


Fig. 2. Enlarged slices of detected bounding boxes for T1 and T2 volumes are shown in (a) and (b) respectively. The corresponding mixture distribution estimated from the underlying intensity distribution is shown in (c) and (d).

i.e. the posterior is proportional to the likelihood multiplied with the prior, and recall

$$\begin{aligned} p(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}) &= p(\mathbf{x} | \mathbf{y}; \boldsymbol{\theta}(\mathbf{x})) p(\mathbf{y}; \boldsymbol{\theta}(\mathbf{x})) \\ &= \exp(-E(\mathbf{y}, \boldsymbol{\theta}(\mathbf{x})) - A(\boldsymbol{\theta}(\mathbf{x}))) \end{aligned} \quad (3)$$

with the energy function

$$\begin{aligned} E(\mathbf{y}, \boldsymbol{\theta}(\mathbf{x})) &= -\langle \mathbf{s}(\mathbf{y}), \boldsymbol{\theta}(\mathbf{x}) \rangle \\ &= \sum_{ij \in \mathcal{E}} \theta_{ij}(y_i, y_j) + \sum_{i \in \mathcal{V}} \theta_i(y_i). \end{aligned} \quad (4)$$

Instead of maximizing Equation (3) using *e.g.* max-product (loopy) belief propagation [11], the energy E can as well be minimized using a linear program [12] or due to binary labels by using a graph-cut formulation [13]. Our choice for a graph-cut formulation is mainly guided by the simplicity of this approach. Note that we use the associated vector $\boldsymbol{\theta} = [\theta_i \ \forall i \in \mathcal{V}, \theta_{ij} \ \forall ij \in \mathcal{E}]^T$, where $[\cdot]^T$ denotes the transpose.

The pairwise potentials $\theta_{ij} \ \forall ij \in \mathcal{E}$ are modeled according to the squared loss $\frac{1}{2\beta} (x_i - x_j)^2$ with regularization parameter $\beta = 16$. When discussing our approach in Section 4 we evaluate two methods for choosing the singleton potentials $\theta_i \ \forall i \in \mathcal{V}$. As a simple idea a region around the center of the estimated box is assigned a high potential of belonging to the bone and the boundary of the estimated box obtains a high potential for the background label. The remaining nodes have equal potential for being bone or background. We name this first method “Simple” and compare it to an “Enhanced” approach, where the parameters $\theta_i \ \forall i \in \mathcal{V}$ are enhanced via a generative Gaussian mixture on the voxel intensities within a bounding box.

To do so we model the intensity distribution within an estimated bounding box as a superposition of Gaussians $\mathcal{N}(x | \mu, \sigma)$ with mean μ and variance σ in the form $p(x_i) = \sum_{k=1}^K \pi_k \mathcal{N}(x_i | \mu_k, \sigma_k)$. The responsibilities are denoted by π_k . For T1 data we choose $K = 2$. Due to bright speckles, also visible in Figure 2(b), $K = 3$ for T2 data. The parameters $\pi_k, \mu_k, \sigma_k \ \forall k$ are estimated using expectation maximization (EM) [14]. For the T1 and T2 volume illustrated in Figures 2(a) and 2(b), we show the relative frequency of the intensity x and a scaled variant of the estimated mixture model in Figures 2(c) and 2(d). The agreement between intensity histograms and estimated distribution is immediately apparent. Before choosing K , we however need to determine the data modality, *i.e.* T1 or T2. A simple yet fast thresholding on the average statistics of the volume under investigation was in our case more than sufficient for that task. To this end, the mean intensity of the center voxels for every box, representing the bone, is compared with the mean intensity of the voxels on the boundaries, mainly representing background. The final classification of the data set is based on majority voting among the eight boxes.

Instead of just assigning the center and the boundary of the estimated bounding box to a particular label as done in the method named “Simple,” we now enhance all potentials via a prior belief. For a voxel with intensity x_i in a T1 volume we use the log odds ratio

$$\log \frac{p(z_i = 0 | x_i)}{p(z_i = 1 | x_i)} = \log \frac{\pi_1 \mathcal{N}(x_i | \mu_1, \sigma_1)}{\pi_2 \mathcal{N}(x_i | \mu_2, \sigma_2)}. \quad (5)$$

Note, that z_i denotes the hidden label of voxel i from the complete data likelihood maximized via EM. For volumes of type T2, we apply the probability ratio

$$\begin{aligned} &\log \frac{p(z_i = 0 | x_i)}{p(z_i = 1 | x_i) + p(z_i = 2 | x_i)} = \\ &\log \frac{\pi_1 \mathcal{N}(x_i | \mu_1, \sigma_1)}{\pi_2 \mathcal{N}(x_i | \mu_2, \sigma_2) + \pi_3 \mathcal{N}(x_i | \mu_3, \sigma_3)}. \end{aligned} \quad (6)$$

Having obtained a binary labeling \mathbf{y} by maximizing Equation (3) using a graph cut approach, we apply a marching cubes algorithm [15] to obtain the meshes shown in Figures 1(b) and 1(d) for T1 and T2 data respectively. In the following section we will quantitatively assess the two methods investigated in this work.

4. EXPERIMENTS AND DISCUSSION

In the previous section we described two approaches for segmenting the carpals once we found a rough estimate for the bounding box. The “Simple” method just uses the information about the boundaries of the bounding box to assign potentials θ_i , whereas the “Enhanced” approach extracts global information from the entire interior to ensure a more accurate segmentation.

Table 1. Segmentation results in terms of area under the receiver operator characteristic (AUC) and Point-to-Mesh Error (PME) divided in (a) and (b) according to the underlying data. The results represent the average per carpal bone, labeled b_1 to b_8 (trapezium, trapezoid, capitate, hamate, scaphoid, lunate, triquetral and pisiform bones), for the 59 T1 and 51 T2 images respectively.

(a) T1 data sets.

| | Method | b_1 | b_2 | b_3 | b_4 | b_5 | b_6 | b_7 | b_8 |
|------------------------------|------------|-------|-------|-------|-------|-------|-------|-------|-------|
| AUC mean in % | “Simple” | 84.22 | 85.29 | 81.63 | 83.46 | 85.49 | 83.99 | 79.91 | 81.25 |
| AUC Standard deviation in % | “Simple” | 12.81 | 14.07 | 14.31 | 15.01 | 13.79 | 13.19 | 15.23 | 15.60 |
| AUC mean in % | “Enhanced” | 89.17 | 90.29 | 86.03 | 88.74 | 90.42 | 87.98 | 84.59 | 86.81 |
| AUC Standard deviation in % | “Enhanced” | 8.25 | 8.49 | 10.43 | 10.44 | 9.83 | 11.17 | 13.24 | 11.41 |
| PME mean in mm | “Simple” | 0.51 | 0.60 | 0.60 | 0.84 | 0.70 | 0.57 | 0.74 | 0.75 |
| PME Standard deviation in mm | “Simple” | 0.39 | 0.65 | 0.49 | 0.94 | 0.80 | 0.48 | 0.65 | 0.69 |
| PME mean in mm | “Enhanced” | 0.37 | 0.42 | 0.48 | 0.54 | 0.49 | 0.47 | 0.56 | 0.51 |
| PME Standard deviation in mm | “Enhanced” | 0.26 | 0.41 | 0.39 | 0.51 | 0.54 | 0.37 | 0.55 | 0.48 |

(b) T2 data sets.

| | Method | b_1 | b_2 | b_3 | b_4 | b_5 | b_6 | b_7 | b_8 |
|------------------------------|------------|-------|-------|-------|-------|-------|-------|-------|-------|
| AUC mean in % | “Simple” | 79.06 | 76.65 | 71.14 | 67.88 | 72.20 | 76.50 | 71.62 | 74.89 |
| AUC Standard deviation in % | “Simple” | 15.46 | 17.15 | 17.01 | 17.09 | 17.44 | 16.09 | 16.82 | 16.17 |
| AUC mean in % | “Enhanced” | 91.17 | 92.64 | 84.76 | 86.21 | 89.08 | 87.61 | 85.78 | 87.91 |
| AUC Standard deviation in % | “Enhanced” | 2.34 | 1.65 | 12.07 | 11.97 | 8.60 | 9.35 | 9.27 | 9.83 |
| PME mean in mm | “Simple” | 0.71 | 1.03 | 1.16 | 1.88 | 1.49 | 0.94 | 1.18 | 1.05 |
| PME Standard deviation in mm | “Simple” | 0.53 | 0.78 | 0.78 | 1.19 | 1.07 | 0.73 | 0.77 | 0.79 |
| PME mean in mm | “Enhanced” | 0.36 | 0.40 | 0.57 | 0.77 | 0.89 | 0.61 | 0.58 | 0.53 |
| PME Standard deviation in mm | “Enhanced” | 0.27 | 0.14 | 0.50 | 0.65 | 0.62 | 0.52 | 0.47 | 0.29 |

The data consists of 59 T1- and 51 T2-weighted medical MRI wrist volumes acquired from three different sites. The data sets are arbitrarily rotated, showing left and right wrists, with a resolution of $0.365 \times 0.365 \times 0.734 \frac{mm^3}{voxel}$. To facilitate a quantitative comparison we got all volumes professionally annotated manually. The evaluation was performed against the ground truth segmentation and conducted separately for each imaging modality. For an assessment we draw conclusions from two different measures. Those are the Point-To-Mesh Error (PME) and the area under the receiver operator characteristic (AUC). Whereas the first measure operates on the final meshes, *i.e.* after conversion of the labeling \mathbf{y} using marching cubes [15], the latter directly compares the labeling \mathbf{y} with the ground truth labeling obtain from the professionally annotated volumes. Note that the AUC measure is used in the statistics, machine learning and medical diagnostics community as shown in [16] and references therein. Point-To-Mesh error is also a commonly used error metric to quantify segmentation accuracy [8]. A brief overview is given below.

We compute the PME by finding the smallest distance between a vertex (point) of the meshed labeling \mathbf{y} , and the ground truth lattice representation (mesh). Hence, the name Point-To-Mesh. The averaged distances for all vertices result in the PME. The AUC, in contrast to PME, has no physical unit. It is a scalar measure for classifier performance within the interval $[0, 1]$ with 1 indicating perfect prediction.

We generally specify the percentage of perfect prediction, *i.e.* 100% corresponds to 1. Due to the page limit we refer the interested reader to [16] for further details.

The results obtained by applying the proposed method to T1- and T2-weighted MRI data sets is shown in Table 1(a) and 1(b) respectively. To abbreviate, we assigned labels b_x to the carpal bones according to the following scheme. The eight bones are arranged in two rows, the distal row contains from lateral to medial side the trapezium, trapezoid, capitate and hamate bones, labeled b_1 to b_4 . The proximal row contains in the same order the scaphoid, lunate, triquetral and pisiform bones, labeled b_5 to b_8 .

In terms of AUC the enhanced approach achieves an average segmentation accuracy of $88.0 \pm 10.6\%$ for T1 data sets and $88.1 \pm 9.2\%$ for T2 data sets. For T1 data sets the simple approach results in an average of $83.2 \pm 14.3\%$, which is about 5% less. However, when processing T2 data sets the segmentation results in terms of AUC drops to an average of $73.7 \pm 16.9\%$, which is about 14% below the results achieved with the enhanced version. For T1 images the PME is on average $0.48 \pm 0.45\text{mm}$ and $0.59 \pm 0.49\text{mm}$ for T2 images, both achieved with the “Enhanced” approach. The results obtained with the “Simple” approach are worse, *i.e.* $0.66 \pm 0.67\text{mm}$ and $1.18 \pm 0.92\text{mm}$ for T1 and T2 images respectively. For difficult cases, like those T2 volumes where intensity values of bone and surrounding tissue hardly differ, the “Enhanced”

method outperforms the “Simple” approach. Taking into account the data set resolution, we obtain reasonable errors in terms of PME, *i.e.* the same order of magnitude.

Duryea *et al.* specify the reader time for the segmentation approach described in [6] to the order of one hour per wrist. Our approach processes (detection and enhanced segmentation) a wrist volume of average size $422 \times 422 \times 51$ voxel in an average time of 6.17 minutes on a standard desktop (2.83 GHz Quadcore). Thus our fully automatic approach is almost 10 times faster than the semi-automatic method proposed in [6]. Unfortunately, the authors of [5–7] do not provide error measures for their segmentation result.

5. CONCLUSION

We obtain accurate results like the ones shown in Figures 1(b) and 1(d) on a diverse data set. Our method allows physicians to follow the course of disease by comparing segmentation results from different acquisitions and facilitates reduction of the search space when looking for inflammations.

To conclude, we presented a framework for *Fully Automatic Segmentation of Wrist Bones for Arthritis Patients*. To this end, we proposed to detect the bounding boxes of the carpals using MSL. For increased accuracy we enhance the segmentation using probability ratios obtained from a mixture of Gaussians. Additionally, we evaluated our approach on a challenging data set consisting of arbitrary rotated 110 volumes showing left and right wrists. The data was acquired from three different clinics and we achieve accurate segmentation deviating from ground truth an average of $0.48 \pm 0.45\text{mm} / 0.59 \pm 0.49\text{mm}$ on T1/T2 modality.

6. REFERENCES

- [1] J. Hootman, J. Bolen, C. Helmick, and G. Langmaid, “Prevalence of Doctor-Diagnosed Arthritis and Arthritis-Attributable Activity Limitation — US 2003–2005,” *Morbidity and Mortality Weekly Report*, vol. 55, no. 40, pp. 1089–1092, 2006.
- [2] K. A. Theis, L. Murphy, J. M. Hootman, C. G. Helmick, and E. Yelin, “Prevalence and correlates of arthritis-attributable work limitation in the US population among persons ages 18–64: 2002 National Health Interview Survey Data,” *Arthritis & Rheumatism*, vol. 57, no. 3, pp. 355–363, 2007.
- [3] J. M. Hootman, M. W. Brault, C. G. Helmick, K. A. Theis, and B. S. Armour, “Prevalence and Most Common Causes of Disability Among Adults – US, 2005,” *Morbidity and Mortality Weekly Report*, vol. 58, no. 16, pp. 421–426, 2009.
- [4] B. Taouli, S. Zaim, C. G. Peterfy, J. A. Lynch, A. Stork, A. Guermazi, B. Fan, K. H. Fye, and H. K. Genant, “Rheumatoid Arthritis of the Hand and Wrist: Comparison of Three Imaging Techniques,” *American Journal of Roentgenology*, vol. 182, pp. 937–943, 2004.
- [5] T. B. Sebastian, H. Tek, J. J. Crisco, and B. B. Kimia, “Segmentation of carpal bones from CT images using skeletally coupled deformable models,” *Medical Image Analysis*, vol. 7, no. 1, pp. 21–45, 2003.
- [6] J. Duryea, M. Magalnick, S. Alli, L. Yao, M. Wilson, and R. Goldbach-Mansky, “Semiautomated three-dimensional segmentation software to quantify carpal bone volume changes on wrist CT scans for arthritis assessment,” *Medical Physics*, vol. 36, 2008.
- [7] Z. Aifeng, A. Gertych, S. Kurowska-Pospiech, B. J. Liu, and H. K. Huang, “Carpal bone analysis in bone age assessment,” *Progress in Biomedical Optics and Imaging*, vol. 7, 2006.
- [8] Y. Zheng, A. Barbu, B. Georgescu, M. Scheuering, and D. Comaniciu, “Four-chamber heart modeling and automatic segmentation for 3D cardiac CT volumes using marginal space learning and steerable features,” *IEEE Transactions on Medical Imaging*, vol. 27, no. 11, pp. 1668–1681, 2008.
- [9] L. Breiman, “Random Forests,” *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [10] M. J. Wainwright and M. I. Jordan, *Graphical Models, Exponential Families, and Variational Inference*, Now Publishers Inc., Hanover, MA, USA, 2008.
- [11] K. P. Murphy, Y. Weiss, and M. I. Jordan, “Loopy Belief Propagation for Approximate Inference: An Empirical Study,” in *Proc. of Uncertainty in Artificial Intelligence*, 1999, pp. 467–475.
- [12] P. Kohli, A. Shekhovtsov, C. Rother, V. Kolmogorov, and P. Torr, “On partial optimality in multi-label MRFs,” in *Int’l Conf. on Machine Learning*, 2008, pp. 480–487.
- [13] V. Kolmogorov and R. Zabini, “What energy functions can be minimized via graph cuts,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 2, 2004.
- [14] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum Likelihood from incomplete data via the EM algorithm,” *Journal of the Royal Statistical Society*, vol. 39, no. 1, pp. 1–38, 1977.
- [15] W. E. Lorensen and H. E. Cline, “Marching Cubes: A high resolution 3D surface construction algorithm,” in *Proc. of Computer Graphics and Interactive Techniques*, 1987.
- [16] J. Fogarty, R. S. Baker, and S. E. Hudson, “Case studies in the use of ROC curve analysis for sensor-based estimates in human computer interaction,” in *Proc. of Graphics Interface*, 2005, pp. 129–136.