A CRF Approach to Fitting a Generalized Hand Skeleton Model

Radu Paul Mihail
Univ. of Kentucky
Dept. of Computer Science
r.p.mihail@uky.edu

Gustav Blomquist
Univ. of Kentucky
Dept. of Radiology
gbl224@email.uky.edu

Nathan Jacobs
Univ. of Kentucky
Dept. of Computer Science
jacobs@cs.uky.edu

Abstract

We present a new point distribution model capable of modeling joint subluxation (shifting) in rheumatoid arthritis (RA) patients and an approach to fitting this model to posteroanterial view hand radiographs. We formulate this shape fitting problem as inference in a conditional random field. This model combines potential functions that focus on specific anatomical structures and a learned shape prior. We evaluate our approach on two datasets: one containing relatively healthy hands and one containing hands of rheumatoid arthritis patients. We provide an empirical analysis of the relative value of different potential functions. We also show how to use the fitted hand skeleton to initialize a process for automatically estimating bone contours, which is a challenging, but important, problem in RA disease progression assessment.

1. Introduction

Imaging of the hand is routinely done to diagnose and asses the severity of diseases that alter the normal appearance of the musculoskeletal system. One such disease is rheumatoid arthritis (RA), a chronic systemic inflammatory autoimmune disease that primarily affects joints. The symptoms are pain, swelling and the loss of the joint function due to inflammatory processes. The underlying cause of RA is multifactorial [16, 8] including genetic susceptibilities, nutrition, lack of exercise and environmental factors. Joint inflammation caused by RA leads to over-vascularization, proliferation and synovial scar formation. The synovial proliferation is most marked at the margins of the joints, where the ticht space leads to bone erosions [3]. The inflammatory processes do not spare the ligaments, tendons and muscles, which leads to weakness, laxity and deformity.

We propose a novel method to automatically fit a skeleton model to a hand radiograph. Our approach builds on a previous model by Fernändez et al. [17], who proposed a point distribution model with landmarks located at joint centers. To support the types of deformation common in

RA, we relax this model by adding additional landmark points. Instead of a single point per joint, we have one located on each of the cortical articular surfaces of adjacent bones in the joint. This modification allows us to model subluxation (dislocation) of bones and supports our long-term goal of automatically measuring inter-joint spacing.

We provide a probabilistic formulation of our approach as a Conditional Random Field (CRF) and show to perform learning and inference with the model. We use a collection of potential functions, each tuned to a particular anatomical feature, such as upper and lower joint surfaces, or bone orientation. Each of these features makes unique contributions to our fitting process. We evaluate these features, and our CRF model, on real data from hands with and without deformation due to RA. Based on an analysis of the relative value of different potential functions, we find that the term that estimates the orientation of the joint makes significant contributions to rough alignment but other terms, such as the upper and lower joint potentials, make significant contributions by enabling more precise positioning of landmark points.

The main contributions of this work are: 1) introducing a new point-distribution model suitable for deformed hands, 2) a CRF framework for fitting this model to hand radiographs, 3) the definition of a set of potential functions that focus on specific anatomical structures in the hand, and 4) the evaluation of the accuracy of the method and the relative value of various potential functions on two datasets of hand radiographs.

2. Related Work

In this section, we describe previous work on vision-based methods for processing and analyzing hand radiographs.

Hand Radiograph Model Fitting Registering a parametric model to a hand radiograph is a key problem in this domain. Fernändez et al. [17] used a landmark-based wire model (which we extend in this paper) to develop a registration algorithm that outperforms thin-plate splines (TPS).

They initialize the wire model through a cascade of image processing routines based on bone axes. Van de Giessen et al. [20] developed a method to register CT scans of wrists by enforcing distances between bone surfaces to remain the same after registration. Bellerini et al. [2] use snakes optimized as initially proposed by Kass et al. [11] using genetic algorithms. They encode the parametric snake as polar coordinates centered at the origin which can be placed arbitrarily on the image. Xu et al. [21] introduced gradient vector flows as external forces, which eliminates the need to know a priori whether the snake will shrink or grow. Our work extends this line of research by generalizing the constrained shape model of Fernändez et al. [17]. This modification leads to the need for improved image feature extraction.

Hand Radiograph Pixel Labeling Numerous approaches have been proposed for pixel-level labeling of hand radiographs, we present several recent examples. Yuksel et al. [22] use a combination of feature classification and morphological operations to segment bone tissue from hand radiographs. Chai et al. [4] use the gray level co-occurrence matrix to segment texture and segment bone tissue from soft tissue. These approaches are similar to our feature extraction approach, but our features were developed to directly aid in model fitting, whereas these were developed for other purposes.

Hand Radiograph Analysis Hand radiographs are used frequently in medical diagnosis because the hands are where the pathology is most evident. We introduce several common medical uses for hand radiographs, each of which could benefit from the improved skeletal model fitting method we propose.

Langs et al. [13] presented a combination of active shape models and active contour models to segment bones and detect erosions on RA patient hand radiographs. Their approach relies on an initialization based on a local linear mapping net. We point out that hand radiographs of late stage RA patients are much more challenging due to severe subluxation, which result in overlapping bones, and joint space narrowing which lead to weak edge information, thus decreasing performance of purely edge-based methods.

In pediatric radiology, skeletal age is an important indicator of a healthy development process. Not only the bone locations and contours are of interest; bone density measurements aid in the diagnosis of skeletal development. One of the first complete descriptions of a system for hand radiograph analysis for skeletal age assessment is presented by Michael et al. [18]. Hue et al. [10] proposed an algorithm to segment hand bones on hand radiographs of children. Their approach relies on an oversegmentation using the watershed algorithm and region of interest extraction

and merging algorithms to segment soft tissue and background from noise. Sotoca et al. [19] proposed a semi-automatic approach where a user places the template at or near the center of a bone and the contour is approximated using active shape models (ASM).

Radiologists rely on expertise to assign a bone maturity score relative to age and gender. The most commonly used method to perform this evaluation is the atlas matching method by Greulich and Pyle (GP method) [2]. This is a time consuming process and correct assessment is highly dependent on the radiologist's experience and expertise, thus automated methods have been proposed. Giordano et al.[9] developed a method to predict bone age using a combination of filtering and Gradient Vector Flow Snakes with accuracy of 90%. Bayesian networks have been used by Mahmoodi et al. [14, 15]. Fuzzy systems have been used for skeletal age assessment by Aja-Fernändez et al. [1].

State-of-the-art bone segmentation and joint space width measurement approaches rely on landmark detection algorithms, usually based on a cascade of image processing techniques. This first step of landmark detection leads to most failures in existing algorithms. Our work fills that gap by accurately computing key anatomical points. roughly centered, currently done manually. Recent work by Davis et al. [7] provides encouraging results on automating this process.

3. Problem Definition

Given a roughly centered hand radiograph, our goal is to estimate landmarks point locations on the edge of cortical articular surfaces along the main axis of long bones. In this section, we formally define our shape model and identify key challenges in solving this problem.

3.1. Shape Model

We represent a shape, s, by a set of n landmark point locations $s = (x_1, x_2, x_3, ..., x_n, y_1, y_2, y_3, ..., y_n)^T$. The choice of landmarks depends on the object of interest and the application, but for hand radiographs they are usually chosen as joint centers and fingertips. A recent example is the work of Fernändez et al. [17] where a shape model is used as an initialization step to an image registration algorithm. We chose to generalize their representation by having two landmarks per joint, one on each side of the joint on the cortical articular surface of the bone, collinear with the bone's main axis, (i.e., the tips of long bones). Figure 1 shows a visualization of this model. This relaxation allows us to model the subluxation deformities that are common in moderate to late stage rheumatoid arthritis patients.

3.2. Key Challenges

Automated methods for radiograph analysis rely on consistent alignment and appearance, which rarely happens in

Figure 1. Our proposed shape model: each segment corresponds to a bone ($B_{1\ldots 19}$). Individual points are indexed as proximal $s_{\{1\ldots19\}p}$ and distal $s_{\{1\ldots19\}d}$.

practice. In this work, we focus on solving the alignment problem by fitting an initial hand model to the radiograph. We describe several important challenges in solving this fitting problem.

Despite attempts to control hand position using clinical protocols, hand radiographs of healthy patients show significant variations in pose. In addition to RA deformities, such as joint fusion, sclerosis and subluxation, other diseases including osteoarthritis may be present further increasing variability in pose. For example, RA damage and pain can prevent patients from flattening their hands on the imaging surface. A solution to the hand model fitting problem must be able to cope with significant changes in pose and joint deformities.

The appearance of bones also varies significantly from patient to patient. This is especially true in RA patients because the disease affects the density and shape of individual bones. Many previous approaches to the hand model fitting problem focus on specific anatomical features for alignment, but this leads to brittle solutions. Therefore, an approach that combines image information from multiple anatomical features is needed. Our work addresses both of these concerns in a consistent, and adaptable probabilistic formulation.

4. Approach

We propose a CRF-based model that combines a shape prior with appearance terms that identify various anatomical structures. The shape prior is based on the distance from

the subspace spanned by a Point Distribution Model (PDM) and the appearance is defined as a collection of likelihood terms that depend on local feature detectors. We use a local optimization strategy to jointly maximize feature responses at landmark locations by minimizing an energy function. Minimizing energy in this context is equivalent to maximizing the posterior distribution over the correct location of landmarks given an image.

Our proposed CRF has the following form:

$$P(s|I,\theta) = \frac{1}{Z} exp\{\sum_i\{\sum_j \Psi_j(s_{i_p}, s_{i_d}, \theta)\}+ \quad (1)$$
$$+ \sum_k \phi_k(s_i, \theta) + \zeta(s, \theta)\}$$

where i is an index over model segments, j and k index our pairwise and unary appearance terms, Z is the partition function, Ψ and ϕ are pairwise and unary appearance terms, ζ is a shape model prior and θ is a weight vector we use to balance the various terms.

4.1. Potential Functions

The data terms in our model are based on a collection of discriminative features that we combine into a set of potential functions.

Discriminative Features The success of the shape fitting process is heavily dependent on a set of features that are highly discriminative. Recently, Cootes et al. [5] showed how regression voting using random forests in the constrained local model (CLM) framework outperforms existing methods on shape fitting. Our approach extends Constrained Local Models (CLM) [6], by formulating the shape fitting problem in a general CRF framework.

We classify each pixel in an image independently using a randomized decision forest (RDF) classifier and use dense SIFT features as input. We train 4 RDF classifiers, one for joint centers (the midpoint between adjacent bones connected by joints), two for cortical articular surface points, one for proximal and one for distal, and a bone tissue classifier.

For a new image, we compute DSIFT descriptors and run them through our RDF classifiers. The output is a likelihood that represents class membership of each pixel. We note that this process implies independence in pixel memberships (e.g., we could potentially have a pixel be labeled as both bone and cortical surface). Let the classifier response for joints be $f_j(s_i)$ where s_i is a landmark point with a corresponding image location. Similarly, we define $f_{pc}(s_i)$ and $f_{dc}(s_i)$ for proximal and distal cortical articular surface pixels. Finally, let $f_b(s_i)$ be the classifier response for bone tissue. Examples of RDF classifiers outputs can be seen in Figure 3.

We then apply a thresholding operation on f_j and f_b to compute binary regions of high probability bf_j and bf_b. Using the thresholded responses, we apply distance transformations, and combine them with values inside the regions to compute $df_j(s_i)$ and $df_b(s_i)$. The distance transformation returns 1 on the region borders and 0 inside. The locations inside the regions are filled with $1 - f_j$ and $1 - f_b$ respectively. This approach increases alignment precision by providing extra information about the most probable location of anatomical interest points. Examples can be seen in Figure 4.

In the next section we show how we convert these low-level image features into potential functions in our CRF model.

Pairwise potentials   The first pairwise potential encodes the compatibility between a segment in our shape model and evidence of bone tissue:

$$\Psi_1(s_{i_p}, s_{i_d}) = \theta_1 \frac{1}{n} \sum_{n=1}^{t} df_b(p_{x_n}, p_{y_n}) \qquad (2)$$

where the summation is over points p sampled along a segment. This function takes as input the distance transformation df_b and is low when a segment is placed over a bone.

The bone evidence from the image can be further exploited by considering segment orientation extracted from the thresholded bone tissue classifier connected components. We define the following potential:

$$\Psi_2(s_{i_p}, s_{i_d}) = \theta_2 \left(tan^{-1}(s_{i_p} - s_{i_d}) - f_o(s_i)\right)^2 \qquad (3)$$

In the above equation, f_o is a function that returns an angle at an image location computed via a weighted averaging of angles of the connected components with respect to the horizontal image axis. Intuitively, if a segment s_i is placed perpendicular to the major axis of a connected component, the potential will be at its maximum. In Figure 4 we show f_o for an image.

We now define a pairwise potential that encodes a prior over adjacent cortical articular surfaces in a joint:

$$\Psi_3(s_{i_d}, s_{i+1_p}) = \theta_3(-log\,\mathcal{N}(0, \Sigma)) \qquad (4)$$

where $\mathcal{N}$ is a Gaussian with full rank covariance Σ computed from the training set. This term constrains joint spaces to be at reasonable distances in order to avoid local minima during optimization.

Unary potentials   We define three terms that encourage the landmark points to be near appropriate anatomical features of the joint. The motivation for our first unary potential is that all points on the model should be in areas of high


Figure 2. Left: in red, segments span the major axis of connected components in a binary image (bf_b) used for initialization. Right: in green, the top 3 models from the training set with the lowest ICP registration error.


Figure 3. Left: joint center pixel probabilities. Middle: color coded probabilities for distal (red channel) and proximal (blue channel) cortical surface pixels. Right: bone pixel probabilities.


Figure 4. Left and middle: Distance transformations $df_j(s_i)$ and $df_b(s_i)$ of thresholded classifier responses. Right: Orientation term f_o.

probability indicated by our joint feature, f_j. To further penalize points from being far from a joint, and to improve optimization performance, we use df_j, which is an augmented version of f_j. This unary potential is defined as follows:

$$\phi_3(s_i) = \theta_6\,df_j(s_i) \qquad (5)$$

The second two unary potential functions encourage landmark points to align to the joint contours. These potentials are defined as follows:

$$\phi_1(s_{i_p}) = \theta_4\,f_{pc}(s_{i_p}) \qquad (6)$$

and

$$\phi_2(s_{i_d}) = \theta_5\,f_{dc}(s_{i_d}). \qquad (7)$$

For this domain using classifiers, f_{dc} and f_{pc}, instead of a generic edge detector is critical because it allows the model to distinguish between the true bone contours used for diagnosis and analysis and apparent edges caused by the radiographic projection of other bone structures.

We find that in practice these terms complement each other. The first is very important for rough initial alignments, while the second two are critical for the precise alignments. See the evaluation section for details.

4.2. Shape Model Prior

The shape prior term $\zeta(s, \theta)$ is used to penalize unlikely shapes. Using Probabilistic Principal Component Analysis (PPCA), we seek to relate shape s to a k-dimensional vector x that is normally distributed with zero mean and covariance $I(k)$:

$$s^T = Wx^T + \bar{s} + \epsilon \qquad (8)$$

where W is the matrix of principal components and $\bar{s}$ is the average shape. ϵ is the model noise component, assumed to be normally distributed $\epsilon \sim \mathcal{N}(0, \sigma^2 I)$.

Under this model, s is normally distributed:

$$P(s) = \mathcal{N}(\bar{s}, WW^T + \sigma^2 I(k)) \qquad (9)$$

so our shape prior is a weighted negative log likelihood of $P(s)$:

$$\zeta(s, \theta) = \theta_7(-\log P(s)) \qquad (10)$$

W and σ^2 are estimated using an Expectation-Maximization algorithm from a training set of hand shapes extracted from hand radiographs.

4.3. Shape Inference

We now describe our strategy for estimating the optimal shape model for a given hand radiograph. We first compute a set of initial models, then use local optimization for each and select the best.

We use thresholded bone tissue classifier, bf_b, for initialization by computing the connected components statistics in the binary image. We model each component as an ellipse and compute its centroid, major axis, orientation and length. The line segments spanning the major axis of the connected components (see Figure 2) form the basis of our initialization scheme. We use a variation of the Iterative Closest Point algorithm to register samples from our training set to the segments extracted from the binary bone tissue classifier. The ICP registration error is then used to select the 3 best shapes that are used as starting points for our local optimization.

Our local shape objective function is the posterior probability of a shape, s, given image data, or equivalently the log likelihood of our CRF model (1), which is defined as follows:

$$E(\hat{s}, \theta, I) = \arg \min_s \sum_i \{ \sum_j \Psi_j(s_{i_p}, s_{i_d}) \} \qquad (11)$$
$$+ \sum_k \phi_k(s_i) + \zeta(s, \theta).$$

To minimize (11) we use coordinate descent with step sizes chosen by an independent local search in each dimension.

4.4. Estimating CRF Weights

We use a supervised learning approach to estimate the model parameters θ. The partition function Z is NP-hard to compute [12]. We overcome the difficulty of computing Z in learning the model parameters by using Pseudo-Likelihood learning, where a uniform prior over model parameters is assumed by setting $\tau = \infty$ in $P(\theta | \tau) = \mathcal{N}(\theta, 0, \tau^2 \mathbf{I})$ where I is the identity matrix. We find the optimal parameters θ by minimizing the difference between our estimated shapes and ground truth shapes over a set of training images:

$$\hat{\theta}^{ML} = \arg \min_\theta \sum_i \left\| \arg \min_s E(s, \theta, I_i) - s_i^{GT} \right\|_2^2. \quad (12)$$

The above minimization is non-convex since we allow s to vary during optimization. We use the simplex method with random restarts to compute model parameters, θ. In the following section we evaluate the model parameters and show results from inference.

5. Evaluation

Datasets We evaluate our method on two datasets with posterior-anterior view hand radiographs: the Digital Hand Atlas Database[1] and a set of 43 radiographs of RA patients from the University of Kentucky Department of Radiology. The second dataset is of a hands in a range of disease stages, from minimal to extreme deformation. For evaluation purposes, we manually determined landmark locations for each image in both datasets. Since radiograph contrast varies due to calibration parameters and noise, we truncate the upper 20% of the intensity histogram.

Quantitative Evaluation We used a set of 20 images from both datasets to train our discriminative classifiers and estimate our remaining model parameters. The remaining images were used for testing and validation.

We evaluate the model by computing the sum of absolute differences between ground truth shapes and results from inference. The model errors for both datasets can be seen in Figure 7 and 6. For the Hand Atlas set, the average per point

---
[1]http://www.ipilab.org/BAAweb/

Figure 5. Results on a representative subset of test images. Top row: healthy radiographs from the Hand Atlas Database. Bottom row: rheumatoid arthritis set.

error was 2.72 pixels, while for the RA dataset it was 2.85 pixels. A comparison to the state of the art is difficult due to our model landmark selection and RA deformity severity. We divide the test set into early (16 images), moderate (11 images) late stage radiographs (11 images), the average per point errors (measured in distance from ground truth, in pixels) are as follows: 2.30, 2.24, and 4.56.

To help understand the failure modes of our approach, we further investigate two radiographs with poor shape estimates. These correspond to images 23 and 27 in Figure 7. We find that by inspecting the optimal fit for both images, shown in Figure 8, that they are both from patients with late-stage RA and have severe deformities and subluxation. In such cases, assistance from a radiologist will be required.

Term Contributions To provide more insight into the model, we estimate the amount each potential function contributes to reducing errors in the RA dataset. We split the dataset into two groups, a training set of size 20 and a testing set of size 12. We use the training dataset to estimate the optimal potential function weights, $\hat{\theta}$, by minimizing (12), as described above, for the full model. Then, for each of the seven potential function, we solve for the optimal set of weights for the model without that potential function, leaving one out. This results in a set of eight different models. For each model, we infer the shape in each image in the testing dataset and sum the absolute pixel error with respect to the ground-truth shape to obtain an error measure. Table 1

Table 1. Percent increase of error for each term when omitted from the model.

Term	Ψ_1	Ψ_2	Ψ_3	ϕ_1	ϕ_2	ϕ_3	ζ
% Error increase	39.14	1.42	1.87	2.19	0.09	10.56	0.27

shows the ratio of the error of a model without the potential function to the error for the full model. Intuitively, an important term will result in a model with significantly higher error if it is removed. We find that there are two dominant terms: Ψ_1 and ϕ_3 that correspond to bone and joint evidence from the feature set. These terms are clearly the most important in gross alignment, however the other terms each make a contribution to reducing errors in the full model.

Example Applications: Initialization for Estimating Bone Contours The weakness of most state-of-the-art approaches for identifying bone contours is the initialization step. As an example application, we propose to fit an active shape model (ASM) to the cortical articular surfaces for each finger joint. This is challenging because ASM models must be initialized very close to the optimal location or they will fall into non-optimal local minima. We use our proposed approach to estimate a model skeleton and use the landmark points and bone segments to align the initial ASM model for each joint. We optimize the ASM shape parameters, using an off-the-shelf software library, and ob-

Figure 6. Hand Atlas Dataset model errors computed as sum of absolute differences from ground truth.


Figure 7. Rheumatoid Arthritis Dataset model errors computed as sum of absolute differences from ground truth.


Figure 8. Failures that correspond to high errors from Figure 6 and 7.

tain the results seen in Figure 9. This demonstrates that our estimated skeleton models are sufficiently accurate to pro-

vide initial conditions for ASM models of joint contours. In combination, such an approach could be used to automatically estimate the joint space width, an important metric for RA progression.

6. Conclusion

We introduced a new method for fitting wireframe-hand models to radiographs. A key innovation in our approach is in fitting a relaxed shape model, with four degrees of freedom at each joint, that is capable of representing the dramatic subluxations present in patients with rheumatoid arthritis (RA). Fitting this model effectively is more challenging than standard models, which only have two degrees of freedom at each joint. We show that our method, which combine low-level discriminative features in a conditional random field framework, is capable of fitting this relaxed model on healthy hands as well as those deformed by RA. We provide quantitative results that highlight which features are most important and show an application of our method to fitting bone contours, which is critical in assessing RA damage.

Acknowledgements

References

[1] S. Aja-Fernández, R. de Luis-Garcıa, M. A. Martın-Fernández, and C. Alberola-López. A computational tw3 classifier for skeletal maturity assessment. a computing with words approach. *Journal of Biomedical Informatics*, 37(2):99–107, 2004. 2

Figure 9. Joint contours (red) estimated by initializing an active shape model based on our initial hand skeleton (green).

[2] L. Ballerini and L. Bocchi. Multiple genetic snakes for bone segmentation. *Applications of Evolutionary Computing*, pages 346–356, 2003. 2

[3] M. Buch and P. Emery. The aetiology and pathogenesis of rheumatoid arthritis. *HOSPITAL PHARMACIST-LONDON-*, 9(1):5–10, 2002. 1

[4] H. Chai, L. Wee, T. Swee, and S. Hussain. Glcm based adaptive crossed reconstructed (acr) k-mean clustering hand bone segmentation. *Recent Researches in Communications, Automation, Signal Processing, Nanotechnology, Astronomy and Nuclear Physics*, pages 192–197, 2011. 2

[5] T. F. Cootes, M. C. Ionita, C. Lindner, and P. Sauer. Robust and accurate shape model fitting using random forest regression voting. In *Computer Vision–ECCV 2012*, pages 278–291. Springer, 2012. 3

[6] D. Cristinacce and T. F. Cootes. Feature detection and tracking with constrained local models. In *BMVC*, volume 17, pages 929–938, 2006. 3

[7] L. Davis, B.-J. Theobald, A. Toms, and A. Bagnall. On the extraction and classification of hand outlines. In *Intelligent Data Engineering and Automated Learning-IDEAL 2011*, pages 92–99. Springer, 2011. 2

[8] K. D. Deane, J. M. Norris, and V. M. Holers. Pre-clinical rheumatoid arthritis: Identification, evaluation and future directions for investigation. *Rheumatic diseases clinics of North America*, 36(2):213, 2010. 1

[9] D. Giordano, C. Spampinato, G. Scarciofalo, and R. Leonardi. An automatic system for skeletal bone age measurement by robust processing of carpal and epiphysial/metaphysial bones. *Instrumentation and Measurement, IEEE Transactions on*, 59(10):2539–2553, 2010. 2

[10] T. Hue, J. Kim, and M. Fahriddin. Hand bone radiograph image segmentation with roi merging. In *Proceedings of the 13th IASME/WSEAS international conference on Mathematical Methods and Computational Techniques in Electrical Engineering conference on Applied Computing*, pages 147–154. World Scientific and Engineering Academy and Society (WSEAS), 2011. 2

[11] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International journal of computer vision*, 1(4):321–331, 1988. 2

[12] S. Kumar and M. Hebert. Discriminative random fields: A discriminative framework for contextual interaction in classification. In *IEEE International Conference on Computer Vision (ICCV)*, 2003. 5

[13] G. Langs, P. Peloschek, and H. Bischof. Asm driven snakes in rheumatoid arthritis assessment. In *Image Analysis*, pages 454–461. Springer, 2003. 2

[14] S. Mahmoodi, B. Sharif, E. Chester, J. Owen, and R. Lee. Automated vision system for skeletal age assessment using knowledge based techniques. In *IEEE International Conference on Image Processing and Its Applications*, 1997. 2

[15] S. Mahmoodi, B. S. Sharif, E. G. Chester, J. P. Owen, and R. Lee. Skeletal growth estimation using radiographic image processing and analysis. *Information Technology in Biomedicine, IEEE Transactions on*, 4(4):292–297, 2000. 2

[16] M. Maravic, C. Berge, J. Daures, M. Boissier, et al. Practices for managing a flare of long-standing rheumatoid arthritis: survey among french rheumatologists. *Clin Exp Rheumatol*, 23:36–42, 2005. 1

[17] M. Martín-Fernández, R. Cárdenes, E. Muñoz-Moreno, R. de Luis-García, M. Martín-Fernández, and C. Alberola-López. Automatic articulated registration of hand radiographs. *Image and Vision Computing*, 27(8):1207–1222, 2009. 1, 2

[18] D. Michael and A. Nelson. Handx: A model-based system for automatic segmentation of bones from digital hand radiographs. *Medical Imaging, IEEE Transactions on*, 8(1):64–69, 1989. 2

[19] J. Sotoca, J. Iñesta, and M. Belmonte. Hand bone segmentation in radioabsorptiometry images for computerised bone mass assessment. *Computerized Medical Imaging and Graphics*, 27(6):459–467, 2003. 2

[20] M. Van De Giessen, G. J. Streekstra, S. D. Strackee, M. Maas, K. A. Grimbergen, L. J. Van Vliet, and F. M. Vos. Constrained registration of the wrist joint. *Medical Imaging, IEEE Transactions on*, 28(12):1861–1869, 2009. 2

[21] C. Xu and J. L. Prince. Generalized gradient vector flow external forces for active contours. *Signal processing*, 71(2):131–139, 1998. 2

[22] A. Yuksel and T. Olmez. Automatic segmentation of bone tissue in x-ray hand images. In *Adaptive and Natural Computing Algorithms*, pages 590–599. Springer, 2009. 2