

Automatic Hand Skeletal Shape Estimation from Radiographs

Radu Paul Mihail

rpmihail@valdosta.edu

Department of Computer Science
Valdosta State University, Valdosta GA

Nathan Jacobs

nathan.jacobs@uky.edu

Department of Computer Science
University of Kentucky, Lexington KY

Abstract—Rheumatoid arthritis (RA) is an autoimmune disease whose common manifestation involves the slow destruction of joint tissue, damage that is visible in a radiograph. Over time, this damage causes pain and loss of functioning which depends, to some extent, on the spatial deformation induced by the joint damage. Building an accurate model of the current deformation and predicting potential future deformations is an important component of treatment planning. Unfortunately, this is currently a time consuming and labor intensive manual process.

To address this problem, we propose a fully automated approach for fitting a shape model to the long bones of the hand from a single radiograph. Critically, our shape model allows sufficient flexibility to be useful for patients in various stages of RA. Our approach uses a deep convolutional neural network to extract low-level features and a conditional random field (CRF) to support shape inference. Our approach is significantly more accurate than previous work that used hand-engineered features. We demonstrate this on two large datasets of hand radiographs and highlight the importance of the low-level features, the relative contribution of different potential functions in the CRF, and the accuracy of the final shape estimates. Our approach is nearly as accurate as a trained radiologist and, because it only requires a few seconds per radiograph, can be applied to large datasets to enable better modeling of disease progression. We will release our code and trained models upon acceptance of this paper.

Index Terms—rheumatoid arthritis, radiograph, conditional random field, convolutional neural network

I. INTRODUCTION

Rheumatoid Arthritis (RA) is an autoimmune disease, with no known cure, that primarily affects synovial joints, especially those of the hands. The disease typically starts with inflammation and swelling, followed by the mutilation of joints through healthy tissue loss and scar tissue formation. The causes of the disease are multiple; genetic susceptibilities, lack of exercise, and environmental factors are considered to play a role in the onset and progression of the disease.

Imaging of the hand is a routine procedure used by radiologists to assess the extent of the damage and to estimate the stage of disease progression. Disease staging can be a complex process, a function of key anatomical features and changes from a healthy baseline. Some of those features include the inter-joint spaces that tend to get smaller as the disease progresses. Recently, Pfeil et al. [20] have investigated using automated estimation of joint spaces as a predictor of RA disease progression with success. This motivates further work in this area, since different combination of medications may be more effective at earlier stages in the disease progression [22].




Fig. 1. Our radiograph shape fitting pipeline. Input image (left) is directly processed by a set of deep convolutional neural networks to produce feature maps (middle) that correspond to key anatomical features. The final shape (right) is inferred using a conditional random field (CRF).

Automated joint space estimation is challenging for several reasons. First, the process for capturing radiographs, which are typically in the posteroanterior (PA) view, results in significant variability in hand placement and joint configuration. Such variability is acceptable for radiologists, but means that automated systems must cope with these differences. Second, appearance variations in radiographs appear due to inconsistent image acquisition parameters from calibration, digitization artifacts from film radiographs (scanner calibration), and anatomical differences. Appearance variations due to anatomy can be classified into two categories: differences from individual to individual and morphological differences due to disease (e.g., rheumatoid arthritis, osteoporosis), including inflammation and surgery.

We propose a fully automatic approach to estimating the configuration of the long bones of the hand from a radiograph. We adopt a previously introduced shape model [19] and combine bottom-up supervised feature extraction with a top-down shape inference process in the form of a conditional random field (CRF). We propose several novel improvements to the previous work on this problem, including improved low-level features and better initialization for the CRF optimization. These changes substantially increase the accuracy and robustness without increasing the run time during inference.

We make the following key contributions: 1) an improved low-level feature extraction process that uses deep fully convolutional neural networks (CNNs), 2) an improved CRF initialization strategy, and 3) an evaluation on a large datasets

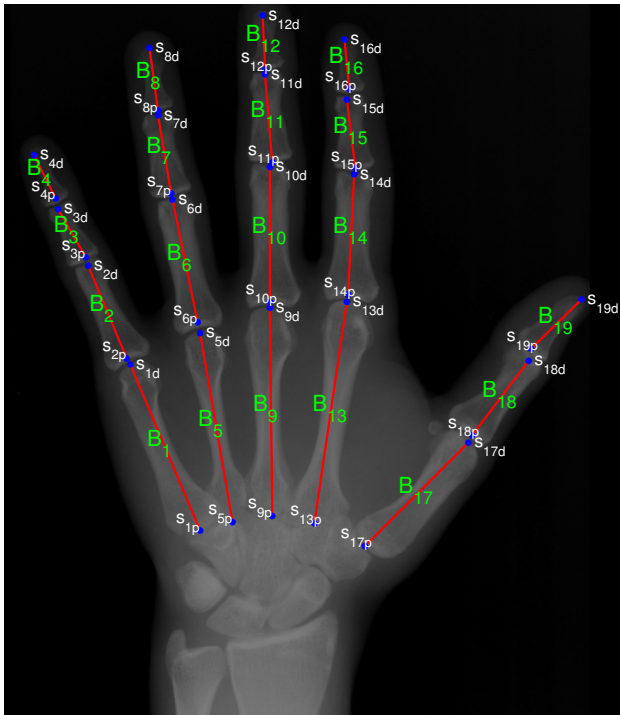


Fig. 2. Our shape model: each segment corresponds to a bone ($B_{1\dots 19}$). Individual points are indexed as proximal $s_{\{1\dots 19\}p}$ and distal $s_{\{1\dots 19\}d}$.

that demonstrates the benefit of our improvements.

II. RELATED WORK

Hand radiograph analysis has been investigated in the past and continues to receive attention from the computer vision and medical imaging research communities. We group related work into three categories: parametric shape model fitting to radiographs (most related to this work), hand radiograph pixel-level labeling algorithms and general radiograph analysis, a combination of parametric and non-parametric models and pixel-level operations. Our problem formulation is similar to face landmark localization and human pose estimation [4], [11], [24]. Recently, significant progress has been made in the above mentioned domains using end-to-end localization using CNNs. This is made possible by extremely large training sets, both real and synthetic. Our proposed method relies on significantly less training data, and thus a direct comparison is not made in this work.

Registration of hand bones has been explored using multiple imaging modalities. Simple radiographs are the least expensive modality and typically the first diagnostic imaging order by rheumatologists, for diagnosis as well as disease tracking [6], [17]. The carpal bones have overlap in 2D views and present significant challenges for vision based algorithms, but have been used with moderate success [3], [25]. Volumetric modalities such as computed tomography (CT or CAT) and magnetic resonance imaging (MRI) have also been used with registration algorithms, for both parametric and non-parametric model fitting. Chen et al. [3] propose a semantic segmentation and

registration pipeline for the carpal bones from volumetric CT data using Grow Cut [25]. Mihail et al. [19] approximate these functions using dense SIFT [16] features and random decision forest (RDF) [5] classifiers. Since the relative size of hands (and bones) are similar in all images, the SIFT features at a fixed scale and orientation performed well with RDF classifiers. Our approach differs from others by not using manually engineered features. Instead, the features we use are learned using deep fully convolutional networks. Many fully convolutional network models have been proposed in the recent past, however, for biomedical image segmentation U-Net [21] is most similar to our design.

Parametric models for shape inference have been used in the past on hand radiographs. These algorithms typically rely on a set of image pre-processing or feature extraction steps for model initialization and inference. This work is closely related to that of Fernández et al. [18], who use a wire model whose wires correspond to the major axis of the hand's metacarpal and phalanges (long bones in the hand). Their method registers this wire model to previously unseen radiographs using a variation of thin plate splines (TPS) algorithm. However such representations are not well suited to modeling the joint displacements in RA patients.

Pediatric skeletal maturity estimates are used to diagnose growth disorders, timing of surgical interventions, and endocrine disorders. This estimate is typically done on hand radiographs using Greulich and Pyle atlas (G&P) [9]. Bunch et al. [2] integrated automated bone age assessment methods to increase clinician reporting quality and speed. Larson et al. [15] directly use a deep learning architecture to estimate skeletal maturity from pediatric radiographs. Guraskin et al. [10] used hand radiographs to pediatric skeletal maturity using a cascade of morphological operations that result in six features used as input to classification algorithms (support vector machines, k-nearest neighbors, decision trees, and naïve Bayes).

In contrast to methods that directly estimate a property from hand radiographs, our method provides very accurate locations of the bones that are typically used in higher level analyses.

III. PROBLEM FORMULATION

Given a roughly centered radiograph in posteroanterior (PA) view, our goal is to estimate the landmark points $s = \{s_1, \dots, s_{19}\}$ that correspond to the epiphyses (end parts of long bones) of the metacarpal and phalanges along their main axis, as seen in Figure 2.

We approach the shape fitting problem by formulating the problem probabilistically. We define the best shape s for an image I , as the shape s that maximizes $P(s|I)$. We look for a factorization of the conditional $P(s|I)$ that is both tractable to approximate and captures conditional dependencies (e.g., shape points are not in the middle of a long bone). To accomplish this, we adopt a conditional random field (CRF) to represent the distribution. We integrate a set of low-level

features, which are described below, in the potential functions of a CRF. The CRF has the following form:

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

where i is an index over model segments, j and k index our binary and unary appearance terms, Z is the partition function, Ψ and ϕ are, respectively, binary and unary appearance terms, ζ is a shape model prior and θ is a weight vector we use to balance the various terms. In the following subsections, we define the potential functions, irrespective of the low-level features, discussed in detail in Section IV-A. In the remainder of this paper, we describe our approach to solving this problem, present evaluation results, and discuss the implications of this work.

IV. APPROACH

Our pipeline consists of two stages: low-level feature extraction, and a CRF whose maximum a posteriori (MAP) inference yields the optimal shape for a given image. In this section, we describe the various components of our approach for estimating the shape s for a given hand radiograph I , including: a deep convolutional neural network for extracting features from the imagery (Section IV-A); our process for converting these features into potential functions for our CRF (Section IV-B); our shape prior (Section IV-C); and the inference process, including shape initialization and CRF optimization (Section V).

A. Convnet for hand segmentation

For our low-level feature extraction process, we first define the following functions of image locations (p is a point on an image):

- 1) $f_j(p)$: a response function where the response is highest over epiphyseal surfaces of long bones aligned with their main axes,
- 2) $f_{pc}(p)$: a response function with maximal activation over proximal (closest to the center of the body) epiphyseal surfaces,
- 3) $f_{dc}(p)$: same as f_{pc} but with maximal responses at the distal (farthest from the center of the body) epiphyseal surfaces,
- 4) $f_{bc1}(p)$: a response function with maximal responses over the main axes of a subset of the long bones (metacarpals and middle phalanges), and
- 5) $f_{bc2}(p)$: a response function with maximal responses over the main axes of long bones (proximal and distal phalanges).

In the remainder of this section, we describe how we compute these functions from image data.

We formulate low-level feature extraction as a semantic segmentation problem. Neural networks have seen a tremendous amount of attention in the ML community over the last decade

due to significant improvements in optimization algorithms and a better understanding of gradient and data propagation through the networks [12], [14], [23]. This formulation allows us to use existing knowledge from the semantic segmentation literature.

The biologically meaningful semantic classes we need are the joints and bones. Joint tissue is further divided into joint spaces, proximal and distal cortical surfaces. We group bones into two groups: 1) metacarpals and middle phalanges and 2) proximal and distal phalanges. Using this taxonomy, we propose to perform semantic segmentation using a novel fully convolutional network architecture.

Given an image I , and semantic label set $\mathcal{L} = \{l_1, l_2, \dots, l_k\}$, a semantic segmentation algorithm will assign one of k labels to each pixel in I . Often, the output of such algorithms is a distribution over the labels, and the label with the highest likelihood is chosen. If there are two classes (e.g., background and bone tissue) $|\mathcal{L}| = 2$, then our previously defined response functions directly correspond to a pixel's segmentation likelihood. We now discuss our proposed convolutional network architecture.

Our network is designed using an encoder/decoder architecture. The encoder learns a low-dimensional representation of the input image. The decoder upsamples this low-dimensional representation to the original resolution, but whose content is semantically meaningful. This approach is similar to that of Badrinarayanan et al. [1] and Ronneberger et al. [21].

We are interested in 4 semantic classes: long bone tissue along its main axis, joint tissue (between long bones), proximal epiphyseal tissue (end of the bone nearest the body) and distal epiphyseal tissue (end of the bone farthest from the body).

Our encoder stage consists of a set of convolutional layers, batch normalization, and maxpool layers. The encoder stage ends at a bottleneck, where the network learns low resolution representations of the image. The encoder is composed of 5 convolutional layers (convolution, batch normalization, and maxpool) with 16 feature maps each. Each downsampling layer in the encoder has an upsampling pair layer in the decoder stage.

The decoder stage begins at the bottleneck, consists of upsampling layers, followed at last by a softmax classifier that makes the final pixel-wise semantic label prediction. The softmax layer has the same resolution as the input image. The upsampling layers use indices from corresponding maxpool layers in the encoder stage. Each convolutional layer is followed by a batch normalization layer and a rectified linear activation (ReLU) layer. In Table I we summarize the network's convolutional layer hyperparameters during training.

B. CRF Potential Functions

The low-level features are combined in various ways to construct a set of binary and unary potential functions that our CRF optimization process uses to fit the final shape.

a) *Unary potential functions*: The unary potential functions encode the compatibility between individual shape points s_{i_p} (proximal), s_{i_d} (distal) and the image, where $i \in$

TABLE I
OUR NETWORK ARCHITECTURE. b IS THE BATCH SIZE AND c IS THE NUMBER OF SEMANTIC SEGMENTATION CLASSES.

Layer	Shape	Kernel
<i>conv0</i>	$b \times 16 \times 600 \times 460$	7×7
<i>pool0</i>	$b \times 16 \times 300 \times 230$	3×3
<i>conv1</i>	$b \times 16 \times 300 \times 230$	7×7
<i>pool1</i>	$b \times 16 \times 150 \times 115$	3×3
<i>conv2</i>	$b \times 16 \times 150 \times 115$	7×7
<i>pool2</i>	$b \times 16 \times 75 \times 57$	3×3
<i>conv3</i>	$b \times 16 \times 75 \times 57$	7×7
<i>pool3</i>	$b \times 16 \times 37 \times 28$	3×3
<i>conv4</i>	$b \times 16 \times 37 \times 28$	7×7
<i>pool4</i>	$b \times 16 \times 18 \times 14$	3×3
<i>upsample4</i>	$b \times 16 \times 37 \times 28$	3×3
<i>conv4_D</i>	$b \times 16 \times 37 \times 28$	7×7
<i>upsample3</i>	$b \times 16 \times 75 \times 57$	3×3
<i>conv3_D</i>	$b \times 16 \times 75 \times 57$	7×7
<i>upsample2</i>	$b \times 16 \times 150 \times 115$	3×3
<i>conv2_D</i>	$b \times 16 \times 150 \times 115$	7×7
<i>upsample1</i>	$b \times 16 \times 300 \times 230$	3×3
<i>conv1_D</i>	$b \times 16 \times 300 \times 230$	7×7
<i>upsample0</i>	$b \times 16 \times 600 \times 460$	3×3
<i>conv0_D</i>	$b \times 16 \times 600 \times 460$	7×7
<i>conv_00</i>	$b \times 16 \times 600 \times 460$	9×9
<i>conv_classifier</i>	$b \times c \times 600 \times 460$	3×3

$\{1, \dots, 19\}$. The first potential function we define encourages shape points to be on or near cortical surfaces of bones (both proximally and distally). This function will make use of the low-level feature function f_j through a function $df_j(p)$ (shown in Figure 3), defined as follows:

$$df_j = \begin{cases} 1 - f_j(p), & \text{if } f_j(p) \text{ is non-zero} \\ d(p, f_j) & \text{otherwise} \end{cases}$$

where $d(p, f)$ is the distance to the closest non-zero point in f from p .

We can now define the unary potential ϕ_1 as follows:

$$\phi_1(s_i) = \theta_1 df_j(s_i)$$

We further define two unary potential functions that encourage shape points to discriminate between proximal and distal cortical surfaces. These functions increase the accuracy of the final fit by encouraging an image-driven joint distance.

Intuitively, the functions ϕ_1 provides a rough alignment with joints, however, it does not discriminate between the proximal and distal bone in a joint. We make use of low-level feature functions f_{dc} and f_{pc} to define augmented distance functions df_{pc} and df_{dc} as follows:

$$df_{pc} = \begin{cases} 1 - f_{pc}(p), & \text{if } f_{pc}(p) \text{ is non-zero} \\ d(p, f_{pc}) & \text{otherwise} \end{cases}$$

and

$$df_{dc} = \begin{cases} 1 - f_{dc}(p), & \text{if } f_{dc}(p) \text{ is non-zero} \\ d(p, f_{dc}) & \text{otherwise} \end{cases}$$

where $d(p, f)$ is the same distance function to a non-zero point p in f as in Equation 2.

Our potential functions are defined as follows:

$$\phi_2(s_{i_d}) = \theta_2 df_{dc}(s_{i_d})$$

and

$$\phi_3(s_{i_p}) = \theta_3 df_{pc}(s_{i_p})$$

b) *Binary potential functions*: The potential functions Ψ_1 and Ψ_2 encode the compatibility between a shape segment i , with points i_p (proximal), i_d (distal) and bone tissue using f_{bc1} and f_{bc2} (f_{bc_x}). These functions should be at a minima when any segment s is directly over the main axis of a bone. We define an augmented distance function $df_{b_x}(p)$ from $f_{bc_x}(p)$ as follows:

$$df_{b_x} = \begin{cases} 1 - f_{bc_x}(p), & \text{if } f_{bc_x}(p) \text{ is non-zero} \\ d(p, f_{bc_x}) & \text{otherwise} \end{cases}$$

If we ensure $\max(f_{bc_x}) \leq 1$ and interpret it as the probability of pixel's p membership to the long bone semantic category, the complement of the collection of non-zero points from f_{bc_x} adds precision to the basins of attraction. Having df_{b_x} defined (shown in Figure 3), we can now define Ψ_1 as follows:

$$\Psi_1(s_{i_p}, s_{i_d}) = \theta_4 \frac{1}{n} \sum_{n=1}^t df_{b_1}(p_{x_n}, p_{y_n})$$

and

$$\Psi_2(s_{i_p}, s_{i_d}) = \theta_5 \frac{1}{n} \sum_{n=1}^t df_{b_2}(p_{x_n}, p_{y_n})$$

C. Hand Shape Potential Function

The shape prior term $\zeta(s, \theta)$ is a global term used to encourage inferred shapes to lie closely in the subspace of known shapes. We model this knowledge using Probabilistic Principal Component analysis, where the goal is to relate shape s to a k -dimensional vector x , where $k \ll \dim(s)$ and x is of zero mean and unit covariance ($I(k)$), such that:

$$s^T = Wx^T + \bar{s} + \epsilon$$

where \bar{s} is the average shape and ϵ is a normally distributed model noise component. Using this model, we can assign a probability to any shape s :

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

and formulate our prior as a weighted negative log likelihood of $P(s)$:

$$\zeta(s, \theta) = \theta_6 (-\log(P(s)))$$

D. Shape Inference

To perform shape inference, we compute the MAP solution to the CRF by taking the log-likelihood of the model, yielding an energy function E :

$$E(\hat{s}, \theta, I) = \arg \min_s \sum_i \left\{ \sum_j \Psi_j(s_{i_p}, s_{i_d}, \theta) + \sum_k \phi_k(s_i, \theta) \right\} + \zeta(s, \theta). \quad (2)$$

As we can see from Equation 2, the optimal shape \hat{s} for image I using model parameters θ is obtained when the




Fig. 3. Visualization of df_{b_1} , df_{b_2} , d_j , df_{pc} and df_{dc} .

energy function is minimized. The minimization problem is non-convex with many possible local minima, hence, a robust initialization procedure is important. The initialization routine computes a shape \hat{s} from which local search is started.

c) Shape initialization: Our initialization process consists of two stages: rough initialization using a variation of the Iterative Closest Point (ICP) algorithm and a fine-tuning local search stage using PCA. For the rough initialization stage, we compute the connected component statistics (centroid and major spanning axis) for our bone tissue responses, f_{bc_1} and f_{bc_2} . These segments are used to register an average hand with the ICP algorithm.

In the fine-tuning stage of our initialization process, we compute the PCA basis of hand shapes in our training set. Let W be the PCA basis and x be coefficients, such that a shape s can be reconstructed from coefficients as follows:

$$s = W^T x + \bar{s}$$

where \bar{s} is the shape average. Given the joint center response centroids J and bone tissue centroids B , we define the following minimization problem:

$$E(\hat{s}, x, I) = \arg \min_x \sum_{i,j} \|J_i - \hat{s}_j\|_2^2 + \sum_{k,l} \|B_k - 0.5(\hat{s}_l + \hat{s}_{l+1})\|_2^2 + \sum |x| \quad (3)$$

where i is the index of the joint centroid nearest to shape point s_j , and k is the index of the bone centroid nearest to the average of a segment $s_l s_{l+1}$, and the last term is a regularization term that encourages the inferred shape to be close to the average shape.

V. EVALUATION

We evaluated our approach on two datasets and found that our method significantly increases the accuracy for joint location estimation relative to the previous work. This result demonstrates that even with relatively small datasets it is possible to use learned feature representations to achieve




Fig. 5. Shapes with highest error in the RA dataset.

better performance than hand engineered features and shallow learning techniques.

A. Datasets

We use two datasets, the Digital Hand Atlas Database¹ and a set of 116 radiographs of RA patients from the University of Kentucky Department of Radiology. The University of Kentucky dataset contains additional annotations of rheumatoid arthritis disease stage. Both datasets have been manually annotated and verified by a radiologist. We use these ground truth shapes to train the convnets, estimate CRF model parameters and quantitatively evaluate inferred shapes and the low-level features used in the CRF inference process. Since the data is similar in both datasets, we alternate using one dataset for training and the other for testing (e.g., we trained on UK and report test on Hand Atlas, and vice-versa).

We augmented both datasets by rotating each image in 10 degree increments from -30 deg to $+30$ deg while maintaining the resolution. This was done to learn invariance to rotation and resulted in a significant improvement in robustness. For evaluation, we look at two aspects of our approach: 1) the estimated PDM and 2) the performance of our low-level feature estimators.

¹<http://ipilab.usc.edu/BAWeb/>




Fig. 4. Sample responses f_{bc_1} , f_{bc_2} , f_j , f_{pc} and f_{dc} . Best seen in color, digital format.

TABLE II
LOW-LEVEL FEATURE QUANTITATIVE EVALUATION (CNN/RDF).

	$f_{bc_1} \& f_{bc_2}$	f_j	f_{pc}	f_{dc}
Average Positive Probability	0.6223 / 0.5069	0.7128 / 0.4301	0.4830 / 0.3263	0.5960 / 0.3669
MCR (% of image)	0.0319 / 0.1297	0.0071 / 0.0252	0.0021 / 0.0137	0.0022 / 0.0088




Fig. 6. RA dataset errors computed as sum of absolute differences from ground truth for healthy, early, moderate and late stage. All models have been trained on the Hand Atlas dataset.

B. Implementation Details

The CNNs were implemented using a customized version of Caffe [13]. The resolution of the input and output feature maps was 460×600 in our experiments. Our networks were trained in CPU on a 40-core Intel blade server using a mini-batch size of 4. All models were trained on one of the datasets and tested on the other.

There are known trade-offs between the filter size and network depth in convolutional layers. We obtained the best performance with 7×7 kernels in the convolutional layers. During the CNN architecture development, we also noted

a significant performance improvement from the addition of batch normalization layers.

Our approach to network design was incremental. We started with a simple encoder-decoder architecture, with a depth of 2 layers for each stage and a fixed kernel size of 7×7 . We gained accuracy increasing the number of layers to 5. A depth of more than 5 layers per stage decreased accuracy. While there is increasing evidence of a relationship between kernel size and network depth, we obtained the best results with the structure we present in this paper.

d) Model parameter optimization: The parameter set θ from Equation 2 is done using standard model selection, as

they encode the relative weight of each feature in the shape inference process. The optimal θ are found by minimizing the global pixel-wise error for the entire training set, as follows:

$$\hat{\theta}^{ML} = \arg \min_{\theta} \sum_i \left\| \arg \min_s E(s, \theta, I_i) - s_i^{GT} \right\|_2^2 \quad (4)$$

The above minimization is non-convex, hence many local minima are possible. We use the simplex method to find the optimal model parameters.

e) Shape inference evaluation: We evaluate our model by computing the sum of absolute differences between ground truth shapes and our model inference results. Shape inference is done by minimizing Equation 2. In practice, we use a coordinate descent algorithm, with a search region per coordinate is 7 pixels. This approach converges in a few seconds on a modern desktop multi-core CPU. This algorithm was implemented in MATLAB.

C. Qualitative and Quantitative Comparisons

By visual inspection, we found the CNN features to be far superior to DSIFT+RDF. In Figure 7, we show samples of each feature function $f_j(p)$, $f_{dc}(p)$, $f_{pc}(p)$, $f_{bc_1}(p)$ for an image using both CNN and DSIFT+RDF. Visually, one can conclude that CNN features are much cleaner, with less noise and increased localization accuracy.

To quantitatively evaluate the features, we look at two performance metrics, summarized in Table II. First, for f_{bc_1} and f_{bc_2} , we take all the pixels along segments $s_{1..19}$ and compare the average probability. This number should be as close to 1 as possible, the higher the better, since responses along the main bone axes should activate. CNN features are roughly 12% more confident. We take the same approach for the other classifiers. CNN exceeds DSIFT+RDF in all cases. As a second metric, we compute the ratio of incorrectly classified pixels, specifically the misclassification rate (MCR), computed as the ratio of incorrectly classified pixels and the total number of pixels in the image (thresholding responses at ≥ 0.25). This metric gives us an insight into the localization precision of the methods. Using this metric, our proposed CNN-based feature extraction method outperforms previous results based on DSIFT+RDF for all instances.

In Figure 8, we show the errors, in pixels, on the RA dataset. The overall average error for all stages is 2.0521 pixels compared to 3.0333 pixels in previous work [19], a significant improvement. In Figure 5 we show the hands that had the highest fitting errors in the RA dataset. The overall error for the Hand Atlas dataset is 1.46 compared to 2.72 in Mihail et al. [19], on the same images with the same resolution.

VI. CONCLUSION AND FUTURE WORK

Simple hand radiographs are used in clinical practice for skeletal development assessment, rheumatic disease progression estimation, among other uses. Automatic estimation of joint spaces and long-bone segments is useful in modeling disease progression, e.g., RA. This paper improves on previous work that fits a shape model to hand radiographs. We propose

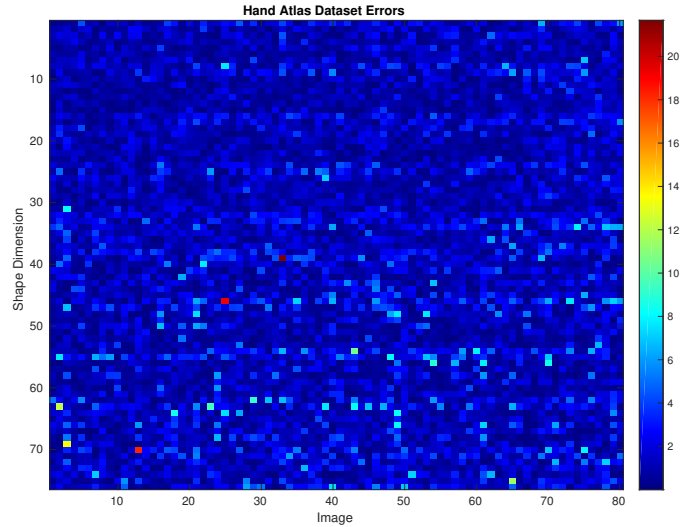


Fig. 8. Hand Atlas dataset errors computed as sum of absolute differences from ground truth. All models have been trained on the UK dataset.

a hand shape inference method from radiographs based on deep convolutional neural network features. The CNN features are superior to those obtained using shallow features, e.g. SIFT [16]. We quantitatively evaluate the features stand-alone, as well as compare the final shape inference results using shallow and deep CNN features. We note an average improvement in the shape inference of roughly 1 pixel per shape dimension compared to the previous work.


The task of shape estimation from hand radiographs is related to non-medical shape estimation tasks, for example the task of facial landmark localization [7], [8], [26]. Recent approaches for this task, and other similar tasks, propose deep learning architectures that are suitable for end-to-end optimization. Our task is potentially well suited for end-to-end optimization, but existing sets of hand radiographs are significantly smaller than those available for other domains. This leads to significant problems with overfitting, which motivated our hybrid approach of using deep learning for low-level features but CRF optimization for top-down inference. We expect that by combining our proposed techniques with larger datasets and end-to-end learning we will be able to make further advances in future work. Upon acceptance of this paper, we will release our fully trained models as well as source code that implements CRF optimization.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. IIS-1553116.

REFERENCES

- [1] V. Badrinarayanan, A. Kendall, and R. Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *arXiv preprint arXiv:1511.00561*, 2015.



Bone Joint Centers Proximal Cortical Surface Distal Cortical Surface

Fig. 7. Top: DSIFT+RDF features. Bottom: convnet features. The difference in quality and precision is evident by simple visual inspection.

- [2] P. M. Bunch, T. A. Altes, J. McIlhenny, J. Patrie, and C. M. Gaskin. Skeletal development of the hand and wrist: digital bone age companion suitable alternative to the greulich and pyle atlas for bone age assessment? *Skeletal Radiology*, 46(6):785–793, 2017.
- [3] X. Chen, J. Graham, and C. Hutchinson. Integrated framework for simultaneous segmentation and registration of carpal bones. In *Image Processing (ICIP), 2011 18th IEEE International Conference on*, pages 433–436. IEEE, 2011.
- [4] X. Chen, E. Zhou, Y. Mo, J. Liu, and Z. Cao. Delving deep into coarse-to-fine framework for facial landmark localization. In *CVPR Workshops*, pages 2088–2095, 2017.
- [5] A. Criminisi, J. Shotton, and E. Konukoglu. Decision forests for classification, regression, density estimation, manifold learning and semi-supervised learning. *Microsoft Research Cambridge, Tech. Rep. MSRTR-2011-114*, 5(6):12, 2011.
- [6] K. D. Deane, J. M. Norris, and V. M. Holers. Preclinical rheumatoid arthritis: identification, evaluation, and future directions for investigation. *Rheumatic Disease Clinics of North America*, 36(2):213–241, 2010.
- [7] Z. Deng, K. Li, Q. Zhao, and H. Chen. Face landmark localization using a single deep network. In *Chinese Conference on Biometric Recognition*, pages 68–76. Springer, 2016.
- [8] H. Fan and E. Zhou. Approaching human level facial landmark localization by deep learning. *Image and Vision Computing*, 47:27–35, 2016.
- [9] W. W. Greulich and S. I. Pyle. Radiographic atlas of skeletal development of the hand and wrist. *The American Journal of the Medical Sciences*, 238(3):393, 1959.
- [10] G. E. Güraskin, H. Uğuz, and Ö. K. Baykan. Bone age determination in young children (newborn to 6 years old) using support vector machines. *Turkish Journal of Electrical Engineering & Computer Sciences*, 24(3):1693–1708, 2016.
- [11] S. Honari, P. Molchanov, S. Tyree, P. Vincent, C. Pal, and J. Kautz. Improving landmark localization with semi-supervised learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [12] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *Advances in Neural Information Processing Systems*, pages 2017–2025, 2015.
- [13] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [15] D. B. Larson, M. C. Chen, M. P. Lungren, S. S. Halabi, N. V. Stence, and C. P. Langlotz. Performance of a deep-learning neural network model in assessing skeletal maturity on pediatric hand radiographs. *Radiology*, page 170236, 2017.
- [16] D. G. Lowe. Object recognition from local scale-invariant features. In *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, volume 2, pages 1150–1157. Ieee, 1999.
- [17] 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(1):36–42, 2005.
- [18] 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.
- [19] R. P. Mihail, G. Blomquist, and N. Jacobs. A crf approach to fitting a generalized hand skeleton model. In *Applications of Computer Vision (WACV), 2014 IEEE Winter Conference on*, pages 409–416. IEEE, 2014.
- [20] A. Pfeil, P. Oelzner, K. Bornholdt, A. Hansch, G. Lehmann, D. M. Renz, G. Wolf, and J. Böttcher. Joint damage in rheumatoid arthritis: assessment of a new scoring method. *Arthritis research & therapy*, 15(1):R27, 2013.
- [21] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [22] J. A. Singh, K. G. Saag, S. L. Bridges, E. A. Akl, R. R. Bannuru, M. C. Sullivan, E. Vaysbrot, C. McNaughton, M. Osani, R. H. Shmerling, et al. 2015 american college of rheumatology guideline for the treatment of rheumatoid arthritis. *Arthritis & rheumatology*, 68(1):1–26, 2016.
- [23] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [24] Z. Tang, X. Peng, S. Geng, L. Wu, S. Zhang, and D. Metaxas. Quantized densely connected u-nets for efficient landmark localization. *arXiv preprint arXiv:1808.02194*, 2018.
- [25] V. Vezhnevets and V. Konouchine. Growcut: Interactive multi-label image segmentation by cellular automata. In *proc. of Graphicon*, volume 1, pages 150–156, 2005.
- [26] L. Wang, X. Yu, and D. N. Metaxas. A coupled encoder-decoder network for joint face detection and landmark localization. In *Automatic Face & Gesture Recognition (FG 2017), 2017 12th IEEE International Conference on*, pages 251–257. IEEE, 2017.