

# PIANO: A Parametric Hand Bone Model from Magnetic Resonance Imaging

Yuwei Li<sup>1,2,3</sup>, Minye Wu<sup>1,2,3</sup>, Yuyao Zhang<sup>1</sup>, Lan Xu<sup>1</sup> and Jingyi Yu<sup>1</sup>

<sup>1</sup>Shanghai Engineering Research Center of Intelligent Vision and Imaging, School of Information Science and Technology, ShanghaiTech University

<sup>2</sup>Shanghai Institute of Microsystem and Information Technology, Chinese Academy of Sciences

<sup>3</sup>University of Chinese Academy of Sciences

{liywu, wumy, zhangyy8, xulan1, yujingyi}@shanghaitech.edu.cn

## Abstract

Hand modeling is critical for immersive VR/AR, action understanding, or human healthcare. Existing parametric models account only for hand shape, pose, or texture, without modeling the anatomical attributes like bone, which is essential for realistic hand biomechanics analysis. In this paper, we present PIANO, the first parametric bone model of human hands from MRI data. Our PIANO model is biologically correct, simple to animate, and differentiable, achieving more anatomically precise modeling of the inner hand kinematic structure in a data-driven manner than the traditional hand models based on the outer surface only. Furthermore, our PIANO model can be applied in neural network layers to enable training with a fine-grained semantic loss, which opens up the new task of data-driven fine-grained hand bone anatomic and semantic understanding from MRI or even RGB images. We make our model publicly available.

## 1 Introduction

Hands are essential for humans to interact with the world and hand modeling enables numerous applications for immersive VR/AR, action analysis and healthcare, etc. However, reconstructing realistic human hands in a biologically and anatomically plausible manner remains challenging and has received substantive attention in both the artificial intelligence and computer vision communities.

In recent years, there has been significant progress in hand pose, shape, geometry [Wang *et al.*, 2020; Smith *et al.*, 2020], and even texture estimation [Qian *et al.*, 2020] in a data-driven manner with the corresponding parametric models [Romero *et al.*, 2017; Moon *et al.*, 2020]. In contrast, there is little work that addresses the parametric modeling and data-driven estimation of human hand anatomy like bone and muscle. However, such fine-grained hand anatomical attribute modeling is important for understanding the complex hand biomechanics and producing realistic and anatomically precise virtual hands. On the other hand, modern 3D medical imaging techniques like Magnetic Resonance Imaging (MRI) and Computed Tomography (CT) are widely used for treating hand injuries clinically [Kapandji and others,




Figure 1: We present the first statistical internal hand bone model. Our model enables anatomically and physically precise modeling of the hand kinematic structure of different individuals, opening up the study of data-driven hand bone anatomic and semantic fine-grained understanding from MRI or even RGB images.

2007] or modeling the function of healthy hands [Miyata *et al.*, 2005]. However, the above-mentioned modeling methods focus on the analysis of only a single individual. Moreover, those CT-based methods [Kurihara and Miyata, 2004; Yang *et al.*, 2020] suffer from ethical issues due to the unacceptably high radiation doses during the scanning of various individuals under various hand poses for building a high-quality parametric hand bone model. The recent MRI-based method [Wang *et al.*, 2019] adopts a molding procedure to build an anatomy-correct hand bone rig of a target performer. However, it suffers from time-consuming and user-specified bone segmentation operations, which is impractical to apply to various individuals for building a parametric model.

To address the gaps above, we present the first statistical hand bone model from MRI data called *PIANO* for *Parametric hANd bOne model*, which enables anatomically and physically precise modeling of the hand kinematic structure of different individuals (see Fig. 1). Our model is biologically correct, simple to animate, and differentiable, opening up the study of data-driven hand bone anatomic and semantic fine-grained understanding from MRI or even RGB images at a level of detail not previously possible.

To this end, we first collect a new large-scale dataset of high-quality MRI hand scans using a clinical 3T scanner and

an efficient simplified molding procedure from [Wang *et al.*, 2019], which consists of 200 scans of 35 subjects in up to 50 poses with rich and accurate 3D joint and bone segmentation annotations from experienced radiologists. Then, we use this data as well as an artist-rigged hand bone template mesh with 20 separated bones to train the statistical hand bone model PIANO similar to the outer human body surface model SMPL [Loper *et al.*, 2015] and hand model MANO [Romero *et al.*, 2017]. Due to the inherent per-bone rigidity in PIANO, we omit the pose blend shape and adopt a pre-defined blending weight. Thus, PIANO inherently factors the per-bone anatomic geometric changes of each individual and has analogous components to those in SMPL and MANO: a hand bone template mesh with kinematic tree, shape blend shape, and a joint regressor. Especially, to build PIANO in an anatomically precise manner, we introduce a novel coarse-to-fine multi-stage registration approach and a bone-aware model learning strategy, which jointly considers the accurate joint and bone segmentation annotations provided by radiologists. Our new parametric hand bone model allows obtaining more anatomically precise personalized rigging than the traditional hand models based on outer surface only. Furthermore, we introduce a novel neural network based on PIANO to produce fine-grained and biologically correct hand bone segmentation from an MRI scan or even a single RGB image input in an analysis-by-synthesis fashion. It enables a fine-grained semantic loss and opens up a new task for data-driven internal hand bone modeling unseen before. To summarize, our main technical contributions include:

- We present the first parametric hand bone model, PIANO, based on a new MRI hand dataset, which enables anatomically precise internal kinematic understanding of human hands in a data-driven manner.
- We propose a novel multi-stage registration approach as well as a bone-aware model learning strategy, which utilizes the joint and bone segmentation annotations in MRI data provided by radiologists to build PIANO.
- We present a proof-of-concept neural network design that enables data-driven fine-grained hand bone segmentation from an MRI scan or RGB image.
- We make available our PIANO model and the new dataset of 200 MRI hand scans of 35 subjects with various poses and rich annotations<sup>1</sup>.

## 2 Related Work

**Parametric Modeling.** Using parametric model for human performance capture on body [Loper *et al.*, 2015] and hand [Romero *et al.*, 2017; Smith *et al.*, 2020] have been widely studied for the past decades. Most of them focus on outer information like skin surface and texture [Qian *et al.*, 2020; Moon *et al.*, 2020; Van Houtte *et al.*, 2018]. [Mirakhorlo *et al.*, 2018] proposed a comprehensive biomechanical hand model based on detailed measurements from a hand specimen, yet it is not differentiable and can not be embedded in deep learning frameworks. [Anas *et al.*, 2016] proposed a statistical wrist shape and bone model for automatic carpal

bone segmentation. On this basis, we utilize medical imaging techniques and build a parametric hand bone model which is anatomically correct, physically precise and differentiable.

**MRI/CT hand analysis.** Common medical imaging techniques that are widely used for hand treatment include Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) [Kapandji and others, 2007]. Pioneering work of [Kurihara and Miyata, 2004] and [Marai *et al.*, 2003] have used CT scans to acquire hand bone skeleton and carpal movements. [Martín-Fernández *et al.*, 2009] proposed automatic articulated registration of hands out of X-ray images. However, CT and X-ray emits ionizing radiation, while MRI is safer and provides more clear contrast between soft body tissue. [Rhee *et al.*, 2007; Miyata *et al.*, 2005] used MRI for understanding the hand anatomy and biomechanics on a small set of poses. Recently, [Wang *et al.*, 2019] adopts a molding procedure for long time MRI scan, enabling anatomical hand bone analysis with complex poses. However, suffering from time-consuming and user-specified operations, these methods only focus on single individuals. Comparably, we employ a simplified molding procedure and a semi-automatic annotation method to build a large-scale dataset of high-quality MRI hand scans with precise annotation on bone and joints.

## 3 Parametric Hand Bone Model

In this work, our goal is to build a parametric internal bone model of human hands, which is anatomically correct, simple to animate and differentiable. First, we build a new MRI hand dataset via an efficient molding procedure, which captures precise anatomical structure of human hands with rich annotations provided by experienced radiologists (Sec. 3.1). Then, similar to outer surface modeling, we train our PIANO model to obtain the hand bone template mesh with kinematic tree, shape blend shape and joint regressor by iterating between model registration and model parameter learning in a flip-flop manner. Our detailed model formulation is provided in Sec. 3.2, followed by the illustration of our novel multi-stage model registration (Sec. 3.3) and our bone-aware parameter learning scheme (Sec. 3.4), so as to utilize the rich annotations in MRI data provided by radiologists.

### 3.1 Hand MRI Data Collection

Magnetic Resonance Imaging (MRI) scanning is performed to capture the anatomical structure under each hand pose. In all, a total of 200 hand volumetric images are captured with 50 different hand posture of 35 individual subjects. The scanning is performed with a 3.0 Tesla MR scanner of United Imaging (uMR780, Shanghai) with spatial resolution  $0.4 \times 0.4 \times 1 \text{ mm}^3$  and T1-weighted sequence. Since MRI is a motion-sensitive imaging technique, during scanning, we adopt a simplified molding procedure to stabilize hand poses similar as that in [Wang *et al.*, 2019]. Briefly, we put hands in a liquid solution which would solidify and encase the hand firmly in less than 3 minutes. We then scan the hand with the mold on for 10 minutes.

The MRI volume is a regular volumetric grid of scalar values representing tissue mass density. To analyze anatomical structure, bone and joint location must be annotated. How-

<sup>1</sup><https://reyuwei.github.io/proj/piano.html>



Figure 2: Left: MRI scanner and participant. Middle: MRI raw slices. Right: Dataset samples. From top to bottom are MRI volume, binary bone mask and bone mesh overlaid with 25 precise joint annotation.

ever, this task is highly cumbersome and requires strong domain knowledge. [Wang *et al.*, 2019] proposed an interactive method that takes 3 hours to segment one volume. However, to label all 200 scans in our dataset would take up to 600 hours. To this end, we adopt a semi-automated human-in-the-loop procedure to label all volumes with a deep neural network. By verifying network predictions and annotating miss labeled samples in an iterative procedure, we only acquire one-third of volume annotations with a moderate manual effort by experienced radiologists. Dataset samples are shown in Fig.2.

### 3.2 Model Formulation

To learn an anatomically correct and simple-to-animate statistical model from our MRI dataset, we follow the outer surface model SMPL [Loper *et al.*, 2015] and MANO [Romero *et al.*, 2017] to formulate our PIANO model into analogous components, including a hand bone template mesh with kinematic tree, shape blend shape and a joint regressor.

As illustrated in Fig. 3, we utilize an artist-rigged hand bone template mesh consisting of 20 separated semantic bone meshes with 3345 vertices, 6610 faces and 19 joints plus the global orientation. Then, our PIANO model defines a function  $\mathcal{M}$  of shape  $\beta$  and pose  $\theta$ :

$$\mathcal{M}(\beta, \theta) = LBS(\bar{T} + B_S(\beta), J(\beta), \theta, \mathcal{W}). \quad (1)$$

Here,  $LBS(\cdot)$  denotes the Linear Blend Skinning (LBS) function applied to the warped bone mesh model  $T = \bar{T} + B_S(\beta)$ , skinning weight  $\mathcal{W}$  and joint rotation  $\theta$ . To model inherent per-bone rigidness, we adopt a fixed blending weight pre-defined by bone label intuitively and omit the pose blend shape. Thus, we utilize the canonical template mesh  $\bar{T} \in \mathbf{R}^{3N}$  of  $N = 3345$  vertices under zero pose and shape, as well as the shape blend shape  $B_S(\beta)$  to factor the per-bone anatomic geometric changes of each individual. In PIANO, we formulate the linear shape blend shape function as  $B_S(\beta; \mathcal{S}) = \mathcal{S}\beta$  to enable the base shape to vary with identity, where  $\beta$  is the shape parameters and matrix  $\mathcal{S} \in \mathbf{R}^{3N \times |\beta|}$  represents orthonormal principal components of shape displacements, which is learned from registered training meshes. Besides, in PIANO we formulate the joint regressor  $J(\beta) = \mathcal{J}(\bar{T} + B_S(\beta))$  to map the various




Figure 3: PIANO Model. (a) Template bone mesh with skinning weight indicated by color and joints. (b) With subject-specific shape blend shape. (c) Reposed with linear blend skinning.

shapes to the corresponding anatomically precise joint positions of human hands, where  $\mathcal{J} \in \mathbf{R}^{3K \times |\beta|}$  is the learned regression matrix that transforms rest pose vertices into rotation joints;  $K = 25$  contains 19 rotation joints, 5 finger tips and 1 wrist keypoint, matching our dataset annotation.

Our final PIANO model is  $\mathcal{M}(\beta, \theta; \Phi) : \mathbf{R}^{|\beta| \times |\theta|} \mapsto \mathbf{R}^{3N}$  which maps the shape and pose parameters to bone vertices of various individuals ( $|\beta| = 10$  and  $|\theta| = 57$  in our setting). The model parameters  $\Phi = \{\bar{T}, \mathcal{S}, \mathcal{J}\}$  includes the canonical template, shape blend shape and joint regressor.

### 3.3 Model Registration

Here, we propose a multi-stage registration scheme to align the bone template to all the MRI scans, bringing them into correspondence. Our scheme encodes the internal structure information in a coarse-to-fine manner to fully utilize the anatomically precise joint and bone segmentation annotations in our MRI dataset. With the joint annotations and bone segmentation, we utilize the radial basis functions (RBF) interpolation [Rhee *et al.*, 2007] to assign per-bone label of each joint and generate fine-grained segmentation of the bone volume. Then, we extract both the iso-surface of internal bones  $b^S$  and the per-bone 3D geometry  $b_l^S$  of the  $l$ -th joint/bone by applying the Marching Cubes [Lorensen and Cline, 1987].

**Coarse Registration.** We first align the warped joints of the hand bone template mesh to the joint annotations by minimizing the per-joint L2 error and solving the inverse kinematics to obtain the initial pose and shape parameters. Since the model parameters  $\Phi$  are unknown during the first iteration of model registration and parameter learning, we downgrade our model from Eq. 1 into  $\mathcal{M}'(\theta) = LBS(\bar{T}, \bar{J}, \theta, \mathcal{W})$ , where we discard shape blend shape function and replace the subject specific parameters  $\{T, J\}$  with a general template  $\{\bar{T}, \bar{J}\}$ .

**Fine Registration.** We further utilize the internal local and global structural information to perform per-bone non-rigid alignment at a finer scale. To this end, for the  $l$ -th bone, we adopt the embedded deformation [Xu *et al.*, 2019b] to register from the initial geometry  $b_l^R$  after coarse registration to the annotated geometry  $b_l^S$  and  $b^S$  from segmentation label. Let  $G$  denote the non-rigid motion field and  $ED(\cdot; G)$  denote the embedded deformation function. Then, our per-bone alignment is formulated as:

$$\mathbf{E}_{\text{fine}}(G) = \mathbf{E}_{\text{data}}(G) + \lambda_{\text{reg}} \mathbf{E}_{\text{reg}}(G). \quad (2)$$

Here, the data term enforces alignment of the warped bone geometry to both current target bone segment  $b_l^S$  and the global bone structure  $b^S$  to enable robust alignment:




Figure 4: Result of our multistage registration. (a) MRI volume and PIANO template. (b) Coarse stage registration to joint annotation. (c) Fine stage registration to fine-grained bone mesh. We show a color coded Hausdorff distance between registration and annotation.

$$E_{\text{data}}(G) = \sum_{(v,u) \in \mathcal{C}} \|ED(v) - u\|_2^2 + \sum_{(v,u) \in \mathcal{P}} \|ED(v) - u\|_2^2, \quad (3)$$

where  $\mathcal{C}$  and  $\mathcal{P}$  denote the the set of correspondences from  $b_i^R$  to  $b_i^S$  and  $b_i^S$ , respectively, in an Iterative Closest Point (ICP) manner. To prevent over-fitting to the noise of annotation, we adopt the same locally as-rigid-as-possible motion regularization  $\lambda_{\text{reg}}$  and  $E_{\text{reg}}$  from [Xu *et al.*, 2019b]. Fig. 4 shows the intermediate alignment results and demonstrates the effectiveness of our registration scheme. After registration, we obtain per-bone aligned meshes with global consistent topology for the following model parameter learning process.

### 3.4 Parameter Training

To train our PIANO model parameters  $\Phi = \{\bar{T}, \mathcal{S}, \mathcal{J}\}$ , similar to previous work [Loper *et al.*, 2015; Romero *et al.*, 2017], we adopt a bone-aware training strategy to model pose and shape separately. Let  $b_{i,j,l}$  denote the registered  $l$ -th bone geometry of the  $i$ -th subject in the  $j$ -th hand pose and  $J_{i,j}^*$  denote the corresponding annotated 3D joint. Then, we optimize the pose parameters  $\theta_{i,j}$  that transforms a general template  $\bar{T}$  to a specific bone model  $b_{i,j}$  with the following energy function:

$$E_{\text{pose}}(\theta_{i,j}) = \lambda_e E_{\text{edge}}(\theta_{i,j}) + \lambda_j E_{\text{joint}}(\theta_{i,j}). \quad (4)$$

Here, similar to SMPL, the edge term measures the edge length difference between two models:

$$E_{\text{edge}}(\theta_{i,j}) = \sum_{ij} \sum_l \|\mathcal{E}(LBS(\bar{T}, \bar{J}, \theta_{i,j}, \mathcal{W})_l) - \mathcal{E}(b_{i,j,l})\|^2, \quad (5)$$

where  $\mathcal{E}(\cdot)$  represents the geometry edges. Such edge difference enables good estimation of pose without knowing the subject specific shape. Then, we add a joint term based on the anatomically precise joint annotation in our dataset:

$$E_{\text{joint}}(\theta_{i,j}) = \sum_{ij} \|\bar{J}_{i,j} - J_{i,j}^*\|_2^2, \quad (6)$$

where  $\bar{J}_{i,j}$  is the joint position after LBS with the pose  $\theta_{i,j}$ .

Next, we fix the above pose parameters and optimize the subject template  $\hat{T}_i$  and joint  $\hat{J}_i$  iteratively with the following




Figure 5: Fine-grained MRI bone segmentation network structure.

vertex difference defined on semantic bone segments:

$$E_{\text{vert}}(\hat{T}_i, \hat{J}_i) = \sum_{ij} \sum_l \|LBS(\hat{T}_i, \hat{J}_i, \theta_{i,j}, \mathcal{W})_l - b_{i,j,l}\|^2, \quad (7)$$

which enforces the posed subject template to match with the aligned bone meshes. We then run principal component analysis (PCA) on  $\hat{T}_i$  to obtain shape space parameters  $\{\bar{T}, \mathcal{S}\}$ , where  $\bar{T}$  is the mean shape of  $\hat{T}_i$  and  $\mathcal{S}$  is a matrix composed of principal component vectors. Also, given subject template  $\hat{J}_i$  and  $\bar{T}_i$  in rest pose, we further compute the joint regression matrix  $\mathcal{J}$  by minimizing  $\|\mathcal{J}\hat{T}_i - \hat{J}_i\|_2^2$ .

To build PIANO, we set  $\lambda_{\text{reg}}$ ,  $\lambda_e$  and  $\lambda_j$  to 5.0, 1.0 and 100.0 respectively and iterate through the whole process of registration and parameter training with Limited-memory BFGS optimizer and PyTorch auto differentiation.

## 4 Fine-grained Semantic Bone Analysis

We now demonstrate the use cases of our statistical model in hand bone anatomic and semantic understanding. We consider the usage of PIANO as a differentiable layer enabling end-to-end MRI fine-grained bone segmentation and joint estimation. Furthermore, we show hand anatomic understanding from a single monocular RGB image.

**PIANO Layer.** We integrate PIANO as a differentiable layer that takes  $\theta$  and  $\beta$  as input and outputs a triangulated semantic bone mesh and hand joints. When training with supervision on PIANO parameters, we define the parameter loss  $\mathcal{L}_{pm}$  on predicted shape and pose  $[\hat{\beta}, \hat{\theta}]$  as:

$$\mathcal{L}_{pm} = \|[\hat{\beta}, \hat{\theta}] - [\beta^*, \theta^*]\|_2^2, \quad (8)$$

where  $[\beta^*, \theta^*]$  denotes ground truth. Alternatively, for training with sparse joint annotations, we define a joint loss as:

$$\mathcal{L}_J = \|\hat{J} - J^*\|_2^2. \quad (9)$$

Here,  $\hat{J}$  denotes the estimated joints and  $J^*$  is the 3D joint annotation. The resulting PIANO loss  $\mathcal{L}_P$  is thus defined as  $\mathcal{L}_P = \mathcal{L}_{pm} + \mathcal{L}_J$ .

**Anatomic inference from MRI.** As illustrated in Fig.5, we propose a novel network structure for fine-grained MRI hand bone segmentation with PIANO layer. Our encoder consists of ResNet3D and three regression branches. Each has two fully connected layers with hidden size 512 and ReLU as activation function after the first layer. The encoder extracts a latent feature from an MRI volume. Then, the extracted feature is passed to regression branches to infer pose  $\theta$  and shape  $\beta$  of PIANO as well as global rotation and translation  $\mathcal{G}$ . PIANO layer then outputs a fine-grained semantic bone mesh.



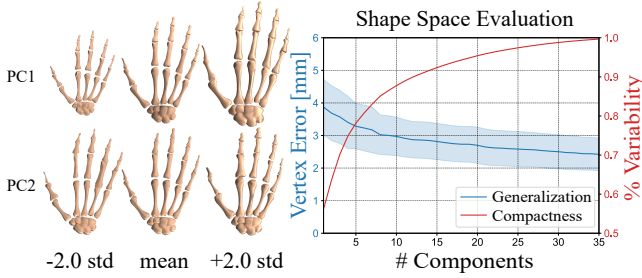


Figure 6: Left: visualization of shape space projected to the first two PCs. Right: generalization and compactness curve of PIANO. We plot the mean and standard deviation of the vertex mean squared error for generalization and the percent of variability recovered with various number of PCs for compactness.

The bone mesh is then voxelized to a regular volumetric grid with scalar values representing semantic bone label. The SegNet takes the concatenation of the voxelized mesh and original volume as input and outputs label probabilities. We train this network on our training set, which contains annotation of 3D joints and fine-grained bone mask. Thus the loss term can be defined as the summation of fine-grained semantic segmentation loss and PIANO loss, where the segmentation loss is the linear combination of multi-class binary cross entropy loss (BCE) and dice loss (DC):

$$\mathcal{L}_{BCE} = -\frac{1}{K} \sum_{l=1}^K (y_l \log \hat{y}_l + (1 - y_l) \log(1 - \hat{y}_l)); \quad (10)$$

$$\mathcal{L}_{DC} = 1 - \frac{2}{K} \sum_{l=1}^K \frac{y_l \hat{y}_l}{y_l + \hat{y}_l}, \quad (11)$$

where  $\hat{y}$  is the sigmoid output of the network,  $y$  is the one hot encoding of the ground truth segmentation.  $K = 20$  is number of fine-grained classes.

**Anatomic inference from RGB.** To extend the application to RGB images, similar to the MRI-based inference, we build a network with an image encoder and two regression branches that infer pose and shape on the FreiHand [Zimmermann *et al.*, 2019] with 3D joint supervision. To account for the mismatch of joint positions, we add an additional linear layer to map from our joint to the dataset joint annotation. Since we only have sparse joint supervision for this task, following [Hasson *et al.*, 2019], we use a regularizer on the hand shape to prevent extreme mesh deformations and formulate it as  $\mathcal{L}_\beta = \|\beta\|^2$ , which constraints the hand shape to be close to the average shape in our MRI dataset, thus the loss is defined as  $\mathcal{L}_J + \mathcal{L}_\beta$ .

## 5 Experiment

In the following section, we present quantitative and qualitative evaluation of PIANO and the corresponding applications.

### 5.1 PIANO Evaluation

We first evaluate the shape space of PIANO qualitatively on the left of Fig.6, which provides a visualization of our shape




Figure 7: Deformation comparison between MANO and PIANO. The interior bone suffers from implausible bending in (a) MANO, while our model is able to maintain the bone rigidity in (b) PIANO. The bottom right is a real MRI slice of the corresponding pose.

Methods	PCK↑	MSE (mm) ↓
FetalPose	0.621	18.529
MANO	0.777	11.167
PIANO	<b>0.907</b>	<b>4.607</b>

Table 1: Comparison with non-model based method and MANO in terms of hand joint detection error from MRI scans. We use 50 mm as the up threshold for the PCK metric.

space projected to the first two principal components (PCs). The effect of each PC is illustrated by adding  $\pm 2$  standard deviations (std) to the mean shape in the center, which depicts that the first component restricts the hand size, while the second component controls the bone thickness. We further evaluate the compactness and generalization of our model quantitatively on the right of Fig.6.

**Compactness.** The red curve in Fig. 6 illustrates the compactness of our model shape space. It shows the shape variance of our dataset recovered by a varying number of components. Note that the first principal component covers over 50% of the shape space. Meanwhile, 10 and 20 components manage to cover 90% and 95% of the complete shape space.

**Generalization.** The blue curve in Fig. 6 illustrates the generalization of our model. We use the leave-one-out evaluation of the shapes in our dataset and report the mean squared vertex error. Such quantitative evaluation depicts the generalization ability of our model to unseen shapes. As the number of principal components increases, the mean error decreases to a minimum error of 2.5 mm achieved by the full space.

### 5.2 Comparison

Here we compare our model on various internal hand bone inference tasks. For the MRI-based tasks, we separate our MRI dataset into 140 training volumes and 60 testing volumes, where testing volumes contain unseen hand shapes and poses. We train our approach with the PIANO model only on the training set and evaluate on the test set.

**Compare with MANO.** We compare the joint estimation ability of our PIANO model against MANO [Romero *et al.*, 2017] by training both models using the encoder network proposed in Sec.4 with joint loss and shape regularization loss. Note that we adopt full dimension for pose parameter and 10 dimension for shape parameters, and use Probability of Correct Keypoint (PCK) and Mean Squared Error (MSE) to measure the joint detection error. As shown in Tab. 1, our model achieves more accurate results consistently on both metrics




Figure 8: Results on fine-grained anatomic and semantic understanding of human hands from MRI scans.




Figure 9: Comparison on fine-grained MRI bone segmentation. (a) Input volume. (b, c, d) Results from 3D-Unet, VNet and KiU-Net. (e) PIANO fitting results. (f) Our network results. (g) Ground truth.



Figure 10: Qualitative results of RGB-based hand bone analysis.

due to our anatomically precise modeling ability. Fig. 7 further reveals that MANO suffers from impractical bone deformations since it is based on outer surface only. In contrast, our PIANO model achieves more physically plausible deformation with internal modeling. For a more thorough comparison, we also compare against the non-model based method FetalPose [Xu *et al.*, 2019a] on MRI joint estimation and report the results in Tab. 1. It can be seen that model based methods out performs FetalPose on both metrics.

**Fine-grained MRI Bone Segmentation.** We adopt PIANO and the network in Fig. 5 for MRI hand bone segmentation and Fig. 8 shows our results on various challenging hand poses. We compare with fully supervised medical segmentation baselines VNet [Milletari *et al.*, 2016], 3D-Unet [Çiçek *et al.*, 2016] and KiU-Net [Valanarasu *et al.*, 2020]. For quantitative analysis, the dice score (DICE) and the symmetric hausdorff distance (Hausdorff) are adopted as metrics. As shown in Fig. 9 and Tab. 2, the baselines suffer from severe local inference noise, while the approach using our PIANO model provides more complete results, significantly improving the performance of fine-grained bone segmentation.

**Anatomic Inference from RGB.** In Fig. 10, we show the

Methods	DICE $\uparrow$	Hausdorff (mm) $\downarrow$
VNet	0.625	27.539
3D-Unet	0.663	23.636
KiU-Net	0.695	22.839
Ours	<b>0.717</b>	<b>4.02</b>

Table 2: Comparison of fine-grained MRI-based bone segmentation.

qualitative results for obtaining PIANO model from a single RGB image, which is taken from the evaluation set of the FreiHand dataset. Note that our PIANO results are overlaid with the MANO surface in Fig. 10. While MANO captures the outer hand surface, our model is able to estimate inner bone structure which is anatomically precise.

## 6 Conclusion

We have presented PIANO, the first parametric hand bone model which enables anatomically precise modeling of the internal kinematic structure of human hands in a data-driven manner. Our model is biologically correct, simple to animate and differentiable, based on a new dataset of high-quality MRI hand scans of various subjects and poses. For model creation, our multi-stage registration approach and the bone-aware model learning strategy enable us to utilize the joint and bone segmentation annotations in MRI data provided by radiologists. Furthermore, we demonstrated that our model enables more anatomically precise personalized rigging and fine-grained hand bone segmentation from an MRI scan or even a single RGB image in a data-driven fashion. We make our model publicly available to stimulate further work and potentially benefit virtual hands in games/film, robotics, human healthcare and medical education. In future, we plan to extend our parametric model to muscles, tendons and fats using our MRI dataset for more ingenious hand modeling.

## Acknowledgments

This work was supported by NSFC programs (61976138, 61977047), the National Key Research and Development Program (2018YFB2100500), STCSM (2015F0203-000-06) and SHMEC (2019-01-07-00-01-E00003). We would like to thank Hui Yang, Xiaozhao Liu, Xin Tang, Zhan Jin, Guohao Jiang, Lang Mei, Lu Zou, Boliang Yu, and Yongshen Li for the help of data acquisition.

## References

- [Anas *et al.*, 2016] E. M. A. Anas, A. Rasoulia, A. Seitel, K. Daras, D. Wilson, P. S. John, D. Pichora, P. Mousavi, R. Rohling, and P. Abolmaesumi. Automatic segmentation of wrist bones in ct using a statistical wrist shape + pose model. *IEEE Transactions on Medical Imaging*, 35(8):1789–1801, 2016.
- [Çiçek *et al.*, 2016] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, pages 424–432. Springer, 2016.
- [Hasson *et al.*, 2019] Yana Hasson, Gul Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J Black, Ivan Laptev, and Cordelia Schmid. Learning joint reconstruction of hands and manipulated objects. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11807–11816, 2019.
- [Kapandji and others, 2007] Ibrahim Adalbert Kapandji et al. *The physiology of the joints*. Churchill Livingstone, 2007.
- [Kurihara and Miyata, 2004] Tsuneya Kurihara and Natsuki Miyata. Modeling deformable human hands from medical images. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 355–363, 2004.
- [Loper *et al.*, 2015] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015.
- [Lorensen and Cline, 1987] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics*, 21(4):163–169, 1987.
- [Marai *et al.*, 2003] Georgeta E Marai, David H Laidlaw, James J Coburn, Mohammad A Upal, and Joseph J Crisco. A 3d method for segmenting and registering carpal bones from ct volume images. In *Proc. of Annual Meeting of the American Society of Biomechanics*, 2003.
- [Martín-Fernández *et al.*, 2009] Miguel Á Martín-Fernández, Rubén Cárdenes, Emma Muñoz-Moreno, Rodrigo de Luis-García, Marcos Martín-Fernández, and Carlos Alberola-López. Automatic articulated registration of hand radiographs. *Image and Vision Computing*, 27(8):1207–1222, 2009.
- [Milletari *et al.*, 2016] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016.
- [Mirakhorlo *et al.*, 2018] M Mirakhorlo, N Van Beek, M Wesseling, H Maas, HEJ Veeger, and I Jonkers. A musculoskeletal model of the hand and wrist: model definition and evaluation. *Computer methods in biomechanics and biomedical engineering*, 21(9):548–557, 2018.
- [Miyata *et al.*, 2005] N. Miyata, M. Kouch, M. Mochimaru, and T. Kurihara. Finger joint kinematics from mr images. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2750–2755, 2005.
- [Moon *et al.*, 2020] Gyeongsik Moon, Takaaki Shiratori, and Kyoung Mu Lee. Deephandmesh: A weakly-supervised deep encoder-decoder framework for high-fidelity hand mesh modeling. In *European Conference on Computer Vision*, pages 440–455. Springer, 2020.
- [Qian *et al.*, 2020] Neng Qian, Jiayi Wang, Franziska Mueller, Florian Bernard, Vladislav Golyanik, and Christian Theobalt. Html: A parametric hand texture model for 3d hand reconstruction and personalization. In *European Conference on Computer Vision*, pages 54–71. Springer, 2020.
- [Rhee *et al.*, 2007] Taehyun Rhee, John P Lewis, Ulrich Neumann, and Krishna Nayak. Soft-tissue deformation for in vivo volume animation. In *15th Pacific Conference on Computer Graphics and Applications (PG’07)*, pages 435–438. IEEE, 2007.
- [Romero *et al.*, 2017] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics (TOG)*, 36(6):245, 2017.
- [Smith *et al.*, 2020] Breannan Smith, Chenglei Wu, He Wen, Patrick Peluse, Yaser Sheikh, Jessica K Hodgins, and Takaaki Shiratori. Constraining dense hand surface tracking with elasticity. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020.
- [Valanarasu *et al.*, 2020] Jeya Maria Jose Valanarasu, Vishwanath A Sindagi, Ilker Hacihaliloglu, and Vishal M Patel. Kiu-net: Towards accurate segmentation of biomedical images using over-complete representations. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part IV 23*, pages 363–373. Springer, 2020.
- [Van Houtte *et al.*, 2018] Jeroen Van Houtte, Kristina Stanković, Brian G Booth, Femke Danckaers, Véronique Bertrand, Frederik Verstrecken, Jan Sijbers, and Toon Huysmans. An articulating statistical shape model of the human hand. In *International Conference on Applied Human Factors and Ergonomics*, pages 433–445. Springer, 2018.
- [Wang *et al.*, 2019] Bohan Wang, George Matcuk, and Jernej Barbič. Hand modeling and simulation using stabilized magnetic resonance imaging. *ACM Transactions on Graphics (TOG)*, 38(4):1–14, 2019.
- [Wang *et al.*, 2020] Jiayi Wang, Franziska Mueller, Florian Bernard, Suzanne Sorli, Oleksandr Sotnychenko, Neng Qian, Miguel A Otaduy, Dan Casas, and Christian Theobalt. Rgb2hands: real-time tracking of 3d hand interactions from monocular rgb video. *ACM Transactions on Graphics (TOG)*, 39(6):1–16, 2020.
- [Xu *et al.*, 2019a] Junshen Xu, Molin Zhang, Esra Abaci Turk, Larry Zhang, P Ellen Grant, Kui Ying, Polina Golland, and Elfar Adalsteinsson. Fetal pose estimation in volumetric mri using a 3d convolution neural network. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 403–410. Springer, 2019.
- [Xu *et al.*, 2019b] Lan Xu, Wei Cheng, Kaiwen Guo, Lei Han, Yebin Liu, and Lu Fang. Flyfusion: Realtime dynamic scene reconstruction using a flying depth camera. *IEEE Transactions on Visualization and Computer Graphics*, 2019.
- [Yang *et al.*, 2020] Xiaopeng Yang, Zhichan Lim, Hayoung Jung, Younggi Hong, Mengfei Zhang, Dougho Park, and Heecheon You. Estimation of finite finger joint centers of rotation using 3d hand skeleton motions reconstructed from ct scans. *Applied Sciences*, 10(24):9129, Dec 2020.
- [Zimmermann *et al.*, 2019] Christian Zimmermann, Duygu Ceylan, Jimei Yang, Bryan Russell, Max Argus, and Thomas Brox. Freihand: A dataset for markerless capture of hand pose and shape from single rgb images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 813–822, 2019.