

# FIND: An Unsupervised Implicit 3D Model of Articulated Human Feet

Oliver Boyne  
ob312@cam.ac.uk  
James Charles  
<http://www.jjcvision.com>  
Roberto Cipolla  
cipolla@eng.cam.ac.uk

Machine Intelligence Lab  
Department of Engineering  
University of Cambridge  
Cambridge, U.K.

## Abstract

In this paper we present a high fidelity and articulated 3D human foot model. The model is parameterised by a disentangled latent code in terms of shape, texture and articulated pose. While high fidelity models are typically created with strong supervision such as 3D keypoint correspondences or pre-registration, we focus on the difficult case of little to no annotation. To this end, we make the following contributions: (i) we develop a **Foot Implicit Neural Deformation** field model, named FIND, capable of tailoring explicit meshes at any resolution *i.e.* for low or high powered devices; (ii) an approach for training our model in various modes of weak supervision with progressively better disentanglement as more labels, such as pose categories, are provided; (iii) a novel unsupervised part-based loss for fitting our model to 2D images which is better than traditional photometric or silhouette losses; (iv) finally, we release a new dataset of high resolution 3D human foot scans, *Foot3D*. On this dataset, we show our model outperforms a strong PCA implementation trained on the same data in terms of shape quality and part correspondences, and that our novel unsupervised part-based loss improves inference on images.

## 1 Introduction

Shape reconstruction from single-view or few-view images is a vital but difficult computer vision task. Current approaches are dependent on strong priors to help constrain this ill-posed problem, usually in the form of well constructed and parameterised 3D models [22, 21]. The process of model creation is also often time consuming, requiring a lot of supervision *e.g.* registration and correspondence annotation. For the case of human feet, the current state-of-the-art in reconstruction involves the use of expensive scanning equipment [11, 8, 9], unavailable to the average consumer at home. Thus, there is growing interest for easier methods of foot reconstruction, particularly for home health monitoring, custom orthotics and the growing online shoe retail industry. As such, these models must be capable of running on low powered devices such as mobile phones, but also displayable at high resolution *e.g.* for a doctor’s inspection. To this end, we develop FIND, a **Foot Implicit Neural Deformation** field

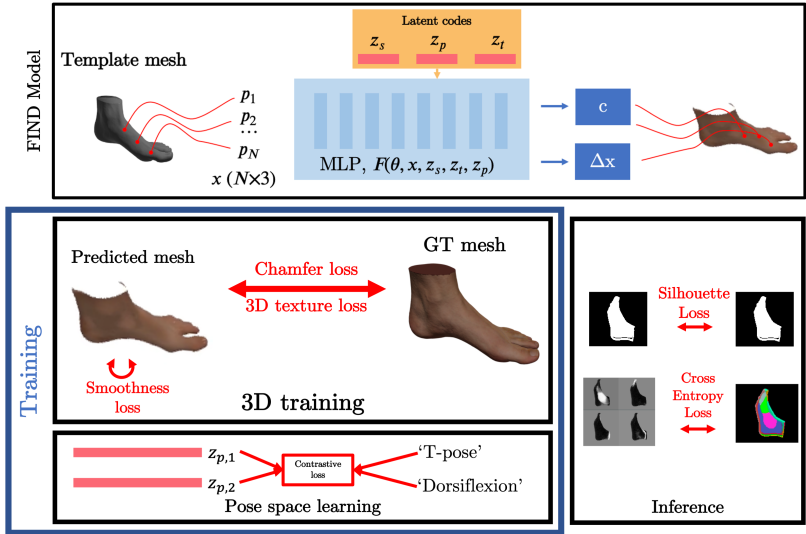


Figure 1: We produce (1) a coordinate based multilayer perceptron (MLP) architecture, FIND, which receives a 3D vertex position on a template mesh, and a shape, pose and texture embedding, and predicts a corresponding vertex deformation and colour. (2) We train our model with 3D based losses against ground truth scans, and learn a sensible pose space using a contrastive loss (3) We use differentiable rendering at inference to enforce a silhouette based loss, in addition to our novel unsupervised part-based loss.

model, which is easy to train as it requires little to no labels and can produce explicit meshes at tailored resolutions, useful for targeting specific hardware or needs. We can improve the model further by introducing minimal extra supervision: (i) we add the weak label of foot identities to disentangle shape and pose; and (ii) pose descriptions to produce a more interpretable pose latent space.

We outline our key contributions as follows: (1) a novel implicit foot model, FIND, capable of producing high fidelity human feet and fully parameterised by a disentangled latent space of shape, texture and pose. The model can also be tuned to specific hardware considerations or needs; (2) an approach to train the model under various levels of weak supervision, with stronger supervision producing better disentanglement; (3) a novel, unsupervised part-based loss for fitting FIND to images using unsupervised feature learning; and (4) we release *Foot3D*, a dataset of high resolution foot scans, providing shape and texture information for training of 3D foot models.

## 2 Related work

**Generative shape models.** These models endeavour to capture the distribution of an object category’s shape in a set of controllable parameters *e.g.* height, width, pose, etc. Much of the work in generative shape modelling looks at human bodies where a number of parameterised models have been built [27, 40]. For construction, the models are typically trained with strong supervision via 2D annotation [4, 20] of rich 3D scans [13, 35]. The building of generative models of more niche object categories often uses similar principles, but usually

on much smaller, custom datasets, often due to the time and expense in collecting the data. For examples of the difficulties involved, analogous generative quadruped models have been constructed from scans of toys of real animals [41] (as real animals don't keep still long enough), and parametric hand models have been constructed from 3D scans [62] and even at larger expense from MRI scans [49]. These models typically use Principal Component Analysis (PCA) to compose a small set of parameters which control linear offsets from a template mesh. We show here, that our implicit surface deformation field performs better against this type of strong PCA baseline on feet.

**Implicit shape representations.** Recent interest has shifted towards implicit representations of shape - functions, often deep networks, that take in coordinates corresponding to a point in space, and return spacial information relating to an object - such as whether the point in space is inside the object [8], or how near that point is to the object's surface [29]. These methods are naturally differentiable and can represent scenes at arbitrary scales, so lend themselves to many reconstruction tasks - for example, reconstructing meshes from point clouds [11] or single view images [36]. One particular method, Neural Radiance Fields (NeRFs) [25], has become increasingly popular due to the very high fidelity in which it can synthesis novel views of a scene. Many extensions to NeRF have investigated how to parameterise shape and texture for generic object classes [14], decompose scenes into individual, controllable components [22], and directly make edits to geometry and texture [20]. This has also been extended to posed humans, with several works [34, 37] showing promise in modelling dynamic human models using NeRF architectures. We take inspiration from this approach and, similar to [24], build our implicit foot model FIND by sampling points on the surface of a manifold, which in our case are vertices of a high resolution template foot mesh, and learn a parameterised deformation field for deforming vertices to different types of foot shapes and poses.

**Unsupervised representation learning for 3D objects.** Generative Adversarial Networks (GANs) [10] have shown promise in learning unsupervised representations of 2D data, even learning information about the underlying 3D geometry [28]. The intermediate feature representations are sufficiently powerful that Zhang *et al.* [39] showed they can be trained to solve downstream tasks, such as semantic segmentations, with very simple MLP classifiers and an incredibly small amount of ground truth labels. We leverage this for our novel unsupervised part-based loss, allowing us to learn foot part classes in 3D (from 2D images) to aid fitting to 2D images.

**Foot reconstruction.** Solutions exist for high resolution, accurate 3D scanners [1, 3, 4]. These are expensive and difficult to use outside of specific environments, so do not lend themselves to the average consumer. While PCA parameterised foot models do exist [6], many existing foot reconstruction solutions instead directly predict an unconstrained mesh from pointclouds or depth maps [23]. Large scale, proprietary datasets of feet, scanned at low resolution [15, 16] only release the population measurement statistics to the public. For our work, we build a new foot scan dataset, *Foot3D*, for model building and evaluation, which we release to the community.




Figure 2: 5 feet from our dataset and corresponding pose descriptions. The pose types are based on foot articulation described in foot anatomy literature [10] - further details in the supplementary




Figure 3: Our implicit surface field is an MLP which takes as input a 3D point queried on a template mesh’s surface, and texture, shape and pose embeddings, and provides as output a vertex colour value and displacement.

## 3 Method

### 3.1 Foot3D Dataset

We produce *Foot3D*, a dataset of high resolution, textured 3D human feet. For acquisition, we use an Artec Leo 3D scanner [11], which has a 3D point accuracy of up to 0.1mm. A total of 61 scans of the left feet on 34 subjects in a variety of poses was collect. To capture the entire surface of the foot, subjects would sit with their leg on a table for stability, and their foot suspended over the edge, this allowed the scanner to view the entire foot surface. Subjects then hold a static pose for approximately 2 minutes while the scan takes place. Details of the nature of the articulation requested can be found in the supplementary.

The raw data is then processed in Artec Studio 16 [12] to produce 100K polygon meshes. These meshes were cut-off at an approximate position on the shin, by a plane approximately perpendicular to the leg. Next, we leverage this slice-plane as a basis for loosely registering all of the meshes in the dataset, such that these planes are parallel to the XY-plane and a vector from the slice plane to the foot’s centroid lies in the X direction. We then slice each foot at a uniform height above the heel to provide a consistent ankle length. Figure 3.1 shows dataset samples, and further details about this process are provided in the supplementary.

### 3.2 The FIND model

The FIND model is controlled using a shape, pose and texture embedding to produce a surface deformation and colour field over a template mesh. This function is implemented as an implicit coordinate based neural network. The architecture is now explained in detail.



Figure 4: We show the effect of slicing and masking on our dataset and our differentiable renderer.

**Positional Encoding.** Texture and shape have both high frequency and low frequency information across a foot’s surface. As a result, we leverage positional encoding to allow for the network to capture the higher frequency signals [29, 80],

$$\gamma_N(x) = \{x, \sin(2\pi x), \cos(2\pi x), \dots, \sin(2N\pi x), \cos(2N\pi x)\} \quad (1)$$

**Network.** After we encode our position with  $N = 10$ , we learn a network  $F$  with weights  $\theta$ , which takes shape, pose and texture encodings  $z_s$ ,  $z_p$  and  $z_t$ , and predicts vertex offsets  $\Delta x$  relative to a template mesh and vertex colours  $c$ ,

$$F(\theta, \gamma(x), z_s, z_p, z_t) \rightarrow (\Delta x, c) \quad (2)$$

The main body of the network is 4  $256 \times 256$  linear layers, encoding the positional encoding  $\gamma(x)$  into a feature,  $g$ .

**Heads.** We concatenate  $g$  with the relevant shape codes and pass to a head specific for each task. The colour head receives  $g$  and the texture code  $z_t$ , and the displacement head receives  $g$  and the shape and pose codes,  $z_s$  and  $z_p$ . Both heads are composed of 3  $256 \times 256$  linear layers. We choose tanh as our activation function, and normalize outputs such that colour values are in the range  $[0, 1]$ , and displacements in the range  $[-0.1, 0.1]$ .

**Defining the surface.** We define the surface over which the network is trained by using a template mesh, which we take as a foot in our dataset outside of our training and validation sets. The vertex positions  $x$  are found by sampling across the surface of a template mesh. At inference time, all of these vertices can be used to output a full mesh.

**Latent vectors.** The shape, texture and pose encodings, all of size 100, are jointly learned during training. We place the constraint that texture and shape codes must be shared between different scans of the same foot. We also learn a per-foot registration to the ground truth: an Euler rotation  $r \in \mathbb{R}^3$ , global translation  $t \in \mathbb{R}^3$ , and global XYZ scaling  $s \in \mathbb{R}^3$ .

### 3.3 Learning part-segmentation

Here we aim to develop an inference-time supervisory loss that learns part-awareness in an unsupervised manner.

**Differentiable rendering.** In order to optimise the model to fit to real images, it is necessary to be able to differentiably render the model. We use PyTorch3D’s [81] differentiable renderer. As feet are not complete, watertight objects on their own, we choose to mask out the slice plane applied to the feet models when rendering, as we argue this better reflects how real feet appear in images (post-segmentation). To achieve this, all faces on this slice



Figure 5: 20 clusters on the 8th feature map of our Restyle encoder.

plane are marked when formulating the dataset, and these faces are masked out of the renders produced by the differentiable renderer. Figure 4 shows the effect of this process.

**Extracting features.** Using the knowledge that 2D GANs can learn 3D geometrical information [28], we train a StyleGAN2 [17] network to generate rendered feet from StyleGAN2 latent codes. We train on 100k 128x128 rendered images from models within Foot3D, with camera viewpoints in a hemisphere above the foot. To use this network as an inference loss, invert StyleGAN to obtain the codes using the Restyle [5] encoder. When passing a rendered image from our differentiable renderer through the encoder-decoder, similar to [39], we can use the features maps after the AdaIN layers of StyleGAN2 as a form of high level information about the foot in the image.

**Learning unsupervised 2D parts.** To learn foot parts from these feature maps, we observe that clustering them with k-means produces meaningful part segmentations, as can be seen in Figure 5, which correspond between different feet. As a result, we train a simple linear classifier to produce these per-pixel part segmentations from a subset of the StyleGAN feature maps. The whole classification process is fully differentiable.

**Moving to 3D.** We then use this classifier to learn a per-vertex part probability on the template mesh *i.e.* a  $C$  length vector  $c \in [0, 1]$  where  $C$  is the total number of classes. These vectors can be directly ‘rendered’ using our differentiable renderer to any new  $C$  channel image producing a per pixel foot part probability map.

**Unsupervised part-based loss.** Using this part classification pipeline, we enforce a cross-entropy loss on our  $C$  cluster classes when optimising to images. Given our rendered part probability map  $P \in \mathbb{R}^{H \times W \times C}$ , and our ‘ground-truth’ labels from our classifier  $L \in \mathbb{R}^{H \times W}$ , we define our loss,

$$Loss_{SCE} = -\frac{1}{H \times W} \sum_{y=1}^H \sum_{x=1}^W \sum_{c=1}^C \ln \frac{\exp P_{y,x,L_{y,x}}}{\sum_{i=1}^C \exp P_{y,x,i}} \quad (3)$$

### 3.4 Training Losses

Our training loop and inference process uses several different losses:

**Chamfer loss.** We sample 5000 vertices uniformly from the surface of each of the ground truth and predicted meshes, to get sampled vertices  $v_{\text{pred}}$  and  $v_{\text{gt}}$ , and evaluate,

$$L_{\text{chamf}} = \sum_j \min_i \|v_{\text{gt},i} - v_{\text{pred},j}\|_2 + \sum_i \min_j \|v_{\text{gt},i} - v_{\text{pred},j}\|_2 \quad (4)$$

**Smoothness loss.** We enforce a loss on the smoothness of the predicted mesh, which is a combination of a Laplacian smoothness [26] metric, and an average edge length regularizer.

**Texture loss.** We sample 1000 vertices  $v_{\text{gt}}$  uniformly over the ground truth mesh surface, collect their colours  $c_{\text{gt}}$ , and enforce a loss encouraging the network to map  $v_{\text{gt}}$  to  $c_{\text{gt}}$ ,

$$L_{\text{tex}} = \|c_{\text{gt}} - F(v_{\text{gt}}, z_s, z_p, z_t)_{\text{col}}\|_2 \quad (5)$$

**Contrastive loss.** To encourage the latent pose space to meaningfully partition the data, we enforce a contrastive loss between selected pairs of pose latent codes  $z_p$ . Given the latent vectors  $z_{p,1}$  and  $z_{p,2}$ , and corresponding ‘label’ vectors  $l_1, l_2$ , we enforce the loss,

$$L_{\text{contr}} = yd^2 + (1 - y) \max(M - d, 0)^2, \quad \text{where } d = \|z_{p,1} - z_{p,2}\|_2, y = l_1 \cdot l_2 \quad (6)$$

where  $M = 0.5$  is a selected margin distance to encourage between dissimilar pose vectors. Details of the formation of the labels  $l$  can be found in the supplementary.

### 3.5 Training

Our full training loop consists of three stages: (1) registration, in which we optimise per-foot registration parameters to align our template mesh to the *Foot3D* dataset; (2) network training, in which we optimise network weights  $\theta$  and latent codes  $z_s, z_p$ , and  $z_t$  in accordance with our training losses; and (3) latent refinement, where we fix network weights  $\theta$  and fully optimise the latent codes to best match the data. In this loop, we minimise:

$$\begin{aligned} E_{\text{registration}}(r, t, s) &= \lambda_{\text{chamf}} L_{\text{chamf}} \\ E_{\text{network training}}(\theta, z_s, z_p, z_t) &= \lambda_{\text{chamf}} L_{\text{chamf}} + \lambda_{\text{smooth}} L_{\text{smooth}} + \lambda_{\text{tex}} L_{\text{tex}} \\ E_{\text{latent refinement}}(z_s, z_p, z_t) &= \lambda_{\text{chamf}} L_{\text{chamf}} + \lambda_{\text{smooth}} L_{\text{smooth}} + \lambda_{\text{tex}} L_{\text{tex}} \end{aligned} \quad (7)$$

We select  $\lambda_{\text{chamf}} = 10^4$ ,  $\lambda_{\text{smooth}} = 10^3$ ,  $\lambda_{\text{tex}} = 1$ , and train our network and latent parameters with an Adam optimizer [18], with a learning rate of  $5 \times 10^{-5}$ .

**Levels of supervision.** We are able to train our model with three levels of supervision, to provide better pose disentanglement and control: (i) Fully unsupervised, training only from mesh data; (ii) Using foot identity as a weak label, fixing shape between feet of the same identity for disentanglement of shape and pose; (iii) Adding pose descriptions as an additional weak label to allow for a more interpretable pose space via our contrastive loss.

## 4 Experiments

**Foot3D dataset.** Our training set contains 27 feet in which the users were instructed to hold a neutral, ‘T-pose’ position, and 51 feet total, including articulated feet. We use a further 8 feet for validation, of which 4 are ‘T-pose’.

**Evaluation metrics.** We evaluate the shape accuracy of our models using three metrics: a mean chamfer loss between a predicted model and the ground truth (as in Equation 4); a mean euclidean distance between 6 hand-labelled keypoints on the model’s template mesh and the ground truth meshes; and an Intersection over Union (IoU) metric on silhouettes produced from 50 rendered viewpoints sampled uniformly in an arc across the foot.




Figure 6: Our network can be sampled at different resolutions for lower-resolution applications (*e.g.* mobile phones), or higher-resolution applications. Query times are shown on a Macbook M1 CPU.



Figure 7: Here we show 5 samples within the latent spaces of shape, pose and texture. Shape variations visible are toe lengths, separation and arch of the foot; pose variations visible are toe flexion and foot plantarflexion.

**Baseline.** In order to compare our method with typical PCA implementations, we generate a PCA model. To do this without dense correspondences, we first use a FoldingNet [38] implementation fitted to all of our data to produce a set of 1600-vertex point clouds with learned weak correspondences, and fit a PCA model to this data.

## 5 Results

**Multi-resolution.** As a coordinate based model, our neural surface deformation field can be sampled at as many 3D vertices as desired to produce an output mesh. As visualised in Figure 6, we leverage this to sample at high resolutions for precise applications, or at lower resolutions for memory and CPU-critical applications, such as mobile phones. This requires template mesh to be reproduced at different resolutions, possible through off-the-shelf mesh simplification and subdivision tools [9].






Figure 8: We visualise the FIND pose space on all training data in its two principal directions. We apply k-means clustering with 4 clusters to show partitioning of the space. We identify feet with poses in one of four pose categories to showcase the effective pose space partitioning due to our contrastive loss.

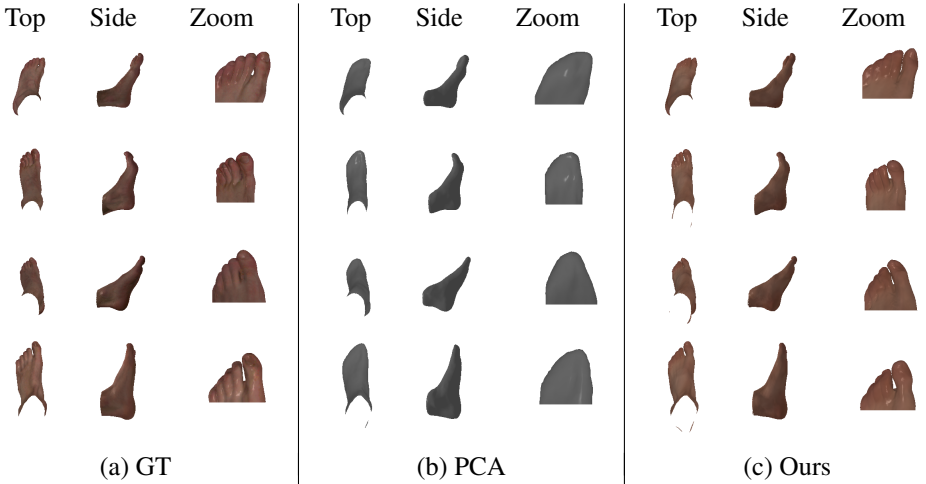


Figure 9: Qualitative results of our 3D fits to 8 validation feet, rendered from 3 views each.

Model	Trained on	Chamfer, $\mu\text{m} \downarrow$	Keypoint, $\text{mm} \downarrow$	IoU $\uparrow$
PCA	T-Pose only	12.9	15.0	0.892
PCA	All	11.7	15.7	0.892
Ours	T-pose only	9.0	6.4	0.923
Ours	All	<b>7.3</b>	<b>5.9</b>	<b>0.931</b>

Table 1: Comparison of FIND and a PCA baseline on our 8 validation feet.

**Latent spaces.** Visualised learned shape, pose, and texture spaces of FIND are shown in Figure 7. Observe that our model has learned to disentangle shape and pose with minimal supervision using the simple constraint of fixing shape and texture codes between scans of feet belonging to the same person. Furthermore, we show in Figure 8 that our contrastive loss ensures pose space is suitably segmented according to pose type.

**3D evaluation.** We compare our model to the PCA baseline. For both, we train 2 models: one on the 27 ‘T-pose’ feet in the training dataset, and one on all 51 feet. Our ‘T-pose’ trained model is restricted from learning a pose space. Table 1 shows a significant quantitative improvement over the baseline, and Figure 9 shows qualitative results. Notice our model does a much better job of toe separation.

Optimisation loss	2 view		5 view	
	Chamfer, $\mu\text{m} \downarrow$	Keypoint, $\text{mm} \downarrow$	Chamfer, $\mu\text{m} \downarrow$	Keypoint, $\text{mm} \downarrow$
Sil	81.8	14.4	16.8	7.7
Sil + VGG	78.7	13.1	15.9	7.3
Sil + CE Loss	<b>45.8</b>	<b>10.3</b>	<b>15.7</b>	<b>6.4</b>

Table 2: Quantitative results on inference pipeline on our validation feet.

Model	Chamfer, $\mu\text{m} \downarrow$	Keypoint, $\text{mm} \downarrow$	IoU $\uparrow$
Full model	<b>7.3</b>	<b>5.9</b>	<b>0.931</b>
– Pose vector	8.0	6.2	0.926
– Texture loss	7.6	6.8	0.928

Table 3: Ablation study.

**Inference evaluation.** We show in Table 2 our results of optimising FIND to renderings of validation feet, using only 2D losses to fit the registration and latent parameters. In our experiments, we sample 2 and 5 viewpoints in an arc around the foot. We compare three training losses: silhouette only; silhouette + an off-the-shelf VGG-16 [53] perceptual loss; and silhouette + our novel Cross Entropy (CE) unsupervised part-based loss. We show that our loss improves the quality of these reconstructions. We find that our loss has a more substantial improvement in the recovered shape when fewer views are available for optimisation - we see this as likely being due to the extra 3D information this loss has available about unseen viewpoints.

## 5.1 Ablation study

We investigate in Table 3 the effect on our full model of several features. We observe that the shape and pose disentanglement improves the shape reconstruction quality. Additionally, the texture loss, which is only able to learn from vertex colours, is also responsible for some of the reconstruction quality. We predict that this might be due to the surface colours providing supervision for vertex correspondences.

## 6 Conclusion

We demonstrate that our neural implicit foot model, FIND, outperforms typical PCA implementations of statistical shape learning, providing a higher fidelity shape space, whilst also able to learn texture and pose with minimal supervision. Furthermore, we have shown that simple foot identity and pose labels can be used to disentangle pose and shape, and produce a meaningful, interpretable pose space. Finally, we have shown that our unsupervised part-based loss improves upon typical shape model inference methods by enforcing part consistency.

## References

- [1] Artec Leo, Artec3D. <https://www.artec3d.com/portable-3d-scanners/artec-leo>, . Accessed: 2022-07-24.
- [2] Artec Studio. <https://www.artec3d.com/3d-software/artec-studio>, . Accessed: 2022-07-25.
- [3] Volumental foot scanner. <https://volumental.com/volumental-retail-scanner>. Accessed: 2022-07-24.
- [4] Vorum YETI 3D foot scanner. <https://vorum.com/yeti-3d-foot-scanner/>. Accessed: 2022-07-24.
- [5] Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. ReStyle: A residual-based Style-GAN encoder via iterative refinement. In *ICCV*, 2021.
- [6] Edmée Amstutz, Tomoaki Teshima, Makoto Kimura, Masaaki Mochimaru, and Hideo Saito. PCA based 3D shape reconstruction of human foot using multiple viewpoint cameras. In *ICVS*, 2008.
- [7] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2D human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014.
- [8] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019.
- [9] Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. MeshLab: an Open-Source Mesh Processing Tool. In Vittorio Scarano, Rosario De Chiara, and Ugo Erra, editors, *Eurographics Italian Chapter Conference*. The Eurographics Association, 2008. ISBN 978-3-905673-68-5. doi: 10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136.
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Y. Bengio. Generative adversarial networks. 2014.
- [11] Rana Hanocka, Gal Metzer, Raja Giryes, and Daniel Cohen-Or. Point2Mesh: A self-prior for deformable meshes. *arXiv preprint arXiv:2005.11084*, 2020.
- [12] Peggy A Houglum and Dolores B Bertoti. *Brunnstrom’s clinical kinesiology*. FA Davis, 2011.
- [13] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7): 1325–1339, jul 2014.
- [14] Wonbong Jang and Lourdes Agapito. CodeNeRF: Disentangled neural radiance fields for object categories. In *ICCV*, 2021.
- [15] Ales Jurca, Tomaz Kolsek, and Tina Vidic. Dorothy mass foot measurement campaign. *Footwear Science*, 3(sup1):S83–S85, 2011.

- [16] Ales Jurca, Jure Žabkar, and Sašo Džeroski. Analysis of 1.2 million foot scans from North America, Europe and Asia. *Scientific reports*, 9(1):1–10, 2019.
- [17] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020.
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [19] Yuwei Li, Minye Wu, Yuyao Zhang, Lan Xu, and Jingyi Yu. PIANO: A parametric hand bone model from magnetic resonance imaging. *arXiv preprint arXiv:2106.10893*, 2021.
- [20] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014.
- [21] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *ICCV*, 2021.
- [22] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *TOG*, 2015.
- [23] Nolan Lunscher and John Zelek. Point cloud completion of foot shape from a single depth map for fit matching using deep learning view synthesis. In *ICCV Workshops*, pages 2300–2305, 2017.
- [24] Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. Text2Mesh: Text-driven neural stylization for meshes. In *CVPR*, 2022.
- [25] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- [26] Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. Laplacian mesh optimization. In *GRAPHITE*, 2006.
- [27] Michael Niemeyer and Andreas Geiger. GIRAFFE: Representing scenes as compositional generative neural feature fields. In *CVPR*, 2021.
- [28] Xingang Pan, Bo Dai, Ziwei Liu, Chen Change Loy, and Ping Luo. Do 2D gans know 3D shape? unsupervised 3D shape reconstruction from 2d image gans. *arXiv preprint arXiv:2011.00844*, 2020.
- [29] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019.
- [30] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *ICML*, 2019.

- [31] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3D deep learning with PyTorch3D. *arXiv:2007.08501*, 2020.
- [32] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *arXiv preprint arXiv:2201.02610*, 2022.
- [33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [34] Shih-Yang Su, Frank Yu, Michael Zollhoefer, and Helge Rhodin. A-NeRF: Surface-free human 3D pose refinement via neural rendering. *arXiv preprint arXiv:2102.06199*, 2021.
- [35] Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3D human pose in the wild using IMUs and a moving camera. In *ECCV*, 2018.
- [36] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2Mesh: Generating 3D mesh models from single RGB images. In *ECCV*, 2018.
- [37] Hongyi Xu, Thiemo Alldieck, and Cristian Sminchisescu. H-NeRF: Neural radiance fields for rendering and temporal reconstruction of humans in motion. *NeurIPS*, 2021.
- [38] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. FoldingNet: Point cloud auto-encoder via deep grid deformation. In *CVPR*, 2018.
- [39] Yuxuan Zhang, Huan Ling, Jun Gao, Kangxue Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, and Sanja Fidler. DatasetGAN: Efficient labeled data factory with minimal human effort. In *CVPR*, 2021.
- [40] Silvia Zuffi and Michael J Black. The stitched puppet: A graphical model of 3D human shape and pose. In *CVPR*, 2015.
- [41] Silvia Zuffi, Angjoo Kanazawa, David W Jacobs, and Michael J Black. 3D menagerie: Modeling the 3D shape and pose of animals. In *CVPR*, 2017.