

# FreiHAND: A Dataset for Markerless Capture of Hand Pose and Shape from Single RGB Images

Christian Zimmermann<sup>1</sup>, Duygu Ceylan<sup>2</sup>, Jimei Yang<sup>2</sup>, Bryan Russell<sup>2</sup>,  
 Max Argus<sup>1</sup>, and Thomas Brox<sup>1</sup>

<sup>1</sup>University of Freiburg

<sup>2</sup>Adobe Research

Project page: <https://lmb.informatik.uni-freiburg.de/projects/freihand/>

## Abstract

Estimating 3D hand pose from single RGB images is a highly ambiguous problem that relies on an unbiased training dataset. In this paper, we analyze cross-dataset generalization when training on existing datasets. We find that approaches perform well on the datasets they are trained on, but do not generalize to other datasets or in-the-wild scenarios. As a consequence, we introduce the first large-scale, multi-view hand dataset that is accompanied by both 3D hand pose and shape annotations. For annotating this real-world dataset, we propose an iterative, semi-automated ‘human-in-the-loop’ approach, which includes hand fitting optimization to infer both the 3D pose and shape for each sample. We show that methods trained on our dataset consistently perform well when tested on other datasets. Moreover, the dataset allows us to train a network that predicts the full articulated hand shape from a single RGB image. The evaluation set can serve as a benchmark for articulated hand shape estimation.

## 1. Introduction

3D hand pose and shape estimation from a single RGB image has a variety of applications in gesture recognition, robotics, and AR. Various deep learning methods have approached this problem, but the quality of their results depends on the availability of training data. Such data is created either by rendering synthetic datasets [4, 6, 19, 20, 33] or by capturing real datasets under controlled settings typically with little variation [7, 22, 27]. Both approaches have limitations, discussed in our related work section.

Synthetic datasets use deformable hand models with texture information and render this model under varying pose configurations. As with all rendered datasets, it is difficult to model the wide set of characteristics of real images, such as varying illumination, camera lens distortion, motion blur, depth of field and debayering. Even more importantly, ren-




Figure 1: We create a hand dataset via a novel iterative procedure that utilizes multiple views and sparse annotation followed by verification. This results in a large scale real world data set with pose and shape labels, which can be used to train single-view networks that have superior cross-dataset generalization performance on pose and shape estimation.

dering of hands requires samples from the true distribution of feasible and realistic hand poses. In contrast to human pose, such distributional data does not exist to the same extent. Consequently, synthetic datasets are either limited in the variety of poses or sample many unrealistic poses.

Capturing a dataset of real human hands requires annotation in a post-processing stage. In single images, manual annotation is difficult and cannot be easily crowd sourced due to occlusions and ambiguities. Moreover, collecting and annotating a large scale dataset is a respectable effort.

In this paper, we analyze how these limitations affect the ability of single-view hand pose estimation to generalize across datasets and to *in-the-wild* real application scenarios. We find that datasets show excellent performance on the respective evaluation split, but have rather poor performance on other datasets, i.e., we see a classical dataset bias.

As a remedy to the dataset bias problem, we created a new large-scale dataset by increasing variation between samples. We collect a real-world dataset and develop a methodology that allows us to automate large parts of the labeling procedure, while manually ensuring very high-fidelity annotations of 3D pose and 3D hand shape. One of



Figure 2: Examples from our proposed dataset showing images (top row) and hand shape annotations (bottom row). The training set contains composited images from green screen recordings, whereas the evaluation set contains images recorded indoors and outdoors. The dataset features several subjects as well as object interactions.

the key aspects is that we record synchronized images from multiple views, an idea already used previously in [2, 22]. The multiple views remove many ambiguities and ease both the manual annotation and automated fitting. The second key aspect of our approach is a semi-automated *human-in-the-loop* labeling procedure with a strong bootstrapping component. Starting from a sparse set of 2D keypoint annotations (e.g., finger tip annotations) and semi-automatically generated segmentation masks, we propose a hand fitting method that fits a deformable hand model [21] to a set of multi-view input. This fitting yields both 3D hand pose and shape annotation for each view. We then train a multi-view 3D hand pose estimation network using these annotations. This network predicts the 3D hand pose for unlabeled samples in our dataset along with a confidence measure. By verifying confident predictions and annotating least-confident samples in an iterative procedure, we acquire 11592 annotations with moderate manual effort by a human annotator.

The dataset spans 32 different people and features fully articulated hand shapes, a high variation in hand poses and also includes interaction with objects. Part of the dataset, which we mark as training set, is captured against a green screen. Thus, samples can easily be composed with varying background images. The test set consists of recordings in different indoor and outdoor environments; see Figure 2 for sample images and the corresponding annotation.

Training on this dataset clearly improves cross-dataset generalization compared to training on existing datasets. Moreover, we are able to train a network for full 3D hand shape estimation from a single RGB image. For this task, there is not yet any publicly available data, neither for training nor for benchmarking. Our dataset is available on our project page and therefore can serve both as training and benchmarking dataset for future research in this field.

## 2. Related Work

Since datasets are crucial for the success of 3D hand pose and shape estimation, there has been much effort on acquiring such data.

In the context of hand shape estimation, the majority of methods fall into the category of model-based techniques. These approaches were developed in a strictly controlled environment and utilize either depth data directly [24, 25, 28] or use multi-view stereo methods for reconstruction [2]. More related to our work are approaches that fit statistical human shape models to observations [3, 17] from *in-the-wild* color images as input. Such methods require semi-automatic methods to acquire annotations such as keypoints or segmentation masks for each input image to guide the fitting process.

Historically, acquisition methods often incorporated markers onto the hand that allow for an easy way to estimate pose afterwards. Common choices are infrared markers [9], color coded gloves [29], or electrical sensing equipment [32]. This alters hand appearance and, hence, makes the data less valuable for training discriminative methods.

Annotations can also be provided manually on hand images [20, 23, 30]. However, the annotation is limited to visible regions of the hand. Thus, either the subject is required to retain from complex hand poses that result in severe self-occlusions, or only a subset of hand joints can be annotated.

To avoid occlusions and annotate data at larger scale, Simon *et al.* [22] leveraged a multi-view recording setup. They proposed an iterative bootstrapping approach to detect hand keypoints in each view and triangulate them to generate 3D point hypotheses. While the spirit of our data collection strategy is similar, we directly incorporate the multi-view information into a neural network for predicting 3D keypoints and our dataset consists of both pose and shape annotations.

eval train	STB	RHD	GAN	PAN	LSMV	FPA	HO-3D	Ours	Average Rank
STB [30]	<b>0.783</b>	0.179	0.067	0.141	0.072	0.061	0.138	0.138	6.0
RHD [33]	0.362	<b>0.767</b>	<b>0.184</b>	<b>0.463</b>	<b>0.544</b>	0.101	<b>0.450</b>	<b>0.508</b>	2.9
GAN [19]	0.110	0.103	<b>0.765</b>	0.092	0.206	<b>0.180</b>	0.087	0.183	5.4
PAN [11]	<b>0.459</b>	<b>0.316</b>	0.136	<b>0.870</b>	0.320	<b>0.184</b>	<b>0.351</b>	<b>0.407</b>	3.0
LSMV [7]	0.086	0.209	0.152	0.189	<b>0.717</b>	0.129	0.251	0.276	4.1
FPA [5]	0.119	0.095	0.084	0.120	0.118	<b>0.777</b>	0.106	0.163	6.0
HO-3D [8]	0.154	0.130	0.091	0.111	0.149	0.073	-	0.169	6.1
<b>Ours</b>	<b>0.473</b>	<b>0.518</b>	<b>0.217</b>	<b>0.562</b>	<b>0.537</b>	0.128	<b>0.557</b>	<b>0.678</b>	2.2

Table 1: This table shows cross-dataset generalization measured as area under the curve (AUC) of percentage of correct keypoints following [33]. Each row represents the training set used and each column the evaluation set. The last column shows the average rank each training set achieved across the different evaluation sets. The top-three ranking training sets for each evaluation set are marked as follows: **first**, **second** or **third**. Note that the evaluation set of *HO-3D* was not available at time of submission, therefore one table entry is missing and the other entries within the respective column report numbers calculated on the training set.

Since capturing real data comes with an expensive annotation setup and process, more methods rather deployed synthetic datasets recently [20, 33].

### 3. Analysis of Existing Datasets

We thoroughly analyze state-of-the-art datasets used for 3D hand pose estimation from single RGB images by testing their ability to generalize to unseen data. We identify seven state-of-the-art datasets that provide samples in the form of an RGB image and the accompanying 3D keypoint information as shown in Table 2.

#### 3.1. Considered Datasets

**Stereo Tracking Benchmark (STB)** [30] dataset is one of the first and most commonly used datasets to report performance of 3D keypoint estimation from a single RGB image. The annotations are acquired manually limiting the setup to hand poses where most regions of the hands are visible. Thus, the dataset shows a unique subject posing in a frontal pose with different background scenarios and without objects.

The **Panoptic (PAN)** dataset [11] was created using a dense multi-view capture setup consisting of 10 RGB-D sensors, 480 VGA and 31 HD cameras. It shows humans performing different tasks and interacting with each other. There are 83 sequences publicly available and 12 of them have hand annotation. We select *171204\_pose3* to serve as evaluation set and use the remaining 11 sequences from the *range motion*, *haggling* and *tools* categories for training.

Garcia *et al.* [5] proposed the **First-person hand action benchmark (FPA)**, a large dataset that is recorded from an egocentric perspective and annotated using magnetic sensors attached to the finger tips of the subjects. Wires run along the fingers of the subject altering the appearance of the hands significantly. 6 DOF sensor measurements are uti-

lized in an inverse kinematics optimization of a given hand model to acquire the full hand pose annotations.

Using the commercial Leap Motion device [18] for keypoint annotation, Gomez *et al.* [7] proposed the **Large-scale Multiview 3D Hand Pose Dataset (LSMV)**. Annotations given by the device are transformed into 4 calibrated cameras that are approximately time synchronized. Due to the limitations of the sensor device, this dataset does not show any hand-object interactions.

The **Rendered Hand Pose Dataset (RHD)** proposed by Zimmermann *et al.* [33] is a synthetic dataset rendered from 20 characters performing 31 different actions in front of a random background image without hand object interaction.

Building on the SynthHands [20] dataset Mueller *et al.* [19] presented the **GANerated (GAN)** dataset. SynthHands was created by retargeting measured human hand articulation to a rigged meshed model in a mixed reality approach. This allowed for hand object interaction to some extent, because the subject could see the rendered scene in real time and pose the hand accordingly. In the following *GANerated* hand dataset, a CycleGAN approach is used to bridge the synthetic to real domain shift.

Recently, Hampali *et al.* [8] proposed an algorithm for dataset creation deploying an elaborate optimization scheme incorporating temporal and physical consistencies, as well as silhouette and depth information. The resulting dataset is referred to as **HO-3D**.

#### 3.2. Evaluation Setup

We trained a state-of-the-art network architecture [10] that takes as input an RGB image and predicts 3D keypoints on the training split of each of the datasets and report its performance on the evaluation split of all other datasets. For each dataset, we either use the standard training/evaluation split reported by the authors or create an 80%/20% split

dataset	num. frames	num. subjects	real	obj- ects	shape	labels
STB [30]	15 k / 3 k	1	✓	✓	✗	manual
PAN [33]	641 k / 34 k	> 10	✓	✓	✗	MVBS [22]
FPA [5]	52 k / 53 k	6	✓	✓	✗	marker
LSMV [7]	117 k / 31 k	21	✓	✗	✗	leapmotion
RHD [33]	41 k / 2.7 k	20	✗	✗	✗	synthetic
GAN [19]	266 k / 66 k	-	✗	✓	✗	synthetic
HO-3D [8]	11 k / -	3	✓	✓	✓	automatic [8]
Ours	33 k / 4 k	32	✓	✓	✓	hybrid

Table 2: State-of-the-art datasets for the task of 3D keypoint estimation from a single color image used in our analysis. We report dataset size in number of frames, number of subjects, if it is real or rendered data, regarding hand object interaction, if shape annotation is provided and which method was used for label generation.

otherwise; see the supplementary material for more details.

The single-view network takes an RGB image  $\mathbf{I}$  as input and infers 3D hand pose  $\mathbf{P} = \{\mathbf{p}_k\}$  with each  $\mathbf{p}_k \in \mathbb{R}^3$ , representing a predefined landmark or keypoint situated on the kinematic skeleton of a human hand. Due to scale ambiguity, the problem to estimate real world 3D keypoint coordinates in a camera centered coordinate frame is ill-posed. Hence, we adopt the problem formulation of [10] to estimate coordinates in a root relative and scale normalized fashion:

$$\mathbf{p}_k = s \cdot \hat{\mathbf{p}}_k = s \cdot \begin{pmatrix} \hat{x}_k \\ \hat{y}_k \\ \hat{z}_k^{\text{rel}} \end{pmatrix} = \begin{pmatrix} \hat{x}_k \\ \hat{y}_k \\ \hat{z}_k^{\text{rel}} + \hat{z}^{\text{root}} \end{pmatrix}, \quad (1)$$

where the normalization factor  $s$  is chosen as the length of one reference bone in the hand skeleton,  $\hat{z}^{\text{root}}$  is the root depth and  $\hat{z}_k^{\text{rel}}$  the relative depth of keypoint  $k$ . We define the resulting 2.5D representation as:

$$\hat{\mathbf{p}}_{\text{rel},k} = (\hat{x}_k, \hat{y}_k, \hat{z}_k^{\text{rel}})^T. \quad (2)$$

Given scale constraints and 2D projections of the points in a calibrated camera, 3D hand pose  $\mathbf{P}$  can be recovered from  $\hat{\mathbf{P}}_{\text{rel}}$ . For details about this procedure we refer to [10].

We train the single-view network using the same hyperparameter choices as Iqbal *et al.* [10]. However, we use only a single stage and reduce the number of channels in the network layers, which leads to a significant speedup in terms of training time at only a marginal decrease in accuracy. We apply standard choices of data augmentation including color, scale and translation augmentation as well as rotation around the optical axis. We apply this augmentation to each of the datasets.

### 3.3. Results

It is expected that the network performs the best on the dataset it was trained on, yet it should also provide reasonable predictions for unseen data when being trained on a

dataset with sufficient variation (e.g., hand pose, viewpoint, shape, existence of objects, etc.).

Table 1 shows for each existing training dataset the network is able to generalize to the respective evaluation split and reaches the best results there. On the other hand, performance drops substantially when the network is tested on other datasets.

Both *GAN* and *FPA* datasets appear to be especially hard to generalize indicating that their data distribution is significantly different from the other datasets. For *FPA* this stems from the appearance change due to the markers used for annotation purposes. The altered appearance gives the network trained on this dataset strong cues to solve the task that are not present for other datasets at evaluation time. Thus, the network trained on *FPA* performs poorly when tested on other datasets. Based on visual inspection of the *GAN* dataset, we hypothesize that subtle changes like missing hand texture and different color distribution are the main reasons for generalization problems. We also observe that while the network trained on *STB* does not perform well on remaining datasets, the networks trained on other datasets show reasonable performance on the evaluation split of *STB*. We conclude that a good performance on *STB* is not a reliable measure for how a method generalizes to unseen data.

Based on the performance of each network, we compute a cumulative ranking score for each dataset that we report in the last column of Table 1. To calculate the cumulative rank we assign ranks for each column of the table separately according to the performance the respective training sets achieve. The cumulative rank is then calculated as average over all evaluation sets, *i.e.* rows of the table. Based on these observations, we conclude that there is a need for a new benchmarking dataset that can provide superior generalization capability.

We present the *FreiHAND* Dataset to achieve this goal. It consists of real images, provides sufficient viewpoint and hand pose variation, and shows samples both with and without object interactions. Consequently, the single-view network trained on this dataset achieves a substantial improvement in terms of ranking for cross-dataset generalization. We next describe how we acquired and annotated this dataset.

## 4. FreiHAND Dataset

The dataset was captured with the multi-view setup shown in Fig. 3. The setup is portable enabling both indoor and outdoor capture. We capture hand poses from 32 subjects of different genders and ethnic backgrounds. Each subject is asked to perform actions with and without objects. To capture hand-object interactions, subjects are given a number of everyday household items that allow for reasonable one-handed manipulation and are asked to





Figure 3: Recording setup with 8 calibrated and temporally synchronized RGB cameras located at the corners of a cube. A green screen background can be mounted into the setup, enabling easier background subtraction.

demonstrate different grasping techniques. More information is provided in the supplementary material.

To preserve the realistic appearance of hands, no markers are used during the capture. Instead we resort to post-processing methods that generate 3D labels. Manual acquisition of 3D annotations is obviously unfeasible. An alternative strategy is to acquire 2D keypoint annotations for each input view and utilize the multi-view camera setup to lift such annotations to 3D similar to Simon *et al.* [22].

We found after initial experiments that current 2D hand pose estimation methods perform poorly, especially in case of challenging hand poses with self- and object occlusions. Manually annotating all 2D keypoints for each view is prohibitively expensive for large-scale data collection. Annotating all 21 keypoints across multiple-views with a specialized tool takes about 15 minutes for each multi-view set. Furthermore, keypoint annotation alone is not sufficient to obtain shape information.

We address this problem with a novel bootstrapping procedure (see Fig. 4) composed of a set of automatic methods that utilize sparse 2D annotations. Since our data is captured against a green screen, the foreground can be extracted automatically. Refinement is needed only to co-align the segmentation mask with the hand model’s wrist. In addition, a sparse set of six 2D keypoints (finger tips and wrist) is manually annotated. These annotations are relatively cheap to acquire at a reasonably high quality. For example, manually correcting a segmentation mask takes on average 12 seconds, whereas annotating a keypoint takes around 2 seconds. Utilizing this information we fit a deformable hand model to multi-view images using a novel fitting process described in Section 4.1. This yields candidates for both 3D hand pose and shape labels. These can-

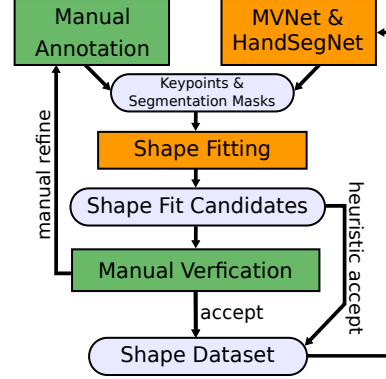


Figure 4: The dataset labeling workflow starts from manual annotation followed by the shape fitting process described in 4.1, which yields candidate shape fits for our data samples. Sample fits are manually verified allowing them to be accepted, rejected or queued for further annotation. Alternatively a heuristic can accept samples without human interaction. The initial dataset allows for training the networks involved, which for subsequent iterations of the procedure, can predict information needed for fitting. The labeling process can be bootstrapped, allowing more accepted samples to accumulate in the dataset.

didates are then manually verified, before being added to a set of labels.

Given an initial set of labels, we train our proposed network, *MVNet*, that takes as inputs multi-view images and predicts 3D keypoint locations along with a confidence score, described in Section 4.2. Keypoint predictions can be used in lieu of manually annotated keypoints as input for the fitting process. This bootstrapping procedure is iterated. The least-confident samples are manually annotated (Section 4.3). With this *human-in-the-loop* process, we quickly obtain a large scale annotated dataset. Next we describe each stage of this procedure in detail.

#### 4.1. Hand Model Fitting with Sparse Annotations

Our goal is to fit a deformable hand shape model to observations from multiple views acquired at the same time. We build on the statistical *MANO* model, proposed by Romero *et al.* [21], which is parameterized by  $\theta \in \mathbb{R}^{61}$ . The model parameters  $\theta = (\alpha, \beta, \gamma)^T$  include shape  $\alpha \in \mathbb{R}^{10}$ , articulation  $\beta \in \mathbb{R}^{45}$  as well as global translation and orientation  $\gamma \in \mathbb{R}^6$ . Using keypoint and segmentation information we optimize a multi-term loss,

$$\mathcal{L} = \mathcal{L}_{\text{kp}}^{2\text{D}} + \mathcal{L}_{\text{kp}}^{3\text{D}} + \mathcal{L}_{\text{seg}} + \mathcal{L}_{\text{shape}} + \mathcal{L}_{\text{pose}}, \quad (3)$$

to estimate the model parameters  $\tilde{\theta}$ , where the tilde indicates variables that are being optimized. We describe each of the terms in (3) next.

**2D Keypoint Loss  $\mathcal{L}_{\text{kp}}^{2\text{D}}$ :** The loss is the sum of distances between the 2D projection  $\Pi^i$  of the models’ 3D keypoints

$\tilde{\mathbf{p}}_k \in \mathbb{R}^3$  to the 2D annotations  $\mathbf{q}_k^i$  over views  $i$  and visible keypoints  $k \in V_i$ :

$$\mathcal{L}_{\text{kp}}^{2\text{D}} = w_{\text{kp}}^{2\text{D}} \cdot \sum_i \sum_{k \in V_i} \cdot \|\mathbf{q}_k^i - \Pi^i(\tilde{\mathbf{p}}_k)\|_2. \quad (4)$$

**3D keypoint Loss  $\mathcal{L}_{\text{kp}}^{3\text{D}}$ :** This loss is defined in a similar manner as (4), but over 3D keypoints. Here,  $\mathbf{p}_k$  denotes the 3D keypoint annotations, whenever such annotations are available (e.g., if predicted by *MVNet*),

$$\mathcal{L}_{\text{kp}}^{3\text{D}} = w_{\text{kp}}^{3\text{D}} \cdot \sum_i \sum_{k \in V_i} \|\mathbf{p}_k - \tilde{\mathbf{p}}_k\|_2. \quad (5)$$

**Segmentation Loss  $\mathcal{L}_{\text{seg}}$ :** For shape optimization we use a sum of  $l_2$  losses between the model dependent mask  $\tilde{\mathbf{M}}^i$  and the manual annotation  $\mathbf{M}^i$  over views  $i$ :

$$\mathcal{L}_{\text{seg}} = w_{\text{seg}} \cdot \sum_i (\|\mathbf{M}^i - \tilde{\mathbf{M}}^i\|_2 + \|\text{EDT}(\mathbf{M}^i) \cdot \tilde{\mathbf{M}}^i\|_2). \quad (6)$$

Additionally, we apply a silhouette term based on the Euclidean Distance Transform (EDT). Specifically, we apply a symmetric EDT to  $\mathbf{M}^i$ , which contains the distance to the closest boundary pixel at every location.

**Shape Prior  $\mathcal{L}_{\text{shape}}$ :** For shape regularization we employ

$$\mathcal{L}_{\text{shape}} = w_{\text{shape}} \cdot \|\tilde{\beta}\|_2, \quad (7)$$

which enforces the predicted shape to stay close to the mean shape of *MANO*.

**Pose Prior  $\mathcal{L}_{\text{pose}}$ :** The pose prior has two terms. The first term applies a regularization on the PCA coefficients  $\mathbf{a}_j$  used to represent the pose  $\tilde{\alpha}$  in terms of PCA basis vectors  $\mathbf{c}_j$  (i.e.,  $\tilde{\alpha} = \sum_j \tilde{a}_j \cdot \mathbf{c}_j$ ). This regularization enforces predicted poses to stay close to likely poses with respect to the PCA pose space of *MANO*. The second term regularizes the distance of the current pose  $\tilde{\alpha}$ , to the  $N$  nearest neighbors of a hand pose dataset acquired from [5]:

$$\mathcal{L}_{\text{pose}} = w_{\text{pose}} \cdot \sum_j \|\tilde{a}_j\|_2 + w_{\text{nn}} \cdot \sum_{n \in N} \|\alpha^n - \tilde{\alpha}\|_2. \quad (8)$$

We implement the fitting process in Tensorflow [1] and use *MANO* to implement a differentiable mapping from  $\tilde{\theta}$  to 3D model keypoints  $\tilde{\mathbf{p}}_k$  and 3D model vertex locations  $\tilde{\mathbf{V}} \in \mathbb{R}^{778 \times 3}$ . We adopt the Neural Renderer [14] to render the segmentation masks  $\tilde{\mathbf{M}}^i$  from the hand model vertices  $\tilde{\mathbf{V}}$  and use the ADAM optimizer [15] to minimize:

$$\theta = \arg \min_{\tilde{\theta}} (\mathcal{L}(\tilde{\theta})) \quad (9)$$




Figure 5: *MVNet* predicts a single hand pose  $\mathbf{P}$  using images of all 8 views (for simplicity only 2 are shown). Each image is processed separately by a 2D CNN that is shared across views. This yields 2D feature maps  $\mathbf{f}_i$ . These are individually reprojected into a common coordinate frame using the known camera calibration to obtain  $\mathbf{F}_i = \Pi^{-1}(\mathbf{f}_i)$ . The  $\mathbf{F}_i$  are aggregated over all views and finally a 3D CNN localizes the 3D keypoints within a voxel representation.

## 4.2. MVNet: Multiview 3D Keypoint Estimation

To automate the fitting process, we seek to estimate 3D keypoints automatically. We propose *MVNet* shown in Fig. 5 that aggregates information from all eight camera images  $\mathbf{I}_i$  and predicts a single hand pose  $\mathbf{P} = \{\mathbf{p}_k\}$ . We use a differentiable unprojection operation, similar to Kar *et al.* [13], to aggregate features from each view into a common 3D volume.

To this end, we formulate the keypoint estimation problem as a voxel-wise regression task:

$$\mathcal{L}_{\text{MVNet}} = \frac{1}{K} \sum_k \|\mathbf{S}_k - \tilde{\mathbf{S}}_k\|_2, \quad (10)$$

where  $\tilde{\mathbf{S}}_k \in \mathbb{R}^{N \times N \times N}$  represents the prediction of the network for keypoint  $k$  and  $\mathbf{S}_k$  is the ground truth estimate we calculate from validated *MANO* fits.  $\mathbf{S}_k$  is defined as a normalized Gaussian distribution centered at the true keypoint location. The predicted point  $\tilde{\mathbf{p}}_k$  is extracted as maximal location in  $\tilde{\mathbf{S}}_k$ . Furthermore, we define the confidence  $c$  of a prediction as maximum along the spatial and average over the keypoint dimension:

$$c = \frac{1}{K} \sum_k (\max_{i,j,l} \tilde{\mathbf{S}}_k(i,j,l)). \quad (11)$$

Additional information can be found in the supplemental material.

## 4.3. Iterative Refinement

In order to generate annotations at large scale, we propose an iterative, *human-in-the-loop* procedure which is visualized in Fig. 4. For initial bootstrapping we use a set of manual annotations to generate the initial dataset  $\mathcal{D}_0$ . In iteration  $i$  we use dataset  $\mathcal{D}_i$ , a set of images and the corresponding *MANO* fits, to train *MVNet* and *HandSegNet* [33]. *MVNet* makes 3D keypoint predictions along with confidence scores for the remaining unlabeled data and *Hand-*

Method	mesh error ↓	F@5mm ↑	F@15mm ↑
Mean shape	1.78	0.300	0.808
MANO Fit	1.45	0.415	0.884
MANO CNN	1.16	0.484	0.925

Table 3: This table shows shape prediction performance on the evaluation split of FreiHAND after rigid alignment. We report two measures: The mean mesh error and the F-score at two different distance thresholds.

*SegNet* predicts hand segmentation masks. Using these predictions, we perform the hand shape fitting process of Section 4.1. Subsequently, we perform *verification* that either accepts, rejects or partially annotates some of these data samples.

**Heuristic Verification.** We define a heuristic consisting of three criteria to identify data samples with good *MANO* fits. First, we require the mean *MVNet* confidence score to be above 0.8 and all individual keypoint confidences to be at least 0.6, which enforces a minimum level of certainty on the 3D keypoint prediction. Second, we define a minimum threshold for the intersection over union (IoU) between predicted segmentation mask and the mask derived from the *MANO* fitting result. We set this threshold to be 0.7 on average across all views while also rejecting samples that have more than 2 views with an IoU below 0.5. Third, we require the mean Euclidean distance between predicted 3D keypoints and the keypoints of the fitted *MANO* to be at most 0.5 cm where no individual keypoint has a Euclidean distance greater than 1 cm. We accept only samples that satisfy all three criteria and add these to the set  $\mathcal{D}_i^h$ .

**Manual Verification and Annotation.** The remaining unaccepted samples are sorted based on the confidence score of *MVNet* and we select samples from the 50<sup>th</sup> percentile upwards. We enforce a minimal temporal distance between samples selected to ensure diversity as well as choosing samples for which the current pose estimates are sufficiently different to a flat hand shape as measured by the Euclidean distance in the pose parameters. We ask the annotators to evaluate the quality of the *MANO* fits for these samples. Any sample that is verified as a good fit is added to the set  $\mathcal{D}_i^m$ . For remaining samples, the annotator has the option of either discarding the sample in which case is marked as *unlabeled* or provide additional annotations (e.g., annotating mislabeled finger tips) to help improve the fit. These additionally annotated samples are added to the set  $\mathcal{D}_i^l$ .

Joining the samples from all streams yields a larger labeled dataset

$$\mathcal{D}_{i+1} = \mathcal{D}_i + \mathcal{D}_i^h + \mathcal{D}_i^m + \mathcal{D}_i^l \quad (12)$$

which allows us to retrain both *HandSegNet* and *MVNet*. We repeated this process 4 times to obtain our final dataset.

## 5. Experiments

### 5.1. Cross-Dataset Generalization of FreiHAND

To evaluate the cross-dataset generalization capability of our dataset and to compare to the results of Table 1, we define the following training and evaluation split: there are samples with and without green screen and we chose to use all green screen recordings for training and the remainder for evaluation. Training and evaluation splits contain data from 24 and 11 subjects, respectively, with only 3 subjects shared across splits. The evaluation split is captured in 2 different indoor and 1 outdoor location. We augmented the training set by leveraging the green screen for easy and effective background subtraction and creating composite images using new backgrounds. To avoid green color bleeding at the hand boundaries we applied the image harmonization method of Tsai *et al.* [26] and the deep image colorization approach by Zhang *et al.* [31] separately to our data. Both the automatic and sampling variant of [31] were used. With the original samples this quadruples the training set size from 33 k unique to 132 k augmented samples. Examples of resulting images are shown in Fig. 2.

Given the training and evaluation split, we train the single view 3D pose estimation network on our data and test it across different datasets. As shown in Table 1, the network achieves strong accuracy across all datasets and ranks first in terms of cross-dataset generalization.

### 5.2. 3D Shape Estimation

Having both pose and shape annotations, our acquired dataset can be used for training shape estimation models in a fully supervised way. In addition, it serves as the first real dataset that can be utilized for evaluating shape estimation methods. Building on the approach of Kanazawa *et al.* [12], we train a network that takes as input a single RGB image and predicts the *MANO* parameters  $\tilde{\theta}$  using the following loss:

$$\mathcal{L} = w_{3D} \|\mathbf{p}_k - \tilde{\mathbf{p}}_k\|_2 + w_{2D} \|\Pi(\mathbf{p}_k) - \Pi(\tilde{\mathbf{p}})\|_2 + w_p \|\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}}\|_2. \quad (13)$$

We deploy  $l_2$  losses for 2D and 3D keypoints as well as the model parameters and chose the weighting to  $w_{3D} = 1000$ ,  $w_{2D} = 10$  and  $w_p = 1$ .

We also provide two baseline methods, constant mean shape prediction, without accounting for articulation changes, and fits of the *MANO* model to the 3D keypoints predicted by our single-view network.

For comparison, we use two scores. The *mesh error* measures the average Euclidean distance between corresponding mesh vertices in the ground truth and the predicted hand shape. We also evaluate the *F-score* [16] which, given a distance threshold, defines the harmonic mean between



recall and precision between two sets of points [16]. In our evaluation, we use two distances:  $F@5\text{mm}$  and  $F@15\text{mm}$  to report the accuracy both at fine and coarse scale. In order to decouple shape evaluation from global rotation and translation, we first align the predicted meshes using Procrustes alignment as a rigid body transformation. Results are summarized in Table 3. Estimating *MANO* parameters directly with a CNN performs better across all measures than the baseline methods. The evaluation reveals that the difference in  $F$ -score is more pronounced in the high accuracy regime. Qualitative results of our network predictions are provided in Fig. 6.




Figure 6: Given a single image (top rows), qualitative results of predicted hand shapes (bottom rows) are shown. Please note that we don’t apply any alignment of the predictions with respect to the ground truth.

### 5.3. Evaluation of Iterative Labeling

In the first step of iterative labeling process, we set  $w_{\text{kp}}^{2\text{D}} = 100$  and  $w_{\text{kp}}^{3\text{D}} = 0$  (since no 3D keypoint annotations are available),  $w_{\text{seg}} = 10.0$ ,  $w_{\text{shape}} = 100.0$ ,  $w_{\text{nn}} = 10.0$ , and  $w_{\text{pose}} = 0.1$ . (For subsequent iterations we set  $w_{\text{kp}}^{2\text{D}} = 50$  and  $w_{\text{kp}}^{3\text{D}} = 1000$ .) Given the fitting results, we train *MVNet* and test it on the remaining dataset. After the first verification step, 302 samples are accepted. Validating a sample takes about 5 seconds and we find that the global pose is captured correctly in most cases, but in order to obtain high quality ground truth, even fits with minor inaccuracies are discarded.

Dataset	$\mathcal{D}_0$	$\mathcal{D}_1$	$\mathcal{D}_2$	$\mathcal{D}_3$	$\mathcal{D}_4$
#samples	302	993	1449	2609	4565
<i>RHD</i>	0.244	0.453	0.493	0.511	0.518
<i>PAN</i>	0.347	0.521	0.521	0.539	0.562

Table 4: Bootstrapping convergence is evaluated by reporting cross-dataset generalization to *RHD* and *PAN*. The measure of performance is AUC, which shows monotonous improvement throughout.

We use the additional accepted samples to retrain *MVNet* and *HandSegNet* and iterate the process. At the end of the first iteration we are able to increase the dataset to 993 samples, 140 of which are automatically accepted by heuristic, and the remainder from verifying 1000 samples. In the second iteration the total dataset size increases to 1449, 289 of which are automatically accepted and the remainder stems from verifying 500 samples. In subsequent iterations the complete dataset size is increased to 2609 and 4565 samples, where heuristic accept yields 347 and 210 samples respectively. This is the dataset we use for the cross-dataset generalization (see Table 1) and shape estimation (see Table 3) experiments.

We evaluate the effectiveness of the iterative labeling process by training a single view 3D keypoint estimation network on different iterations of our dataset. For this purpose, we chose two evaluation datasets that reached a good average rank in Table 1. Table 4 reports the results and shows a steady increase for both iterations as our dataset grows. More experiments on the iterative procedure are located in the supplemental material.

## 6. Conclusion

We presented *FreiHAND*, the largest RGB dataset with hand pose and shape labels of real images available to date. We capture this dataset using a novel iterative procedure. The dataset allows us improve generalization performance for the task of 3D hand pose estimation from a single image, as well as supervised learning of monocular hand shape estimation.

To facilitate research on hand shape estimation, we plan to extend our dataset even further to provide the community with a challenging benchmark that takes a big step towards evaluation under realistic *in-the-wild* conditions.

## Acknowledgements

We gratefully acknowledge funding by the Baden-Württemberg Stiftung as part of the RatTrack project. Work was partially done during Christian’s internship at Adobe Research.



## References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016. 6
- [2] Luca Ballan, Aparna Taneja, Jürgen Gall, Luc Van Gool, and Marc Pollefeys. Motion capture of hands in action using discriminative salient points. In *European Conference on Computer Vision*, pages 640–653. Springer, 2012. 2
- [3] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *Proc. of the Europ. Conf. on Computer Vision (ECCV)*, pages 561–578. Springer, 2016. 2
- [4] Adnane Boukhayma, Rodrigo de Bem, and Philip HS Torr. 3d hand shape and pose from images in the wild. *arXiv preprint arXiv:1902.03451*, 2019. 1
- [5] Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, and Tae-Kyun Kim. First-person hand action benchmark with rgb-d videos and 3d hand pose annotations. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 409–419, 2018. 3, 4, 6
- [6] Lihao Ge, Zhou Ren, Yuncheng Li, Zehao Xue, Yingying Wang, Jianfei Cai, and Junsong Yuan. 3d hand shape and pose estimation from a single rgb image. *arXiv preprint arXiv:1903.00812*, 2019. 1
- [7] Francisco Gomez-Donoso, Sergio Orts-Escolano, and Miguel Cazorla. Large-scale multiview 3d hand pose dataset. *Image and Vision Computing*, 81:25–33, 2019. 1, 3, 4
- [8] Shreyas Hampali, Markus Oberweger, Mahdi Rad, and Vincent Lepetit. Ho-3d: A multi-user, multi-object dataset for joint 3d hand-object pose estimation. *arXiv preprint arXiv:1907.01481*, 2019. 3, 4
- [9] Gerrit Hillebrand, Martin Bauer, Kurt Achatz, Gudrun Klinker, and Am Oferl. Inverse kinematic infrared optical finger tracking. In *Proceedings of the 9th International Conference on Humans and Computers (HC 2006)*, Aizu, Japan, pages 6–9. Citeseer, 2006. 2
- [10] Umar Iqbal, Pavlo Molchanov, Thomas Breuel, Juergen Gall, and Jan Kautz. Hand pose estimation via latent 2.5 d heatmap regression. In *Proc. of the Europ. Conf. on Computer Vision (ECCV)*, pages 118–134, 2018. 3, 4
- [11] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 8320–8329, 2018. 3
- [12] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 7122–7131, 2018. 7
- [13] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. In *Proc. of Int. Conf. on Neural Information Processing Systems (NIPS)*, pages 365–376, 2017. 6
- [14] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3907–3916, 2018. 6
- [15] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. 6
- [16] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):78, 2017. 7, 8
- [17] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J. Black, and Peter V. Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 2
- [18] Leap Motion. <https://www.leapmotion.com>. 3
- [19] Franziska Mueller, Florian Bernard, Oleksandr Sotnychenko, Dushyant Mehta, Srinath Sridhar, Dan Casas, and Christian Theobalt. Generated hands for real-time 3d hand tracking from monocular rgb. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 49–59, 2018. 1, 3, 4
- [20] Franziska Mueller, Dushyant Mehta, Oleksandr Sotnychenko, Srinath Sridhar, Dan Casas, and Christian Theobalt. Real-time hand tracking under occlusion from an egocentric rgb-d sensor. In *Int. Conf. on Computer Vision (ICCV)*, October 2017. 1, 2, 3
- [21] 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. 2, 5
- [22] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1145–1153, 2017. 1, 2, 4, 5
- [23] Srinath Sridhar, Franziska Mueller, Michael Zollhoefer, Dan Casas, Antti Oulasvirta, and Christian Theobalt. Real-time joint tracking of a hand manipulating an object from rgb-d input. In *Proc. of the Europ. Conf. on Computer Vision (ECCV)*, October 2016. 2
- [24] Anastasia Tkach, Mark Pauly, and Andrea Tagliasacchi. Sphere-meshes for real-time hand modeling and tracking. *ACM Transactions on Graphics (ToG)*, 35(6):222, 2016. 2
- [25] Anastasia Tkach, Andrea Tagliasacchi, Edoardo Remelli, Mark Pauly, and Andrew Fitzgibbon. Online generative model personalization for hand tracking. *ACM Transactions on Graphics (ToG)*, 36(6):243, 2017. 2
- [26] Yi-Hsuan Tsai, Xiaohui Shen, Zhe Lin, Kalyan Sunkavalli, Xin Lu, and Ming-Hsuan Yang. Deep image harmonization. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3789–3797, 2017. 7
- [27] Dimitrios Tzionas, Luca Ballan, Abhilash Srikantha, Pablo Aponte, Marc Pollefeys, and Juergen Gall. Capturing hands in action using discriminative salient points and physics simulation. *Int. Journal of Computer Vision*, 118(2):172–193, 2016. 1
- [28] Dimitrios Tzionas, Abhilash Srikantha, Pablo Aponte, and Juergen Gall. Capturing hand motion with an rgb-d sensor,

- fusing a generative model with salient points. In *Proc. of the German Conf. on Computer Vision (GCPR)*, pages 277–289. Springer, 2014. [2](#)
- [29] Robert Y Wang and Jovan Popović. Real-time hand-tracking with a color glove. *ACM Transactions on Graphics (ToG)*, 28(3):63, 2009. [2](#)
- [30] Jiawei Zhang, Jianbo Jiao, Mingliang Chen, Liangqiong Qu, Xiaobin Xu, and Qingxiong Yang. 3d Hand Pose Tracking and Estimation Using Stereo Matching. *arXiv preprint arXiv:1610.07214*, 2016. [2](#), [3](#), [4](#)
- [31] Richard Zhang, Jun-Yan Zhu, Phillip Isola, Xinyang Geng, Angela S Lin, Tianhe Yu, and Alexei A Efros. Real-time user-guided image colorization with learned deep priors. *arXiv preprint arXiv:1705.02999*, 2017. [7](#)
- [32] Thomas G Zimmerman, Jaron Lanier, Chuck Blanchard, Steve Bryson, and Young Harvill. A hand gesture interface device. In *ACM SIGCHI Bulletin*, pages 189–192. ACM, 1987. [2](#)
- [33] Christian Zimmermann and Thomas Brox. Learning to estimate 3d hand pose from single rgb images. In *Int. Conf. on Computer Vision (ICCV)*, 2017. <https://arxiv.org/abs/1705.01389>. [1](#), [3](#), [4](#), [6](#)