

Two-hand Global 3D Pose Estimation using Monocular RGB

Fanqing Lin Connor Wilhelm Tony Martinez

Department of Computer Science, Brigham Young University

Abstract

We tackle the challenging task of estimating global 3D joint locations for both hands via only monocular RGB input images. We propose a novel multi-stage convolutional neural network based pipeline that accurately segments and locates the hands despite occlusion between two hands and complex background noise and estimates the 2D and 3D canonical joint locations without any depth information. Global joint locations with respect to the camera origin are computed using the hand pose estimations and the actual length of the key bone with a novel projection algorithm. To train the CNNs for this new task, we introduce a large-scale synthetic 3D hand pose dataset. We demonstrate that our system outperforms previous works on 3D canonical hand pose estimation benchmark datasets with RGB-only information. Additionally, we present the first work that achieves accurate global 3D hand tracking on both hands using RGB-only inputs and provide extensive quantitative and qualitative evaluation.

1. Introduction

As the primary operating tool for human activities, the hands play a significant role in applications such as gesture control, action recognition, human-computer interaction and VR/AR. As the field of computer vision advances, commercial systems [3, 1, 2, 4] are shifting from marker/glove-based methods to vision-based hand tracking and pose estimation. However, accurate hand pose estimation from camera inputs remains challenging due to the possible heavy occlusion from the hand itself, the other hand or objects, complex background noise and the large pose space.

Most contemporary vision-based markerless works tackling the task of 3D hand pose estimation rely on depth information, requiring either multi-view setup or depth cameras. However, such hardware requirements add severe limitations to the possible applications by significantly increasing the setup overhead and cost. Depth cameras also only work in indoor scenes and have relatively high-power consump-

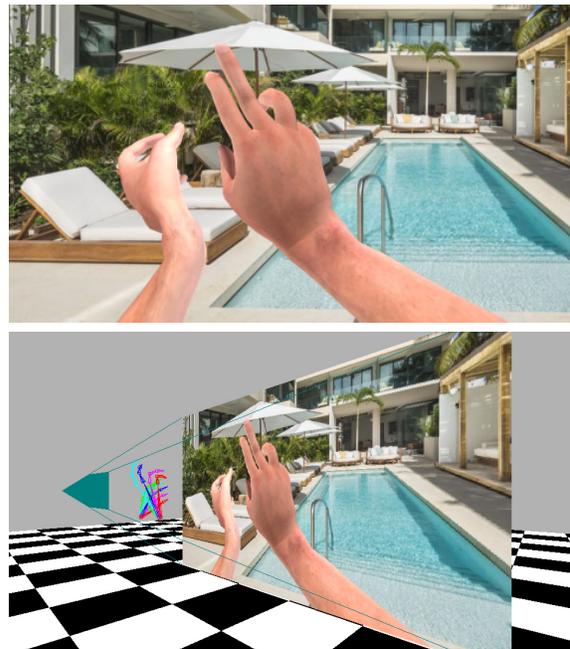


Figure 1: We present an approach to estimate the global 3D hand poses for both hands from a single RGB image, a new task that is particularly challenging not only due to complex background noise and various types of occlusion, but also the lack of depth information for estimating the distances. Given a RGB image (top), we show the side view of our estimation of global 3D hand poses (bottom).

tion. To circumvent this problem, some recent approaches tackle 3D canonical one-hand pose estimation using deep CNNs with only RGB-based inputs and show good results.

In this paper, we present the first algorithm that simultaneously estimates the 3D global joint locations of both hands with respect to the camera origin using monocular RGB inputs (Fig. 1), which is an essential step towards the next generation gesture control and pose recognition systems. Our pipeline consists of 4 major components: (1) hand segmentation and detection, (2) 2D hand pose estimation, (3) 3D canonical hand pose estimation and (4) 3D

global hand pose estimation.

The first challenge of training our pipeline is the lack of annotated data from existing benchmark datasets. Real-world markerless 3D hand tracking data collected using multiple cameras or RGB-D camera setup inevitably has tracking error. Manual annotation is time-consuming and infeasible for large-scale data collection. Consequently, many recent benchmark datasets provide synthetic data with perfect ground truth annotation of the joint locations. However, synthetic images have different statistical distributions than real-world images and knowledge learned on synthetic data does not always transfer to the real-world domain, since CNNs are sensitive to textural information.

Since it is currently infeasible to collect real-world two-hand pose data with accurate 2D and 3D joint annotations at a large-scale with sufficient variety, we create a novel high-quality synthetic 3D hand dataset suitable for training and evaluating the networks and focus on the challenging task of two-hand global 3D pose estimation using RGB.

In summary, our main contributions include:

- The first system capable of estimating 3D global joint locations for both hands using monocular RGB inputs. We introduce a viewpoint-invariant global projection algorithm capable of perfectly reconstructing the absolute 3D joint locations and the evaluation protocol for the new task.
- A novel egocentric RGB-D two-part (static + dynamic) synthetic dataset for the task of two-hand 3D global pose estimation, which introduces unique challenges and also benefits researches in hand segmentation and detection, 2D and 3D canonical pose estimation.
- Extensive evaluation on both two-hand 3D global and single-hand 3D canonical hand pose estimation on 4 target datasets. Our networks outperform the current state-of-the-art canonical methods with less information (only RGB) and additionally achieve promising results for global pose estimation.

2. Related Work

Hand pose estimation is a long-standing research area due to its wide range of applications. Compared to the popular task of body pose estimation, vision-based 3D hand pose estimation has more complex articulation, heavier occlusion and more restricted availability of data. We first review the most relevant previous methods that utilize depth information, then shift our emphasis to approaches that use RGB-only data.

Depth-based methods. Oberweger *et al.* [16] and Zhou *et al.* [36] introduced CNN architectures that regressed 3D joint locations from depth images directly. Ge *et al.* [8] proposed to project the depth image onto three orthogonal planes and fuse the corresponding 2D joint locations for the

final 3D joint locations. Cai *et al.* [6] and Iqbal *et al.* [10] proposed models capable of training on RGB-D data and evaluating on RGB inputs.

Multiple-camera methods. Many methods use multiple RGB cameras to gain additional information from different viewpoints that can help with resolving the depth ambiguity and heavy occlusion. Wang *et al.* [30] used 2 cameras and estimated 3D hand pose by matching with instances in a hand database. Oikonomidis *et al.* [17] demonstrated the tracking of both the hand and an interacting object in 3D with 8 surrounding fixed cameras. Sridhar *et al.* [23, 24] estimated the hand pose by using generative approach on inputs of multiple RGB cameras and a depth sensor. Zhang *et al.* [35] introduced 3D hand pose estimation using matching algorithm on inputs from stereo cameras.

Single-camera methods. Due to the significantly higher setup overhead and costs introduced by depth sensors and multiple calibrated cameras, some methods use a single RGB image to estimate the 3D hand poses. Zimmermann and Brox [37] proposed a CNN-based pipeline that estimates the 2D joint locations and lifts the 2D heatmaps to 3D canonical joint locations. Mueller *et al.* [13] introduced a model that estimates both 2D and 3D canonical joint locations with kinematic skeleton fitting to better address physical constraints and temporal smoothness. Spurr *et al.* [21] and Yang *et al.* [33] proposed to use variational autoencoders to learn a latent space for hand poses, which are capable of estimating 3D hand poses from RGB inputs. Some methods [5, 32, 9] estimated the low-dimensional parameters for a 3D deformable hand model [19] to fit the RGB inputs in order to retrieve the 3D canonical hand poses.

Two-hand methods. It is more natural yet difficult to estimate poses for two interacting hands due to inter-hand occlusion. Tzionas *et al.* [29], Taylor *et al.* [27] and Mueller *et al.* [14] achieved promising results using energy optimization to fit parametric hand models using depth data. Note that 3D pose estimation using RGB data is much more challenging due to the lack of depth information and the additional noise from images in the wild. To the best of our knowledge, our work is the first to address two-hand global pose estimation using RGB data from a single camera.

3. Datasets for Hand Pose Estimation

For depth-based hand pose estimation, [28, 26, 25, 34, 7] presented large-scale datasets consisting of real depth images with estimated ground-truth joint locations. For RGB-D hand pose estimation, due to the need to manually annotate joint locations, small-scale datasets [15, 22, 35] with limited variation were presented with real RGB and depth data. For RGB-based hand pose estimation, [20, 38] performed extensive manual annotation and provided a decent amount of labeled real-world instances. Note that accurately annotated hand data with sufficient variation is

necessary for learning-based approaches, and datasets with real-world RGB/depth data can only provide estimation of the joint locations as the ground truth and some severely lacks sufficient variety. Consequently, synthetic datasets [13, 15, 37] with large amount of color, depth images and perfect annotation are introduced for advancing research in the field. It is worth mentioning that existing RGB-based datasets and methods are designed to estimate the 3D hand poses in a canonical (localized) frame for a single hand. Therefore, it is necessary to generate a new dataset suitable for the task of RGB-based two-hand global pose estimation.

3.1. Ego3DHands Dataset

We introduce the first dataset for the task of two-hand global 3D pose estimation from an egocentric view. Following [37], the dataset is generated using rendering from *Blender*¹, which enables us to obtain the segmentation masks of hand parts as well as the annotated 2D and 3D joint locations (infeasible to obtain on real hand data at large-scale with variety). We utilize a single character from *Mixamo*² to keep the bone ratios of the hands consistent for global 3D pose reconstruction. The dataset includes two versions for static and dynamic hand pose estimation respectively. Despite the domain gap between synthetic and real-world data, this dataset enables training for learning-based approaches and quantitative analysis for a new task.

Data Representation. As illustrated in Fig. 2, the dataset provides 7 segmentation masks for each hand. The 2D joint locations are normalized values ranging from (0, 0) at the top left to (1, 1) at the bottom right. The global 3D coordinates are represented in the camera space. We scale the 3D coordinates so that the bone length from the wrist to the middle metacarpophalangeal (mMCP) is 10 cm. Each hand consists of 21 joints: wrist and 4 joints for each finger. The Depth map is also provided but not used in this work.

Static Ego3DHands. To capture the images for static hand poses, we set the camera to be between the eyes of the character facing forward. We keep the hands inside a fix-sized bounding box in front of the character so the targets stay in sight for estimation. Each hand has a 10% drop rate for single-hand scenarios. The rotational angles for arm and hand joints are randomized within reasonable rotational ranges to obtain vast variety in the pose space. We include 4 light sources with slightly randomized color, brightness and position for illumination. Additionally, for the background of the hand pose images, we selected 100 unique scene topics and collected 20,000 images from online sources, on which we further applied random color augmentation and horizontal flips. We create 50,000 instances for the training set and 5,000 instances for the test set.

Dynamic Ego3DHands. For global dynamic two-hand



Figure 2: Our dataset provides a total of 14 segmentation masks (right) for the fingers, palm and arm along with the 2D and 3D joint annotations.

3D tracking, we introduce an additional dataset with 100 sequences for the training set, and 10 sequences for the test set. Each sequence consists of 500 frames where we randomize independent motion for both hands. For background sequences, we selected 110 short videos with variety from Pexels³, so each hand pose sequence has a unique corresponding background sequence. This dataset enables researchers to explore methods for 3D global hand pose estimation that utilize temporal consistency. We report our baseline results in Section 5 for future comparison.

4. Method

In this paper, we present the first algorithm capable of estimating the global 3D poses of both hands from a monocular RGB image. The overall system is demonstrated in Fig. 3. Given a single RGB image as input, we use *HandSegNet* to simultaneously obtain the segmentation masks and the heatmap energy of both hands. The hand heatmap energy indicates the approximate locations of the hands despite occlusion, which are used to detect and provide a cropped image for each hand. The cropped RGB hand images are then processed using the corresponding segmentation masks for the next stage. To estimate the 2D joint locations, we present *PoseNet_{2D}* that estimates and refines the 2D heatmaps of the joints in multiple stages. To lift the 2D heatmaps to a 3D pose estimation, we present *PoseNet_{3D}* that takes the heatmaps as input and regresses the 3D canonical joint locations which we define in detail in Section 4.3. Finally, we present a novel algorithm that accurately estimates the 3D global hand joint locations in the spherical coordinate system using the obtained 2D and 3D canonical information, the actual length of key bone and the camera intrinsics. Our method can theoretically be applied to estimate the global location and pose of any object given the aforementioned information.

4.1. Two-hand Segmentation and Detection

Unlike existing methods that perform pose estimation on a single cropped hand, we need to first distinguish between left and right hand by estimating the individual hand locations. For the task of hand segmentation and detection, we use a deep convolutional neural network trained to predict

¹www.blender.org

²www.mixamo.com

³www.pexels.com

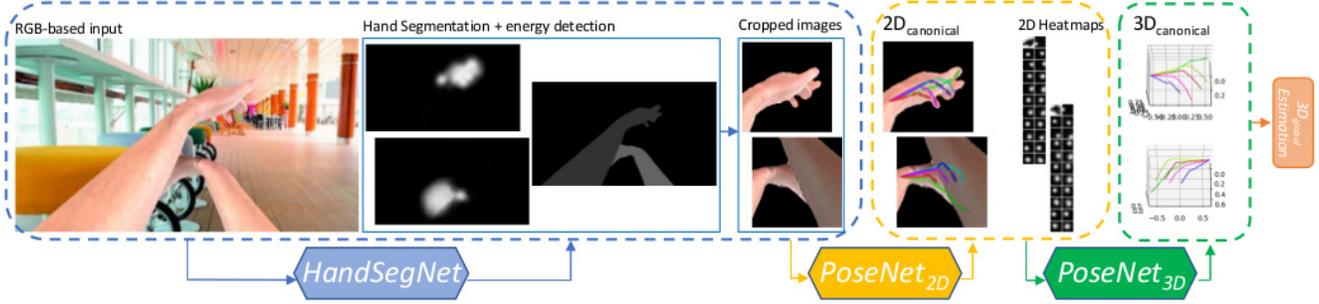


Figure 3: Overview of proposed pipeline for two-hand global pose estimation from monocular RGB. Given a RGB image, we segment and detect the hands, crop and process the hand images for 2D canonical pose estimation despite inter-hand occlusion, and use the 2D heatmaps to estimate 3D canonical poses. For the final step, We introduce a novel algorithm for computing the 3D global hand poses using the 2D and 3D canonical estimation as well as the actual bone lengths.

both the segmentation masks and the location of hands in form of heatmap energy. We show in Section 4.2 that the accurate segmentation of hands is necessary information for 2D hand pose estimation in the presence of occlusion from the other hand.

For the architecture of *HandSegNet*, we use a residual network for the task of semantic segmentation and hand detection. It consists of 4 downsampling and 4 upsampling layers comprised of 16 residual blocks. For the output layer, we have 3 channels for the task of segmentation (2 objects and background) and 2 additional channels for estimating the heatmap energy of the left and right hand. The heatmap energy is capable of providing high activation at locations of partially or even completely occluded hands. To generate the cropped bounding boxes, we apply Otsu Thresholding on the hand energy for selecting the high activation area. In the case of very low activation, we classify the corresponding hand(s) as being absent and drop the absent hand(s) in the subsequent stages.

4.2. 2D Canonical Hand Pose Estimation

The goal of *PoseNet_{2D}* is to estimate the 2D joint locations given a cropped hand image. We use a variant of Convolutional Pose Machines (CPM) [31] as our base model, with batch normalization layers inserted after the convolutional layers for better adaptation to the vast RGB image space. The 2D joint locations are represented as heatmaps and CPM refines the output heatmaps in progressive stages. Since the left and right hand have different articulation, we horizontally flip the cropped images of the right hand so the learned articulation remains consistent for the model. We resize the cropped input hand images from *HandSegNet* to 256x256. The output of the CPM consists of 21 heatmaps with size of 32x32. We generate stronger heatmap energy for closer joints so that depth information is encoded into the 2D heatmaps. We show that this heatmap generation technique (we refer to as *z-heatmaps*) improves accuracy for 3D canonical hand pose estimation in Section 5.

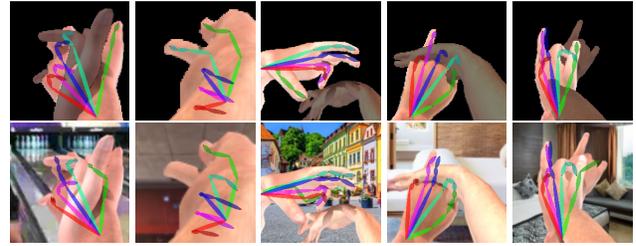


Figure 4: We show that 2D canonical hand pose estimation with segmentation information (top) better resolves inter-hand occlusion from the secondary hand.

Previous work [13] showed good performance on RGB-based 2D hand pose estimation with occlusion introduced by an interacting object. However, the method fails when the background has similar appearance as the hand. With the presence of both hands, accurate 2D hand pose estimation becomes more challenging due to the similar-object occlusion introduced by the secondary hand. We show in Fig. 4 that we successfully address this issue by providing the segmentation information necessary for distinguishing the left and right hand. By using the segmentation masks of the two hands, we simplify the input image space by removing the background noise; additionally, we differentiate the color space between two hands by reducing the brightness of the secondary hand by a factor of 0.5. As a result, it is very important for *HandSegNet* to produce accurate segmentation masks for the two hands. The output heatmaps are used as inputs for *PoseNet_{3D}* and we retrieve the resulting 2D global joint locations using the bounding box coordinates from *HandSegNet*. For a set of 2D global joint locations, we use $\mathbf{p}_j = (\rho_j^r, \rho_j^c)$, where ρ^r and ρ^c represent the corresponding row and column position of the j^{th} joint in form of percentages.

4.3. 3D Canonical Hand Joints Regression

The 3D canonical frame for a single hand is defined such that the middle metacarpophalangeal (mMCP) joint

is at the origin and the distance between wrist and mMCP is 1 [13]. The defined canonical frame requires the target hand to be in the middle of the cropped image so the z-axis aligns with the camera direction. Therefore, we generate the 3D canonical joint annotation for training by rotating the original global 3D joint locations of both hands to the center of the image; zero-centering on the mMCP and normalization is applied afterwards. Thus, for a set of annotated canonical 3D Cartesian coordinates represented as $\mathbf{w}_j = (x_j, y_j, z_j)$,

$$\begin{aligned} \mathbf{w}^{center} &= \mathcal{R} \cdot \mathbf{w}^{glob} \\ \mathbf{w}^{can} &= (\mathbf{w}^{center} - \mathbf{w}_{mmcp}^{center})/d \end{aligned} \quad (1)$$

where \mathcal{R} is the 3D rotational matrix for centering \mathbf{w}_{mmcp}^{can} and d is the distance between the wrist joint and the mMCP. As a result, our 3D canonical hand poses are consistent with the visual representation of the hands, which is necessary for estimating the global joint locations and will be explained in Section 4.4.

For the architecture of *PoseNet_{3D}*, we use a small residual network comprised of 8 residual blocks with 2 fully connected layers before the output layer. The input heatmaps are upscaled by a factor of 2 to the size of 64x64 for better performance. The model estimates the root-relative 3D coordinates of 21 joints for each hand.

To enforce physical constraints and encourage 2D pose consistency between w^{can} and \mathbf{p} , we employ the following loss function for training *PoseNet_{3D}*,

$$\mathcal{L}_{3d} = \mathcal{L}_j + \mathcal{L}_{bone} + \mathcal{L}_{proj} \quad (2)$$

where \mathcal{L}_j is the Mean Squared Error (MSE) loss for joint regression. In addition, we introduce \mathcal{L}_{bone} that indicates the MSE between the ground truth and the predicted bone lengths. \mathcal{L}_{proj} indicates the MSE between the (x, y) component of \mathbf{w}^{can} and \mathbf{p} projected into the same canonical frame. The overall \mathcal{L}_{3d} aims to produce physically plausible 3D canonical hand poses consistent with the 2D poses.

For the task of 3D canonical hand pose estimation only (Section 5.1), we expect *PoseNet_{3D}* to refine the potentially inaccurate 2D pose estimations and drop \mathcal{L}_{proj} . However, for the task of 3D global hand pose estimation in the next stage, we replace \mathbf{p} from *PoseNet_{2D}* with the (x, y) component of \mathbf{w}^{can} projected back into the pixel space since the consistency between 2D and 3D poses is important for our projection algorithm (Section 4.4).

4.4. 3D Global Hand Pose Estimation

The problem of 3D global hand pose estimation in Cartesian coordinate system introduces different challenges compared to conventional 3D canonical hand pose estimation. First, the same canonical hand pose in different global positions is visually rotated, thus introducing rotational

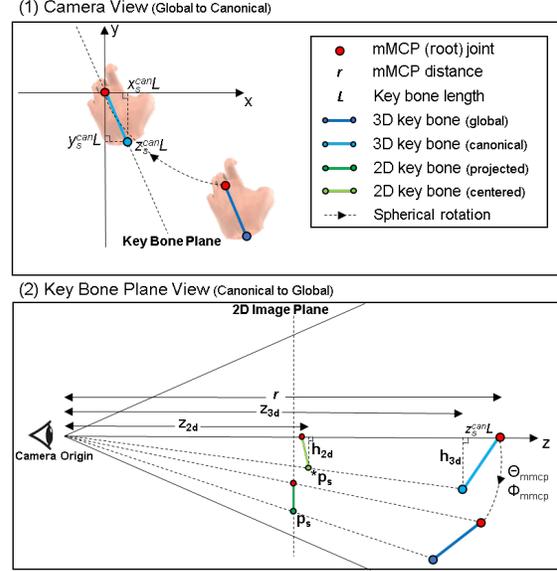


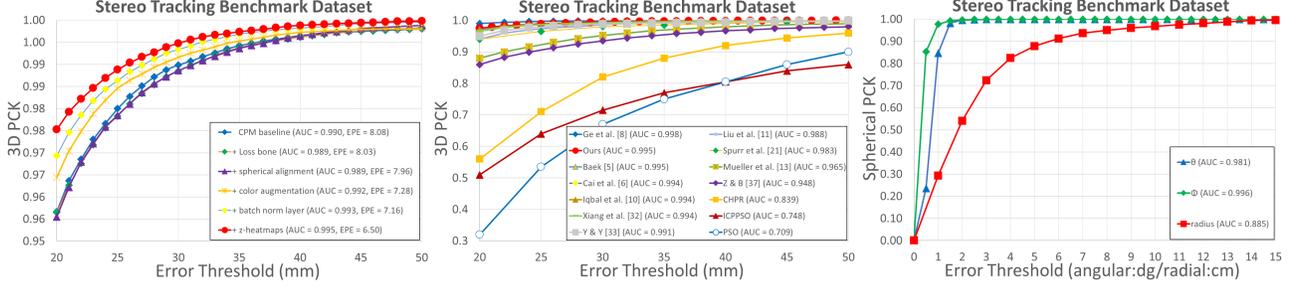
Figure 5: The camera view (top) shows how we treat the hand as if it is in the center of the image and generate the corresponding 3D canonical pose using Eq. (1) for training. The bone plane view (bottom) illustrates how we compute the absolute distance r of the centered root joint and subsequently return the pose to its original global 3D position by rotation in spherical coordinate system.

ambiguity. Second, the size of the hand correlates not to the z-value of its global Cartesian 3D position, but to its absolute distance to the camera origin, which introduces depth ambiguity. As a result, we propose a novel algorithm for global 3D hand pose estimation using the spherical coordinate system. As illustrated in Fig. 5, in order to transform the 3D canonical hand pose back to its original 3D global position, we first scale it by the known actual key bone length L to the real-world size, then translate it by r cm in the positive direction along the z-axis, and finally apply a 3D global rotation with θ_{mmcp} and ϕ_{mmcp} respectively in the spherical coordinate system. To compute a set of global 3D Cartesian coordinates \mathbf{w}^{glob} given \mathbf{w}^{can} in 3D and \mathbf{p} in 2D, we find the absolute spherical coordinate of the mMCP $\mathbf{v}_{mmcp} = (r_{mmcp}, \theta_{mmcp}, \phi_{mmcp})$. For the rotational angles,

$$\begin{aligned} \theta_{mmcp} &= \text{atan}(((\mathbf{p}_{mmcp}^r \cdot \mathbf{H}) - \mathbf{H}/2)/pxcm, foc) \\ \phi_{mmcp} &= \text{atan}(((\mathbf{p}_{mmcp}^c \cdot \mathbf{W}) - \mathbf{W}/2)/pxcm, foc) \end{aligned} \quad (3)$$

where \mathbf{H} and \mathbf{W} represent the height and width of the RGB input image, $pxcm$ is a constant conversion factor for converting from image pixels to centimeters and foc is the camera focal length.

In order to compute r_{mmcp} , we need to apply the Side Splitter Theorem on the right-angled similar triangles shown on the key bone plane (Fig. 5 (2)),



(a) Self-comparisons. (b) Comparison with the state-of-the-art. (c) Spherical PCK for global estimation.

Figure 6: Self-comparisons (left) and comparison with the state-of-the-art (middle) for 3D canonical hand pose estimation on the STB dataset. + indicates that the feature is applied incrementally. Spherical PCK (right) without alignment with the ground truth root joint is also reported.

$$z_{3d} = z_{2d} \cdot h_{3d}/h_{2d} \quad (4)$$

where z_{2d} and h_{2d} are computed by rotating \mathbf{p}_s (the secondary joint s that forms the key bone with mMCP) with $-\theta_{mmcp}$ and $-\phi_{mmcp}$. For h_{3d} , the segment of key bone perpendicular to the z-axis is

$$h_{3d} = \sqrt{x_s^{can2} + y_s^{can2}} \cdot \mathbf{L} \quad (5)$$

where $(x_s^{can}, y_s^{can}, z_s^{can})$ is the 3D canonical position for joint s . Finally, we compute the spherical radius of the mMCP,

$$r_{mmcp} = z_{3d} - z_s^{can} \cdot \mathbf{L} \quad (6)$$

with z_s^{can} being positive if the key bone extends away from the camera.

Since the key bone needs to have sufficient length in 2D images for accurate projection and estimation, we use mMCP as the primary joint for the key bone, and select either the wrist or the pinky MCP as the secondary joint. By selecting the longer one of the two "bones" as the key bone, we guarantee its validity since the two "bones" can never be parallel and therefore can never both point in the direction of z-axis in 2D images.

For tracking both hands in 3D global space in video settings with temporal smoothness, we apply polynomial regression for the estimation of r_{mmcp} . We use two queues (for left and right hands) to store the most recent r_{mmcp} values to estimate the current r_{mmcp} .

Our 3D global pose estimation algorithm can be applied generally to estimate the 3D global location of any objects given the 2D, 3D canonical information, actual key bone length and camera intrinsics. Unlike other methods [12, 18] that attempt to estimate the approximate global poses, our algorithm is, to our knowledge, the first capable of perfectly reconstructing the global 3D poses given accurate input information. We show its effectiveness in Section 5.2 on several hand pose datasets (both synthetic and real) with annotated 3D global joint locations and different viewpoints.

5. Experiments

We first compare the performance of $PoseNet_{2D}$ and $PoseNet_{3D}$ for single-hand 3D canonical pose estimation with the current state-of-the-art methods on two popular benchmark datasets: Stereo Tracking Benchmark Dataset (STB) [35] and Rendered Hand Pose Dataset (RHP) [37]. For two-hand global 3D pose estimation, we evaluate our method on the test sets of both the static ($Ego3D_s$) and the dynamic ($Ego3D_d$) versions of Ego3DHands. To demonstrate the effectiveness of our global pose estimation algorithm, we show quantitative and qualitative results for global hand pose estimation on all 4 target datasets. Training details are included in the supplementary document.

5.1. Single-hand Canonical Pose Estimation

Stereo Tracking Benchmark Dataset (STB) consists of 12 sequences (1500 frames per sequence) of captured single-hand motion of 1 subject with 6 different backgrounds and lighting. Stereo and depth images are provided, but only the RGB images from the left camera are used in this work. We follow the same evaluation protocol as [37], training on 10 and testing on the other 2 sequences. For evaluation of the 3D pose estimation, the 3D canonical pose needs to be scaled to its actual size and transformed to its global position using the ground truth root joint. Other methods scale \mathbf{w}^{can} by \mathbf{L} and simply translate \mathbf{w}_{mmcp}^{can} to \mathbf{w}_{mmcp}^{glob} (Cartesian alignment), which assumes that there is no rotational discrepancy since the hand is relatively close to the center of the camera. We align our canonical hand poses to global hand poses by spherical alignment. Specifically, we scale \mathbf{w}^{can} by \mathbf{L} , apply translation in the z-axis and rotation in the spherical coordinate system to align \mathbf{w}_{mmcp}^{can} with \mathbf{w}_{mmcp}^{glob} . In Fig. 6a, we perform various experiments for self-comparisons to justify our design choices, reporting the Area Under Curve (AUC) computed using the Percentage of Correct Keypoints (PCK), where a predicted keypoint is correct if its location is within the threshold radius around the ground truth. We also report the End Point Error (EPE) in mm. For comparison with other methods, we show

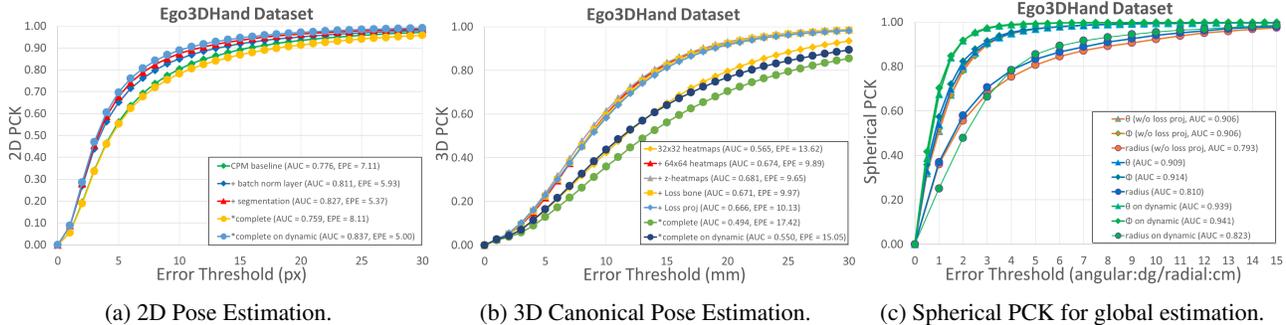


Figure 7: Quantitative results on the $Ego3D_s$ and $Ego3D_d$ for 2D (a), 3D canonical (b) and 3D global (c) hand pose estimation. 2D PCK is computed using image size of (270x480). Results are reported given the ground truth input for isolated studies unless marked as "*complete" in (a) and (b). "dynamic" indicates that experiments are performed on $Ego3D_d$.

Method	3D Hand Pose Estimation		
	AUC \uparrow	EPE (mm)	
		median \downarrow	mean \downarrow
Iqbal <i>et al.</i> [10]	0.940	11.33	13.41
Baek <i>et al.</i> [5]	0.926	–	–
Ge <i>et al.</i> [9]	0.920	–	–
Cai [6]	0.887	–	–
Yang <i>et al.</i> [33]	0.849	–	19.95
Spurr <i>et al.</i> [21]	0.849	–	19.73
Z & B [37]	–	18.80	–
Ours	0.929	11.86	13.47
Ours + seg	0.942	11.14	12.47

Table 1: Comparison with the state-of-the-art methods on the RHP dataset. 3D AUCs are computed over an error range from 20 to 50mm.

in Fig. 6b that we outperform most state-of-the-art methods with an AUC = 0.995. It is worth mentioning that fair comparison cannot be made since many methods utilized depth information [9, 6, 10], deformable hand model (guaranteed physical plausibility) [9, 5] or additional datasets [9, 5, 32] during training while we simply trained on the left RGB images in the training set of STB with color augmentation.

Rendered Hand Pose Dataset (RHP) provides 41258 images for training and 2728 images for evaluation. Each rendered image contains a single character performing 1 of 39 gestures and the view is focused on one of the two hands. The training and test sets contain 31 and 8 distinct gestures respectively. We use the same setting as our method for the STB dataset. Since the pose variation in this dataset is very limited, we perform data augmentation by rotating the images with random angles. As shown in Tab. 1, we achieve an AUC of 0.929 and a top AUC of 0.942 with utilization of the segmentation information. Note that [10, 6] leverage on depth data (more information than segmentation since the background has infinite depth) for training and [9, 5] utilize a deformable hand model and additional datasets. Our RGB-only method outperforms other state-of-the-art meth-

ods that utilize various additional information.

5.2. Two-hand Global Pose Estimation

We first provide quantitative results on $Ego3D_s$ for all relevant subparts using the ground truth inputs for isolated analysis, then provide evaluation on the complete cascaded pipeline on both $Ego3D_s$ and $Ego3D_d$. Results for the two hands are combined by taking the average for simplicity. Extensive ablation studies are also performed.

HandSegNet. For the segmentation of the two hands, we report mean Intersection over Union (mIoU) of 0.955 and 0.962 on $Ego3D_s$ -test and $Ego3D_d$ -test respectively. For hand detection, the ground truth bounding boxes are determined by the annotated 2D joint locations. We report 2 metrics for the task of hand detection: the hand detection accuracy for how well the model correctly classifies the presence of the left and right hand; the bounding box detection accuracy for how accurately the model determines the location of the left and right hands. We report a hand detection accuracy of 1.00 and bounding box detection accuracy of 0.965 and 0.982 for $Ego3D_s$ -test and $Ego3D_d$ -test respectively, where a positive is scored when the IoU between the ground truth and the predicted bounding boxes is greater than 0.5.

PoseNet_{2D}. We compute 2D PCK using the ground truth and the predicted global 2D joint pixel locations. Fig. 7a shows the 2D PCK of $PoseNet_{2D}$ on $Ego3D_s$ -test and $Ego3D_d$ -test. Additionally, we perform ablation studies by comparing the 2D PCK of $PoseNet_{2D}$ on various settings. We show that the hand segmentation and inserted batch normalization layers both lead to noticeable improvement.

PoseNet_{3D}. We show the canonical 3D PCK of $PoseNet_{3D}$ on $Ego3D_s$ -test and $Ego3D_d$ -test in Fig. 7b. We transform the 3D canonical hand poses to the global 3D space with spherical alignment for evaluation. We also perform ablation studies by comparing the results of training using different training losses and heatmaps. We point out that \mathcal{L}_{bone} decreases the average bone length error from 3.7mm to 1.8mm despite having slightly worse AUC and EPE.

3D Global Pose Estimation. For this new task, it is nec-

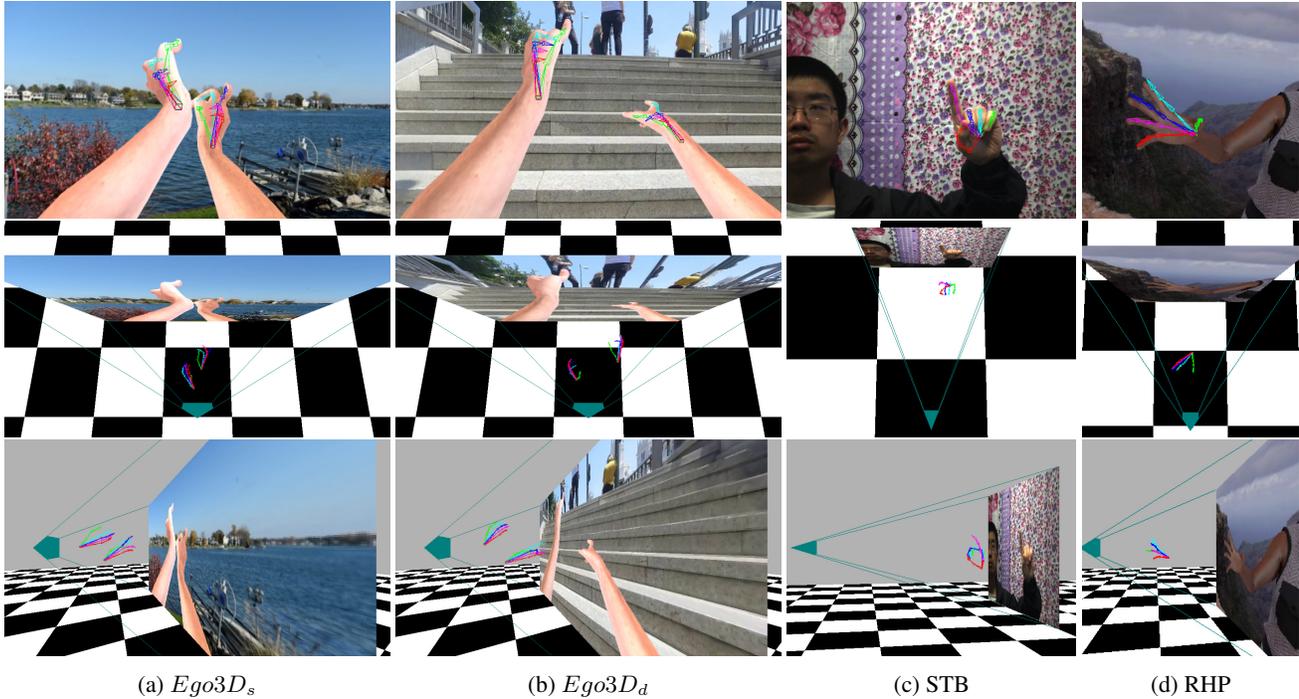


Figure 8: Qualitative results for 3D global hand pose estimation on 4 datasets. Top row visualizes the 3D global hand poses from the center camera view. Middle and bottom rows show the top and side views respectively.

essary that we evaluate the global pose estimation accuracy using PCK in the spherical coordinate system for more intuitive results. Specifically, the spherical PCK evaluates directional accuracy and distance accuracy of the root joint (mMCP). We claim that the spherical PCK on the root joint and PCK for the 3D canonical pose estimation together produce comprehensive evaluation results for the task of 3D global hand pose estimation. We skip the spherical PCK plot for isolated study since our 3D global pose estimation algorithm perfectly reconstructs the global 3D poses given the ground truth \mathbf{w}^{can} in 3D and \mathbf{p} in 2D.

The task for the complete cascaded pipeline is particularly difficult not only due to each module being dependent on the accuracy of the prior estimation, but also the fact that any error on the 2D or 3D canonical pose estimation can directly impact the global projection and decrease the final accuracy. Fig. 7c shows the spherical PCK of our complete pipeline on *Ego3D_s-test* and *Ego3D_d-test*. Note that \mathcal{L}_{proj} improved the overall spherical PCK for 3D global pose estimation despite leading to slightly worse performance in 3D canonical pose estimation.

We demonstrate that global hand pose estimation through monocular RGB input is achievable and we show promising results. For *Ego3D_s* and *Ego3D_d*, our method achieves a directional and distance accuracy of 0.90 approximately at an angular threshold of 3 degrees and a radius threshold of 7 cm respectively. Note that hand poses with differences of 7 cm in distance with respect to the camera origin show little difference visually in 2D images but there

is definitely room for improvement.

Our global estimation algorithm also generalizes for other datasets. We report the spherical PCK on STB in Fig. 6c for global pose estimation without accessing the ground truth location of the root joint. Our results on STB indicates that our method is capable of accurate global pose estimation on real-world data as well if sufficient training data is available. For RHP, we report an AUC of 0.960, 0.958 and 0.690 for the spherical PCK of θ , ϕ and radius respectively. This dataset is challenging for accurate distance estimation due to its low image resolution. We show qualitative results for 3D global hand pose estimation on 4 datasets in Fig. 8. In addition, we demonstrate global two-hand tracking on *Ego3D_d* and STB and show preliminary evaluation on real-world data in the supplementary material.

6. Conclusion

In this work we present the first method that estimates the 3D global poses for both hands given only a single RGB image. We contribute a large-scale synthetic egocentric hand pose dataset for training and evaluation of the networks. We show that our approach outperforms methods that utilize additional information for single-hand 3D canonical hand pose estimation and further achieves promising results for two-hand 3D global hand pose estimation. Evaluation on real-world data remains as necessary future work, which requires sufficiently annotated training data in the real-world domain.

References

- [1] Hololens2. <https://www.microsoft.com/en-us/hololens>, 2019.
- [2] HTC Vive. <https://developer.vive.com/us/viveport>, 2019.
- [3] LeapMotion. <https://developer.leapmotion.com/orion>, 2019.
- [4] Oculus Quest. <https://www.oculus.com/quest/>, 2019.
- [5] S. Baek, K. I. Kim, and T.-K. Kim. Pushing the Envelope for RGB-Based Dense 3D Hand Pose Estimation via Neural Rendering. *In CVPR*, 2019.
- [6] Y. Cai, L. Ge, J. Cai, and J. Yuan. Weakly-Supervised 3D Hand Pose Estimation from Monocular RGB Images. *In ECCV*, 2018.
- [7] G. Garcia-Hernando, S. Yuan, S. Baek, and T.-K. Kim. First-person Hand Action Benchmark with RGB-D Videos and 3D Hand Pose Annotations. *In CVPR*, 2017.
- [8] L. Ge, H. Liang, J. Yuan, and D. Thalmann. Robust 3D Hand Pose Estimation in Single Depth Images: from Single-View CNN to Multi-View CNNs. *In IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [9] L. Ge, Z. Ren, Y. Li, Z. Xue, Y. Wang, J. Cai, and J. Yuan. 3D Hand Shape and Pose Estimation from a Single RGB Image. *In CVPR*, 2019.
- [10] U. Iqbal, P. Molchanov, T. Breuel, J. Gall, and J. Kautz. Hand Pose Estimation via Latent 2.5D Heatmap Regression. *In ECCV*, 2018.
- [11] S. Liu, X. Zhang, J. Wangni, and J. Shi. Normalized Diver-sification. *In CVPR*, 2019.
- [12] G. Moon, J. Y. Chang, and K. M. Lee. Camera Distance-aware Top-down Approach for 3D Multi-person Pose Estimation from a Single RGB Image. *In ICCV*, 2019.
- [13] F. Mueller, F. Bernard, O. Sotnychenko, D. Mehta, S. Sridhar, D. Casas, and C. Theobalt. GANerated Hands for Real-Time 3D Hand Tracking from Monocular RGB. *In CVPR*, 2018.
- [14] F. Mueller, M. Davis, F. Bernard, O. Sotnychenko, M. Verschoor, M. A. Otaduy, D. Casas, and C. Theobalt. Real-time Pose and Shape Reconstruction of Two Interacting Hands With a Single Depth Camera. *ACM Transactions on Graphics (TOG)*, 2019.
- [15] F. Mueller, D. Mehta, O. Sotnychenko, S. Sridhar, D. Casas, and C. Theobalt. Real-time Hand Tracking under Occlusion from an Egocentric RGB-D Sensor. *In ICCV*, 2017.
- [16] M. Oberweger, P. Wohlhart, and V. Lepetit. Hands Deep in Deep Learning for Hand Pose Estimation. *arXiv preprint arXiv:1502.06807*, 2015.
- [17] I. Oikonomidis, N. Kyriazis, and A. A. Argyros. Full DOF Tracking of a Hand Interacting with an Object by Modeling Occlusions and Physical Constraints. *In ICCV, IEEE Computer Society*, pages 2088–2095, 2011.
- [18] G. Rogez, P. Weinzaepfel, and C. Schmid. LCR-Net: Localization-Classification-Regression for Human Pose. *In CVPR*, 2017.
- [19] J. Romero, D. Tzionas, and M. J. Black. Embodied Hands: Modeling and Capturing Hands and Bodies Together. *ACM Transactions on Graphics (TOG)*, 2017.
- [20] T. Simon, H. Joo, I. Matthews, and Y. Sheikh. Hand Key-point detection in single images using multiview bootstrapping. *In CVPR*, 2017.
- [21] A. Spurr, J. Song, S. Park, and O. Hilliges. Cross-modal Deep Variational Hand Pose Estimation. *In CVPR*, 2018.
- [22] S. Sridhar, F. Mueller, M. Zollhoefer, D. Casas, A. Oulasvirta, and C. Theobalt. Real-time Joint Tracking of a Hand Manipulating an Object from RGB-D Input. *In ECCV*, 2016.
- [23] S. Sridhar, A. Oulasvirta, and C. Theobalt. Interactive Markerless Articulated Hand Motion Tracking Using RGB and Depth data. *In IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2456–2463, 2013.
- [24] S. Sridhar, H. Rhodin, H.-P. Seidel, A. Oulasvirta, and C. Theobalt. Real-time Hand Tracking Using a Sum of Anisotropic Gaussians Model. *3DV, IEEE Computer Society*, pages 319–326, 2014.
- [25] X. Sun, Y. Wei, S. Liang, X. Tang, and J. Sun. Cascaded Hand Pose Regression. *In CVPR*, 2015.
- [26] D. Tang, H. J. Chang, A. Tejani, and T.-K. Kim. Latent Regression Forest: Structured Estimation of 3D Articulated Hand Posture. *In CVPR*, 2014.
- [27] J. Taylor, V. Tankovich, D. Tang, C. Keskin, D. Kim, P. Davidson, A. Kowdle, and S. Izadi. Articulated Distance Fields for Ultra-Fast Tracking of Hands Interacting. *SIG-GRAPH Asia*, 2017.
- [28] J. Tompson, M. Stein, Y. Lecun, and K. Perlin. Real-Time Continuous Pose Recovery of Human Hands Using Convolutional Networks. *ACM Transactions on Graphics*, 33, 2014.
- [29] D. Tzionas, L. Ballan, A. Srikantha, P. Aponte, M. Pollefeys, and J. Gall. Capturing Hands in Action Using Discriminative Salient Points and Physics Simulation. *IJCV*, 2016.
- [30] R. Wang, S. Paris, and J. Popovic. 6D Hands: Markerless Hand-tracking for Computer Aided Design. *In Proc. of UIST*, pages 549–558, 2011.
- [31] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional Pose Machines. *In CVPR*, 2016.
- [32] D. Xiang, H. Joo, and Y. Sheikh. Monocular Total Capture: Posing Face, Body, and Hands in the Wild. *In CVPR*, 2019.
- [33] L. Yang and A. Yao. Disentangling Latent Hands for Image Synthesis and Pose Estimation. *In CVPR*, 2019.
- [34] S. Yuan, Q. Ye, B. Stenger, S. Jain, and T.-K. Kim. Big-Hand2.2M Benchmark: Hand Pose Dataset and State of the Art Analysis. *In CVPR*, 2017.
- [35] J. Zhang, J. Jiao, M. Chen, L. Qu, X. Xu, and Q. Yang. 3D Hand Pose Tracking and Estimation Using Stereo Matching. *arXiv preprint arXiv:1610.07214*, 2016.
- [36] X. Zhou, Q. Wan, W. Zhang, X. Xue, and Y. Wei. Model-based Deep Hand Pose Estimation. *arXiv preprint arXiv:1606.06854*, 2016.
- [37] C. Zimmermann and T. Brox. Learning to Estimate 3D Hand Pose from Single RGB Images. *In ICCV*, 2017.
- [38] C. Zimmermann, D. Ceylan, J. Yang, B. Russell, M. Argus, and T. Brox. FreiHAND: A Dataset for Markerless Capture of Hand Pose and Shape From Single RGB Images. *In ICCV*, 2019.