# SeqHAND: RGB-Sequence-Based 3D Hand Pose and Shape Estimation

John Yang<sup>1\*</sup>, Hyung Jin Chang<sup>2\*</sup>, Seungeui Lee<sup>1</sup>, and Nojun Kwak<sup>1\*\*</sup>

<sup>1</sup> Seoul National University, South Korea {yjohn,seungeui.lee,nojunk}@snu.ac.kr <sup>2</sup> University of Birmingham, UK h.j.chang@bham.ac.uk

Abstract. 3D hand pose estimation based on RGB images has been studied for a long time. Most of the studies, however, have performed frame-by-frame estimation based on independent static images. In this paper, we attempt to not only consider the appearance of a hand but incorporate the temporal movement information of a hand in motion into the learning framework, which leads to the necessity of a large-scale dataset with sequential RGB hand images. We propose a novel method that generates a synthetic dataset that mimics natural human hand movements by re-engineering annotations of an extant static hand pose dataset into *pose-flows*. With the generated dataset, we train a newly proposed recurrent framework, exploiting visuo-temporal features from sequential synthetic hand images and emphasizing smoothness of estimations with temporal consistency constraints. Our novel training strategy of detaching the recurrent layer of the framework during domain finetuning from synthetic to real allows preservation of the visuo-temporal features learned from sequential synthetic hand images. Hand poses that are sequentially estimated consequently produce natural and smooth hand movements which lead to more robust estimations. Utilizing temporal information for 3D hand pose estimation significantly enhances general pose estimations by outperforming state-of-the-art methods in our experiments on hand pose estimation benchmarks.

**Keywords:** 3D Hand Pose Estimations, Pose-flow Generation, Syntheticto-real domain gap reduction, Synthetic hand motion dataset

# 1 Introduction

Since expressions of hands reflect much of human behavioral features in a daily basis, hand pose estimations are essential for many human-computer interactions, such as augmented reality (AR), virtual reality (VR) [16] and computer vision tasks that require gesture tracking [8]. Hand pose estimations conventionally struggle from an extensive space of pose articulations and occlusions including

 $<sup>^{\</sup>star}$  Equal Contribution

<sup>\*\*</sup> Corresponding Author



ated by the proposed method.

Fig. 1: Illustrations of sequential 2D im- Fig. 2: Each frame of sequential hand ages of hand pose-flows that are gener- motion videos is composed of varying poses and moving backgrounds.

self-occlusions. Some recent 3D hand pose estimators that take sequential depth image frames as inputs have tried to enhance their performance considering temporal information of hand motions [14, 22, 26, 43]. Motion context provides temporal features for narrower search space, hand personalizing, robustness to occlusion and refinement of estimations. In this paper we focus on the hand pose estimation considering its movements using only RGB image sequences for better inference of 3D spatial information.

Although the problem of estimating a hand pose in a single RGB image is an ill-posed problem, its performance is rapidly improving due to the development of various deep learning networks [3, 11, 24]. However, most studies have focused on accurately estimating 3D joint locations for each image without considering motion tendency. Pose of hands changes very quickly and in many cases contains more information on the movements of the successive poses than on the momentary ones. In addition, the current pose is greatly affected by the pose from the previous frames. Until now, there has been a lack of research on the estimation network considering the continuous changes of poses. The main reason that conventional RGB-based deep 3D hand pose estimators [1, 3, 24, 49] have only proposed frameworks with per-frame pose estimation approaches is that any large scale RGB sequential hand image dataset has not been available unlike the datasets with static images of hand poses. The diversity and the authenticity of hand motions along with generalization over skin colors, backgrounds and occlusions is a challenging factor for a dataset to be assured.

In this paper, we present a novel perspective on hand pose and shape estimation tasks and propose to consider temporal movements of hands as well as their appearances for more accurate 3D estimations of hand poses based on RGB image inputs. In order to train a framework that exploits visuo-temporal features to manage successive hand pose images, we are required to have sufficient pose data samples that are sequentially correlated. We thus propose a new generation method of dataset, SeqHAND dataset, with sequential synthetic RGB images of natural hand movements, re-engineering extant static hand pose annotations of

BigHand2.2M dataset [46].To effectively test our generated dataset, we extend the framework of [3] with a recurrent layer based on empirical validity of its structure. Also since it is widely accepted that models trained with synthetic images perform poorly on real images [24], we present a new training pipeline to preserve pre-trained image-level temporal mapping during synthetic-real domain transition. Our contributions to this end are as follows :

- We design a new generation method for sequential RGB image dataset with realistic hand motions that allows 3D hand pose and shape estimators to learn the dynamics of hand pose variations (See Figure 1) by proposing a pose-flow generation procedure.
- We propose a new recurrent framework with convolution-LSTM layer to directly exploit visuo-temporal information from hand pose and shape variations in image space and map to 3D space.
- We present a novel training pipeline of preserving extracted spatio-temporal features from sequential RGB hand images during domain finetuning from synthetic to real.
- Our approach achieves not only state-of-the-art performance in standard 3D hand pose estimation dataset benchmarks, but also smooth human-like 3D pose fittings for the image sequences.

To the best of our knowledge, we propose the first deep-learning based 3D hand pose and shape estimator without any external 2D pose estimator that exploits temporal information directly from sequential RGB images.

# 2 Related Works

Many approaches of hand pose estimation (HPE) have been actively studied. To acquire hand information, the literature of single hand 3D pose estimation has been mainly based on visual inputs of depth sensors and/or RGB cameras.

**Per-frame RGB-based 3D HPE.** As views of a single 3D scene in multiple perspectives are correlated, efforts of 3D estimation based on multiple RGB images of a hand have also been introduced [6, 12, 27, 34, 37]. Multi-view camera setups allow refinements against occlusions, segmentation enhancements and better sense of depth. In the work of [34], bootstrapping pose estimations among images from multiple perspectives help the estimator to retrain badly annotated data samples and refine against occlusions. A pair of stereo images provides similar effects in a more limited setting. Integration of paired stereo images has yielded better 3D hand pose estimations through manipulations of disparity between paired images [28, 30, 32, 48].

Monocular RGB-only setup is even more challenging because it only provides visual 2D vision of hand poses. With deep learning methods that have allowed successful achievements of hand detection [13, 19], deep pose estimators have recently been able to concentrate on per-frame hand 3D pose estimation problems [51]. To overcome the lack of 3D spatial information from the 2D inputs, there are needs of constraints and guidance to infer 3D hand postures [29]. Most recently,

works of [1,3,49] employ a prior hand model of MANO [31] and have achieved significant performance improvement in the RGB-only setup.

Temporal information in 3D HPE. Considering temporal features of depth maps, sequential data of hand pose depth images [25, 26, 46, 48] have been trained with hand pose estimators. The temporal features of hand pose variations are used for encoding temporal variations of hand poses with recurrent structure of a model [14, 43], modeling of hand shape space [17], and refinement of current estimations [22, 26]. With sequential monocular RGB-D inputs, Taylor et al. [41] optimize surface hand shape models, updating subdivision surfaces on corresponding 3D hand geometric models. Temporal feature exploitation has not been done for deep-learning based 3D hand pose estimators that take color images as inputs because large scale sequential RGB hand pose datasets have not been available in the literature. We share the essential motivation with the work of [5], but believe that, even without the assistance of 2D pose estimation results, sequential RGB images provide sufficient temporal information and spatial constraints for better 3D hand pose inference with robustness to occlusions.

**Synthetic hand data generations.** Since RGB images also consist of background noise and color diversity of hands that distract pose estimations, synthetic RGB data samples are generated from the hand model to incite the robustness of models [3, 4, 11, 25]. In [35, 45], cross-modal data is embedded in a latent space, which allows 3D pose labeling of unlabeled samples generated from (disentangled) latent factor traverses. Mueller et al. [24] had applied cycleGAN [50] for realistic appearances of generated synthetic samples to reduce the synthetic-real domain gap. While there have been recent attempts to solve an issue of lacking reliable RGB datasets through generations of hand images [3,4,24,35,51], most of the works have focused on generation of realistic appearances of hands that are not in motions. To strictly imitate human perception of hand poses, it is critical for RGB-based hand pose estimators to understand the dynamics of pose variations in a spatio-temporal space. We further consider that synthetic hand pose dataset in realistic motions provides efficient information for pose estimations as much as appearances.

### 3 SeqHAND Dataset

#### 3.1 Preliminary: MANO Hand Model

MANO hand model [31] is a mesh deformation model that takes two lowdimensional parameters  $\theta$  and  $\beta$  as inputs for controlling the pose and the shape, respectively, of the 3D hand mesh outputs. With a given mean template  $\overline{T}$ , the rigid hand mesh is defined as:

$$M(\theta, \beta) = W(T(\theta, \beta, \overline{T}), J(\beta), \theta, \omega)$$
(1)

where  $T(\cdot)$  defines the overall shape for the mesh model based on pre-defined deformation criteria with pose and shape, and  $J(\cdot)$  yields 3D joint locations using a kinematic tree.  $W(\cdot)$  represents the linear blend skinning function that is applied with blend weights  $\omega$ . MANO model may take up to 45-dimensional pose parameters  $\theta$  and 10-dimensional shape parameters  $\beta$  while the original MANO framework uses 6-dimensional PCA (principal component analysis) subspace of  $\theta$ for computational efficiency.

**2D Reprojection of MANO Hands:** The location of joints  $J(\beta)$  can be globally rotated based on the pose  $\theta$ , denoted as  $R_{\theta}$ , to obtain a hand posture P with corresponding 3D coordinates of 21 joints:

$$P = J(\theta, \beta) = R_{\theta}(J(\beta)).$$
<sup>(2)</sup>

After 3D estimations for mesh vertices  $M(\theta, \beta)$  and joints  $J(\theta, \beta)$  are computed by MANO model, in [3], 3D estimations are re-projected to 2D image plane with a weak-perspective camera model to acquire 2D estimations with a given rotation matrix  $R \in SO(3)$ , a translation  $t \in \mathbb{R}^2$  and a scaling factor  $s \in \mathbb{R}^+$ :

$$M_{2D} = s\Pi RM(\theta, \beta) + t \tag{3}$$

$$J_{2D} = s\Pi R J(\theta, \beta) + t \tag{4}$$

where  $\Pi$  represents orthographic projections. Hand mesh  $M(\theta, \beta)$  is composed of 1,538 mesh faces and defined by 3D coordinates of 778 vertices, and joint locations  $J(\theta, \beta)$  are represented by 3D coordinates of 21 joints. The re-projected 2D coordinates of  $M_{2D}$  and  $J_{2D}$  are represented in 2D locations in the image coordinates. We have utilized MANO hand model in both synthetic hand motion data generation and the proposed pose and shape estimator.

### 3.2 Generation of SeqHAND Dataset

Although the potential of temporal features have been shown promising results for 3D HPE tasks [5, 26, 41], large scale RGB sequential hand image datasets have not been available during recent years in the literature of RGB-based 3D HPE. In this section, we describe a new generation method of hand motions that consist of sequential RGB frames of synthetic hands.

To generate sequential RGB image data with human-like hand motions, all poses during the variation from an initial pose to a final pose need to be realistic. We thus utilize BigHand2.2M (BH) [46] for sequential hand motion image dataset generation. BH dataset consists of 2.2 million pose samples with 3D annotations for joint locations acquired from 2 hour-long hand motions collected from 10 real subjects. With BH datasets, the generated samples are expected to inherit the manifold of its real human hand articulation space and kinematics of real hand postures. As 3D mapping of BH samples using t-SNE [21] in Figure 3 shows, BH is known that the pose samples are densely and widely collected.

**Pose-flow generation:** We firstly define a *pose-flow*, a set of poses at each time step during the variation. Putting gradually changing poses in a sequential manner, we newly propose a *pose-flow* generation method. For each pose-flow generation, an initial and a final poses,  $P_{initial}$  and  $P_{final}$ , are independently and randomly selected from BH dataset. While varying from the initial to the





Fig. 3: 3D t-SNE visualization of 10,000 of BH data samples randomly selected. BH dataset completes a pose space that covers previously reported datasets, having a dense pool of related neighboring poses.

Fig. 4: An illustration of the pose-flow generation procedure. All poses per pose-flow are selected from the annotations of BH dataset. At each frame, a current pose is updated by the difference between the previous pose and the final pose. The pose nearest to the updated pose is then selected for the frame.

final pose during *n* frames, the coordinates of joints are updated by  $\alpha/n$  of the difference between the current coordinates and the ones of the final pose.<sup>3</sup> The update size  $\alpha$  is empirically chosen for the desirable speed of pose variations. A pose  $P_i^{BH}$  from BH dataset that is the nearest to the updated pose in terms of Euclidean distance is then newly selected as the current pose for the *k*-th frame:

$$P_0 = P_{initial}^{BH} \tag{5}$$

$$P_{updated} = P_{k-1} - \frac{\alpha}{n} (P_{k-1} - P_{final}^{BH}) \tag{6}$$

$$P_k = P_i^{BH} \quad \text{s.t.} \quad \min_i ||P_{updated} - P_i^{BH}||. \tag{7}$$

The overall procedure of the Pose-flow generation is summarized in Figure 4. The intermediate pose ( $P_{updated}$ ) is calculated as stochastic update. Such stochasticity of our pose updates helps avoiding strict updates of pose gradients and encourages wandering more within the pose space. Pose selections from the BH annotations, again, allows assurance on the authencity of hand poses during the variation.

To generate RGB images for a pose, a shallow four-layer network with which takes inputs of 3D coordinates for joints of BH annotations is trained to output corresponding pose parameters  $\theta$  for MANO hand model. For each pose at a frame, we feed corresponding 21 joint location coordinates to the this network to acquire a hand mesh model in the desired pose, which is then re-projected to an image plane. As done in [3], we assign each vertex in a mesh the RGB

<sup>&</sup>lt;sup>3</sup> Note that direct random samplings from continuous pose parameter space  $\theta \in \mathbb{R}$  does not assure diversity and authenticity of poses [31].

Table 1: Among the contemporary 3D hand pose datasets, SeqHAND dataset is the first dataset for 3D hand pose estimations that provides sequential RGB hand image frames along with stable annotations in both  $3^{rd}$ -person and egocentric perspectives.

| Datasets        | RGB/Depth | Real/Synth | Static/Sequential | $3^{rd}$ /Ego view | # of frames |
|-----------------|-----------|------------|-------------------|--------------------|-------------|
| SynthHands [25] | RGB+Depth | Synth      | Static            | Ego                | 63k         |
| RHD [51]        | RGB+Depth | Synth      | Static            | $3^{rd}$           | 43.7k       |
| FHAD [10]       | RGB+Depth | Real       | Sequential        | Ego                | 100k        |
| NYU [42]        | Depth     | Real       | Sequential        | $3^{rd}$           | 80k         |
| ICVL [40]       | Depth     | Real       | Sequential        | $3^{rd}$           | 332.5k      |
| MSRA15 [38]     | Depth     | Real       | Sequential        | $3^{rd}$           | 76,375      |
| MSRC [33]       | Depth     | Synth      | Sequential        | $3^{rd}$ +Ego      | 100k        |
| SynHand5M [23]  | Depth     | Synth      | Sequential        | $3^{rd}$           | 5M          |
| GANerated [24]  | RGB       | Synth      | Static            | Ego                | 330k        |
| SeqHAND         | RGB       | Synth      | Sequential        | $3^{rd}$ +Ego      | 410k        |

value of predefined color templates of hands to create appearances of hands. Sampled hand shape parameter  $\beta \in [-2, 2]^{10}$  and selected color template are set unchanged along per flow. Camera parameters of rotation R, scale s and translation t factors are independently sampled for initial and final poses and updated at each frame in the same way as the poses are. All frames are in the size of w and h. Figure 1 depicts illustrations of our generated pose-flows.

Further mimicking images of hand motions in the wild, we sample two (initial and ending) random patches from VOC2012 data [9] with the size of w and h and move the location of the patch for backgrounds along the frames. As Table 1 denotes, the generated SeqHAND dataset provides not only both  $3^{rd}$ -person and egocentric viewpoints of hand postures but also sequential RGB images of hand poses that firstly allow data-hungry neural networks to exploit visuo-temporal features directly from RGB inputs.<sup>4</sup>

## 4 SeqHand-Net for Visuo-Temporal Feature Exploitation

With SeqHAND dataset, we are able to overcome the scarcity of sequential RGB dataset which limits conventional RGB-based 3D HPE methods from exploiting temporal image features. Motivated by [3], we design sequential hand pose and shape estimation network (SeqHAND-Net). On top of the encoder network of [3], we incorporate convolution-LSTM (ConvLSTM) layer [44] to capture sequential relationship between consecutive hand poses. Our method does not consider additional hand 2D joint locations as inputs, and purely performs 3D hand pose estimation based on sequentially streaming RGB images in an effort to overcome the dependency on external 2D pose estimators. We also propose, in this section, a training pipeline for domain adaptation from synthetic to real, adapting low-level features with real hand images while preserving high-level visuo-temporal features of hand motions.

<sup>&</sup>lt;sup>4</sup> Although we can generate as many synthetic data as we want, our SeqHand dataset contains 400K/10K samples used for training/validation.

From each frame, a cropped hand image is fed into SeqHAND-Net as illustrated in Figure 5. Our problem scope is to better perform hand pose estimations on streaming cropped frames that are unseen by the estimator. The encoder of our SeqHAND-Net has the backbone structure of ResNet-50 [13] and, for training, expects sequential inputs with k frames. A single ConvLSTM is implemented right before the last layer as a recurrent visual feature extractor so that the dynamics of hand motions are embedded in the highest-level latent space. Learning of hand motion sequential dynamics in the high-level space is important since low-level visual features are changed with the ConvLSTM layer fixed during finetuning for real hand images. After the recurrent layer, a simple linear mapping layer from hidden features to the output vector is set.



Fig. 5: Our training strategy for preservation of temporal features during domain adaptation synth-to-real. SeqHAND-Net is created from [3] with an extra visuo-temporal feature exploitation layer for sequential RGB inputs. During finetuning, SeqHAND-Net considers a static real data as 1-frame-long sequence. The temporal high-level feature encoding is preserved while low-level image feature encoding layers are finetuned.

The encoder's resultant vector consists of parameters for pose  $\theta \in \mathbb{R}^{10}$ , shape  $\beta \in \mathbb{R}^{10}$ , scale  $s \in \mathbb{R}^+$ , translation  $t \in \mathbb{R}^2$  and rotation  $r \in \mathbb{R}^3$  which turns into a matrix  $R \in SO(3)$  through Rodrigues rotation formula for Eqs (3) and (4).

Synth-to-Real Domain Transfer with Preservation of Temporal Features As mentioned earlier, many recent researches have used synthetic hand images for pre-training and finetuned into real domain to overcome the scarcity of real hand images. While finetuning into real domain may allow faster training convergence, further training with a smaller dataset not only causes overfitting and may result in catastrophic forgetting [18]. To preserve visuo-temporal features learned from synthetic hand motions of SeqHAND dataset, we exclude the ConvLSTM layer of SeqHAND-Net from domain transfer to real hand images, allowing the network to only finetune low-level image features. Only the 'Encoder' and 'MLP' layers from Figure 5 are finetuned with a real static hand image dataset (e.g. FreiHand [52]). SeqHAND-Net is therefore trained, considering each image sample as 1-frame-long sequential image during domain transition to real.

**Training Objectives** The followings are the types of criteria used for training our proposed framework to consider visuo-temporal features and emphasize the temporal smoothness of estimations :

2D joint regression loss: The re-projected 2D joint loss is represented as :

$$L_{2D}^J = ||J_{2D} - x_{2D}^J||_1, (8)$$

where  $x_{2D}$  represents the ground-truth 2D locations of hand joints within a frame image. We have used the L1 loss because of inaccuracies in annotations in the training datasets.

**3D joint regression loss:** The ground-truth joint locations and the ones predicted are regressed to be the same using the following loss:

$$L_{3D}^{J} = ||RJ(\theta, \beta) - x_{3D}^{J}||_{2}^{2},$$
(9)

where  $x_{3D}^J$  represents ground-truth 3D joint coordinates. If a dataset provides ground-truth coordinates of 3D vertex points (e.g. FreiHand dataset), the 3D coordinates of each vertex predicted and the ones of ground-truth is minimized as done for 3D joint loss, based on the following loss:

$$L_{3D}^{M} = ||RM(\theta, \beta) - x_{3D}^{M}||_{2}^{2}$$
(10)

where  $x_{3D}^M$  represents ground-truth 3D mesh vertex coordinates.

Hand mask fitting loss: The hand mask loss is proposed in [3] to fit the shape and pose predictions in the binary mask of hands in the image plane. This loss ensures predicted coordinates of mesh vertices to be inside of a hand region when re-projected:

$$L_{mask} = 1 - \frac{1}{N} \sum_{i} H(M_{2D}^{i}), \quad H(x) = \begin{cases} 1, & \text{if } x \text{ inside a hand region.} \\ 0, & \text{otherwise.} \end{cases}$$
(11)

where H is a hand mask indicator function that tells if vertex point x is inside the hand region or not. The loss represents the percentage of vertices that are outside the region.

**Temporal consistency loss:** For pre-training on SeqHAND dataset, our method needs to be constrained with temporal consistency to ensure smoothness of pose and shape predictions. Similar to [5], we have adopted the temporal consistency loss for smoothness of temporal variation of poses:

$$L_{temp} = ||\beta_{t-1} - \beta_t||_2^2 + \lambda_{temp}^{\theta} ||\theta_{t-1} - \theta_t||_2^2.$$
(12)

While finely penalizing current estimations with the previous ones, this loss allows the reduction of search space and natural 3D hand motion estimations.

**Camera parameter regression loss:** During training with SeqHAND dataset where all ground-truths for pose, shape and viewpoint parameters  $\{\theta, \beta, r, t, s\}$  are available, our model is trained with L2-norm loss between predictions and the ground-truth.

$$L_{cam} = \sum_{i \in \{\theta, \beta, r, t, s\}} ||\hat{i} - i||_2^2$$
(13)

where  $\hat{i}$  and i respectively refer to predicted and ground-truth parameters for pose, shape and viewpoint.

## 5 Experiments

**Datasets for Training.** For visuo-temporal feature encodings of sequential RGB hand images, we pretrain SeqHAND-Net with our SeqHAND dataset. We have generated 40,000 sequences for training and 1,000 for validate samples each of which is 10-frames-long. All images are generated in the size of 224×224 for ResNet-50 input size. SeqHAND data samples are exemplified in Figure 1 and 2.

To finetune SeqHAND-Net for synthetic-real domain gap reduction, we have used STB (Stereo Hand Pose Tracking Benchmark) [48] and FR (FreiHand) [52] datasets. STB dataset consists of real hand images captured in a sequential manner during 18,000 frames with 6 different lighting conditions and backgrounds. Each frame image is labeled with 2D and 3D annotations of 21 joints. Since STB dataset has annotations for joint locations of palm centers instead of wrist, we have interpolated related mesh vertices of MANO hand model to mach the annotation of STB dataset. The dataset is divided into training and testing sets as done in [3].

FR dataset has 130,240 data samples that are made up of 32,560 non-sequential real hand images with four different backgrounds. Since the dataset has hands that are centered within the image planes, we have modified each sample by repositioning the hand randomly within the image for more robust training results. FR dataset provides MANO-friendly annotations of 21 joint 3D/2D locations along with 778 vertex ground-truth 2D/3D coordinates with hand masks.

We have finetuned the SeqHAND-Net pretrained on SeqHAND dataset with real-hand image datasets mentioned above in a non-sequential manner while conserving hand motion dynamic features detached from further learning.

**Datasets for Evaluation.** We evaluate various framework structures that consider temporal features on the validation set of SeqHAND dataset for the logical framework choice. For the comparison against other state-of-the-art methods, we have selected standard hand pose estimation datasets of the splitted test set of STB, EgoDexter(ED) [25] and Dexter+Obeject(DO) [36] in which there exists temporal relations among data samples since our network requires sequential RGB inputs for fair comparisons. While STB and DO datasets consist of real hand images in  $3^{rd}$ -person viewpoints, ED dataset has samples that are in egocentric perspective. For all datasets, our method is evaluated on every frame of input sequences.

**Metrics.** For evaluation results, we measure the percentage of correct key-points for 3D joint locations (3D-PCK) along with the area under the curve (AUC) of various thresholds. In addition, we provide average Euclidean distance error for all 2D/3D joint key-points so that more absolute comparisons can be made.

Hand Localizations. For all experiments, we have used MobileNet+SSD version of hand detection implementation [19] trained with a hand segmentation dataset [2] for providing sequential cropped hand images to SeqHAND-Net. For localized hands with tight bounding rectangular boxes, we choose the longer edge with a length size l and crop the region based on the center point of boxes so that the cropped images have a square ratio with width and height size of 2.2 \* l, as done in [3].

| Frameworks                                                                  |       | AUC   |      | (px/mm) | # parame |  |
|-----------------------------------------------------------------------------|-------|-------|------|---------|----------|--|
|                                                                             |       | 3D    | 2D   | 3D      |          |  |
| ResNet50-Encoder (baseline) [3]                                             | 0.855 | 0.979 | 3.44 | 7.85    | 28.8M    |  |
| ResNet101-Encoder [3]                                                       | 0.861 | 0.981 | 3.31 | 7.54    | 47.8M    |  |
| I3D-Encoder [7]                                                             | 0.831 | 0.967 | 4.19 | 9.24    | 31.5M    |  |
| MFNet-Encoder [20]                                                          | 0.818 | 0.912 | 5.48 | 10.54   | 41.7M    |  |
| ResNet50-Encoder+LSTM                                                       | 0.826 | 0.956 | 4.64 | 9.63    | 39.3M    |  |
| $\hline \mathbf{ResNet50}\text{-}\mathbf{Encoder}\text{+}\mathbf{ConvLSTM}$ | 0.873 | 0.986 | 3.17 | 7.18    | 43.2M    |  |

Table 2: Ablation study results of various structures of the framework proposed in [3] for sequential inputs.

Table 3: Performances of differently (partially) trained models on ED, DO, STB datasets.

| Methods                                                     | AUC   |       |       | Avg. 3D Error (mm) |        |       |
|-------------------------------------------------------------|-------|-------|-------|--------------------|--------|-------|
| Methods                                                     |       | DO    | STB   | ED                 | DO     | STB   |
| Encoder + Train(SynthHAND)                                  | 0.350 | 0.095 | 0.140 | 52.11              | 100.84 | 68.86 |
| Encoder + Train(SynthHAND) + Train(FH + STB)                | 0.397 | 0.516 | 0.985 | 49.18              | 33.12  | 9.80  |
| Encoder + ConvLSTM + Train(SeqHAND)                         | 0.373 | 0.151 | 0.121 | 52.18              | 81.51  | 71.10 |
| Encoder + ConvLSTM + Train(SeqHAND) + Train(FH + STB)       | 0.444 | 0.581 | 0.981 | 40.94              | 29.41  | 9.82  |
| $Encoder + ConvLSTM + Train(SeqHAND) + Train_{C}(FH + STB)$ | 0.766 | 0.843 | 0.978 | 17.16              | 18.12  | 9.87  |

#### 5.1 Ablation Study

**Framework Selection:** To show the logic behind the selection of the proposed framework, we evaluate various forms of extended baseline model [3] shown in Table 2 for managing sequential inputs on our newly generated SeqHAND dataset. The extended versions of baseline encoder (ResNet-50) include the baseline model with a LSTM layer [39], the baseline model with a ConvLSTM layer [44], the baseline encoder with the structure of I3D [7] and the baseline encoder with the structure of MF-Net [20]. Both I3D-Encoder and MFNet-Encoder represent methods that incorporate sequential inputs with 3D convolutional neural network. For I3D, we have changed few features from the original form of I3D so that its structure fits into the hand pose estimation task. The original backbone structure of I3D with Inception modules have changed into ResNet-50 for a fair comparison. MFNet is another examplary 3D convolution network proposed specifically for motion feature extractions. Of the candidates, the encoder with a ConvLSTM layer has performed the best.

Effectiveness of SeqHAND-Net and SeqHAND Dataset: To clarify the effectiveness of our proposed framework and our generated dataset, variations of the proposed method and the baseline model are investigated. We report AUCs of 3D PCK curves and average 3D joint location errors for ED, DO and the evaluation set of STB datasets. In Table 3, 'Encoder' denotes the baseline model with ResNet50 backbone structure while 'Encoder + ConvLSTM' denotes our proposed framework SeqHAND-Net. 'Train(SynthHAND)' and 'Train(SeqHAND)' represent training a model with synthetic hand image dataset respectively in non-sequential and sequential manner. 'Train(FH + STB)' and 'Train<sub>C</sub>(FH +

ž

Table 4: Average 3D joint distance (mm) to ground-truth for RGB Sequence datasets hand pose benchmarks.

|                     | Avg. 3D Error (mm) |       |      |  |
|---------------------|--------------------|-------|------|--|
|                     | ED                 | DO    | STB  |  |
| Our Method          | 17.16              | 18.12 | 9.87 |  |
| Bouk. et al. (RGB)  | 51.87              | 33.16 | 9.76 |  |
| Bouk. et al. (Best) | 45.33              | 25.53 | 9.76 |  |
| Spurr et al.        | 56.92              | 40.20 | -    |  |
| Zimmer. et al.      | 52.77              | 34.75 | -    |  |





STB)' refers to training with STB and FreiHand datasets for the synthetic-real domain transfer with the ConvLSTM layer, respectively, attached and detached from finetuning.

We show in the Table 3 how much performance enhancement can be obtained with SeqHAND dataset and our proposed domain adaptation strategy. Encoder with the ConvLSTM layer finetuned to real domain consequently performs similar to the encoder that does not consider visuo-temporal correlations. If the ConvLSTM layer is detached from finetuning and visuo-temporal features learned are preserved, the performance significantly improves. Also, SeqHAND dataset does not consist with any occluded hands except for self-occlusions. With training for FH dataset, our method is able to learn the visual features of not only real hands but also occluded real hands since FH dataset's augmentations consist of occlusions. Due to the temporal constraint that penalizes large difference among sequential estimations, per-frame estimation performs slightly better for the STB dataset.

## 5.2 Comparison with the State-of-the-art Methods

In Figures 6, 7 and 8, we have plotted 3D-PCK graph with various thresholds for STB, ED and DO datasets. For STB dataset, deep-learning based works



Fig. 9: Qualitative Results of SeqHAND-Net.

of [3,4,15,24,35,51] and approaches from [29,47] are compared. Many previous methods have reached near the maximum performance for STB dataset. With our temporal constraints and fixing the ConvLSTM layer during finetuning, our method reaches a competitive performance. For both ED and DO datasets, our method outperforms other methods. For ED dataset, contemporary works of [3,15,35,51] are compared to our method. The best performance of our baseline [3] is reached with inputs of RGB and 2D pose estimations provided by an external 2D pose estimator. Our method results in outstanding performance against other compared methods [3,15,24,35,51] for DO dataset with heavy occlusions, which shows that the learning of pose-flow continuity enhances robustness to occlusions. Temporal information exploitation from sequential RGB images affect our model to be robust against dynamically moving scene. For more absolute comparisons, we provide our average 3D error of joint location in Table 4.

We provide qualitative results in Figure 9 for visual comparison against a frame-by-frame 3D pose estimator, our reproduced work of [3]. All images in the figure are sequentially inputted to both estimators from left to right. Per-frame estimations that fit postures at each frame result in unnatural 3D hand posture changes over a sequence while our method's leaning trajectories biased by previous frames produces natural hand motions and robust estimations to frames that lack visual information of hand postures. Our method models estimated hand shapes as consistent as possible per sequence. During the qualitative evaluation on a RGB image sequence of a single real hand, our method's average difference among temporal changes of shape parameters  $\beta_{t-1} - \beta_t$  is  $4.16e^{-11}$  while that of the frame-by-frame estimator is  $2.38e^{-5}$ . The average difference among temporal changes of the pose parameters  $\theta_{t-1} - \theta_t$  are  $1.88e^{-6}$  for our method and  $6.90e^{-6}$  for the other.

## 6 Conclusion

In this paper, we have addressed and tackled the scarcity of sequential RGB dataset which limits conventional methods from exploiting temporal features for 3D HPE. We have proposed a novel method to generate SeqHAND dataset, a dataset with sequential RGB image frames of synthetic hand poses in motions that are interpolated from existing static pose annotations. We have then also proposed a framework that exploits visuo-temporal features for 3D hand pose estimations in a recurrent manner. Our proposed method outperforms other existing approaches that take RGB-only inputs that are based on solely appearance-based methods, and consequently produces pose-flow estimations that mimick natural movements of human hands. We plan to enable the framework to solve (self-)occlusion problems more robustly.

Acknowledgement This work was supported by IITP grant funded by the Korea government (MSIT) (No.2019-0-01367, Babymind) and Next-Generation Information Computing Development Program through the NRF of Korea (2017M3C4A7077582).

## References

- Baek, S., Kim, K.I., Kim, T.K.: Pushing the envelope for RGB-based dense 3D hand pose estimation via neural rendering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1067–1076 (2019)
- Bambach, S., Lee, S., Crandall, D.J., Yu, C.: Lending a hand: Detecting hands and recognizing activities in complex egocentric interactions. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1949–1957 (2015)
- Boukhayma, A., Bem, R.d., Torr, P.H.: 3d hand shape and pose from images in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10843–10852 (2019)
- Cai, Y., Ge, L., Cai, J., Yuan, J.: Weakly-supervised 3D hand pose estimation from monocular RGB images pp. 666–682 (2018)
- 5. Cai, Y., Ge, L., Liu, J., Cai, J., Cham, T.J., Yuan, J., Thalmann, N.M.: Exploiting spatial-temporal relationships for 3D pose estimation via graph convolutional networks (2019)
- de Campos, T.E., Murray, D.W.: Regression-based hand pose estimation from multiple cameras 1, 782–789 (2006)
- Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017)
- Chang, H.J., Garcia-Hernando, G., Tang, D., Kim, T.K.: Spatio-temporal hough forest for efficient detection–localisation–recognition of fingerwriting in egocentric camera. Computer Vision and Image Understanding 148, 87 – 96 (2016)
- Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PAS-CAL Visual Object Classes Challenge 2012 (VOC2012) Results. http://www.pascalnetwork.org/challenges/VOC/voc2012/workshop/index.html
- Garcia-Hernando, G., Yuan, S., Baek, S., Kim, T.K.: First-person hand action benchmark with rgb-d videos and 3d hand pose annotations. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 409–419 (2018)
- Ge, L., Ren, Z., Li, Y., Xue, Z., Wang, Y., Cai, J., Yuan, J.: 3D hand shape and pose estimation from a single RGB image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10833–10842 (2019)
- Gomez-Donoso, F., Orts-Escolano, S., Cazorla, M.: Large-scale multiview 3D hand pose dataset. arXiv preprint arXiv:1707.03742 (2017)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Hu, Z., Hu, Y., Liu, J., Wu, B., Han, D., Kurfess, T.: A crnn module for hand pose estimation. Neurocomputing **333**, 157–168 (2019)
- Iqbal, U., Molchanov, P., Breuel Juergen Gall, T., Kautz, J.: Hand pose estimation via latent 2.5 d heatmap regression pp. 118–134 (2018)
- Jang, Y., Noh, S., Chang, H.J., Kim, T., Woo, W.: 3D finger cape: Clicking action and position estimation under self-occlusions in egocentric viewpoint. IEEE Transactions on Visualization and Computer Graphics 21(4), 501–510 (April 2015)
- 17. Khamis, S., Taylor, J., Shotton, J., Keskin, C., Izadi, S., Fitzgibbon, A.: Learning an efficient model of hand shape variation from depth images pp. 2540–2548 (2015)
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming catastrophic forgetting in neural networks. Proceedings of the national academy of sciences 114(13), 3521–3526 (2017)

- 16 J. Yang et al.
- Le, T.H.N., Quach, K.G., Zhu, C., Duong, C.N., Luu, K., Savvides, M.: Robust hand detection and classification in vehicles and in the wild. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 1203–1210. IEEE (2017)
- Lee, M., Lee, S., Son, S., Park, G., Kwak, N.: Motion feature network: Fixed motion filter for action recognition. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 387–403 (2018)
- Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research 9(Nov), 2579–2605 (2008)
- Madadi, M., Escalera, S., Carruesco, A., Andujar, C., Baró, X., Gonzàlez, J.: Topdown model fitting for hand pose recovery in sequences of depth images. Image and Vision Computing 79, 63–75 (2018)
- Malik, J., Elhayek, A., Nunnari, F., Varanasi, K., Tamaddon, K., Heloir, A., Stricker, D.: Deephps: End-to-end estimation of 3d hand pose and shape by learning from synthetic depth. In: 2018 International Conference on 3D Vision (3DV). pp. 110–119. IEEE (2018)
- Mueller, F., Bernard, F., Sotnychenko, O., Mehta, D., Sridhar, S., Casas, D., Theobalt, C.: Ganerated hands for real-time 3D hand tracking from monocular RGB pp. 49–59 (2018)
- Mueller, F., Mehta, D., Sotnychenko, O., Sridhar, S., Casas, D., Theobalt, C.: Real-time hand tracking under occlusion from an egocentric RGB-d sensor pp. 1284–1293 (2017)
- Oberweger, M., Riegler, G., Wohlhart, P., Lepetit, V.: Efficiently creating 3D training data for fine hand pose estimation pp. 4957–4965 (2016)
- Oikonomidis, I., Kyriazis, N., Argyros, A.A.: Full dof tracking of a hand interacting with an object by modeling occlusions and physical constraints pp. 2088–2095 (2011)
- Panteleris, P., Argyros, A.: Back to RGB: 3D tracking of hands and hand-object interactions based on short-baseline stereo pp. 575–584 (2017)
- Panteleris, P., Oikonomidis, I., Argyros, A.: Using a single RGB frame for real time 3D hand pose estimation in the wild. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 436–445. IEEE (2018)
- Remilekun Basaru, R., Slabaugh, G., Alonso, E., Child, C.: Hand pose estimation using deep stereovision and markov-chain monte carlo pp. 595–603 (2017)
- Romero, J., Tzionas, D., Black, M.J.: Embodied hands: Modeling and capturing hands and bodies together. ACM Transactions on Graphics (TOG) 36(6), 245 (2017)
- Rosales, R., Athitsos, V., Sigal, L., Sclaroff, S.: 3D hand pose reconstruction using specialized mappings 1, 378–385 (2001)
- 33. Sharp, T., Keskin, C., Robertson, D., Taylor, J., Shotton, J., Kim, D., Rhemann, C., Leichter, I., Vinnikov, A., Wei, Y., et al.: Accurate, robust, and flexible real-time hand tracking. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems. pp. 3633–3642 (2015)
- Simon, T., Joo, H., Matthews, I., Sheikh, Y.: Hand keypoint detection in single images using multiview bootstrapping pp. 1145–1153 (2017)
- Spurr, A., Song, J., Park, S., Hilliges, O.: Cross-modal deep variational hand pose estimation pp. 89–98 (2018)
- 36. Sridhar, S., Mueller, F., Zollhöfer, M., Casas, D., Oulasvirta, A., Theobalt, C.: Real-time joint tracking of a hand manipulating an object from RGB-d input. In: European Conference on Computer Vision. pp. 294–310. Springer (2016)

- 37. Sridhar, S., Rhodin, H., Seidel, H.P., Oulasvirta, A., Theobalt, C.: Real-time hand tracking using a sum of anisotropic gaussians model **1**, 319–326 (2014)
- Sun, X., Wei, Y., Liang, S., Tang, X., Sun, J.: Cascaded hand pose regression. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 824–832 (2015)
- Sundermeyer, M., Schlüter, R., Ney, H.: Lstm neural networks for language modeling. In: Thirteenth annual conference of the international speech communication association (2012)
- 40. Tang, D., Jin Chang, H., Tejani, A., Kim, T.K.: Latent regression forest: Structured estimation of 3D articulated hand posture pp. 3786–3793 (2014)
- Taylor, J., Stebbing, R., Ramakrishna, V., Keskin, C., Shotton, J., Izadi, S., Hertzmann, A., Fitzgibbon, A.: User-specific hand modeling from monocular depth sequences. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 644–651 (2014)
- Tompson, J., Stein, M., Lecun, Y., Perlin, K.: Real-time continuous pose recovery of human hands using convolutional networks. ACM Transactions on Graphics (ToG) 33(5), 169 (2014)
- 43. Wu, Y., Ji, W., Li, X., Wang, G., Yin, J., Wu, F.: Context-aware deep spatiotemporal network for hand pose estimation from depth images. IEEE transactions on cybernetics (2018)
- 44. Xingjian, S., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.K., Woo, W.c.: Convolutional lstm network: A machine learning approach for precipitation nowcasting. In: Advances in neural information processing systems. pp. 802–810 (2015)
- 45. Yang, L., Yao, A.: Disentangling latent hands for image synthesis and pose estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9877–9886 (2019)
- 46. Yuan, S., Ye, Q., Stenger, B., Jain, S., Kim, T.K.: Bighand2. 2m benchmark: Hand pose dataset and state of the art analysis pp. 4866–4874 (2017)
- 47. Zhang, J., Jiao, J., Chen, M., Qu, L., Xu, X., Yang, Q.: 3D hand pose tracking and estimation using stereo matching. arXiv preprint arXiv:1610.07214 (2016)
- 48. Zhang, J., Jiao, J., Chen, M., Qu, L., Xu, X., Yang, Q.: A hand pose tracking benchmark from stereo matching pp. 982–986 (2017)
- Zhang, X., Li, Q., Zhang, W., Zheng, W.: End-to-end hand mesh recovery from a monocular RGB image. arXiv preprint arXiv:1902.09305 (2019)
- Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017)
- 51. Zimmermann, C., Brox, T.: Learning to estimate 3D hand pose from single RGB images pp. 4903–4911 (2017)
- Zimmermann, C., Ceylan, D., Yang, J., Russell, B., Argus, M., Brox, T.: Freihand: A dataset for markerless capture of hand pose and shape from single RGB images. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 813–822 (2019)