

Temporal-Aware Self-Supervised Learning for 3D Hand Pose and Mesh Estimation in Videos

Liangjian Chen¹, Shih-Yao Lin², Yusheng Xie^{*3}, Yen-Yu Lin⁴, and Xiaohui Xie¹

¹University of California, Irvine, ²Tencent America, ³Amazon, ⁴National Chiao Tung University, {liangjc2,xhx}@ics.uci.edu, mike.lin@ieee.org, yushx@amazon.com, lin@cs.nctu.edu.tw

Abstract

Estimating 3D hand pose directly from RGB images is challenging but has gained steady progress recently by training deep models with annotated 3D poses. However annotating 3D poses is difficult and as such only a few 3D hand pose datasets are available, all with limited sample sizes. In this study, we propose a new framework of training 3D pose estimation models from RGB images without using explicit 3D annotations, i.e., trained with only 2D information. Our framework is motivated by two observations: 1) Videos provide richer information for estimating 3D poses as opposed to static images; 2) Estimated 3D poses ought to be consistent whether the videos are viewed in the forward order or reverse order. We leverage these two observations to develop a self-supervised learning model called temporal-aware self-supervised network (TASSN). By enforcing temporal consistency constraints, TASSN learns 3D hand poses and meshes from videos with only 2D keypoint position annotations. Experiments show that our model achieves surprisingly good results, with 3D estimation accuracy on par with the state-of-the-art models trained with 3D annotations, highlighting the benefit of the temporal consistency in constraining 3D prediction models.

1. Introduction

3D hand estimation is an important research topic in computer vision due to a wide range of potential applications, such as sign language translation [45], robotics [1], movement disorder detection and monitoring, and human-computer interaction (HCI) [29, 18, 28].

Depth sensors and RGB cameras are popular devices for collecting hand data. However, depth sensors are not as widely available as RGB cameras and are much more expensive, which has limited the applicability of hand pose estimation methods developed upon depth images. Recent research interests have shifted toward estimating 3D hand poses directly from RGB images by utilizing color, texture,

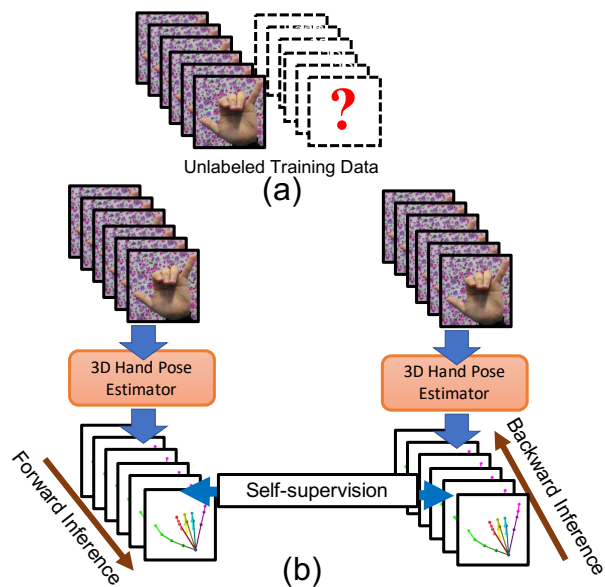


Figure 1. Motivation and idea: (a) Training a robust 3D hand pose estimator from RGB images relies on plenty images with 3D hand pose annotations, but obtaining 3D annotations from 2D images is quite difficult; (b) We leverage bi-directional temporal consistency in videos and enable hand pose estimators to make more plausible predictions. It turns out that the hand pose estimator can be derived in a self-supervised fashion without using 3D annotations.

and shape information contained in RGB images. Some methods carried out 3D hand pose estimation from monocular RGB images [5, 20, 48]. More recently, progresses have been made on estimating 3D hand shape and mesh from RGB images [2, 3, 15, 47, 42, 7, 26, 25]. Compared to poses, hand meshes provide richer information required by many immersive VR and AR applications. Despite the advances, 3D hand pose estimation remains a challenging problem due to the lack of accurate, large-scale 3D pose annotations.

In this work, we develop a new approach to 3D hand

*Work done outside of Amazon.

pose and mesh estimation by taking the following two observations into account. First, most existing methods rely on training data with 3D information, but capturing 3D information from 2D images is intrinsically difficult. Although there are a few datasets providing annotated 3D hand joints, the amount is too small to train a robust hand pose estimator. Second, most studies focus on hand pose estimation from a single image. Nevertheless, important applications based on 3D hand poses, such as augmented reality (AR), virtual reality (VR), and sign language recognition, are usually carried out in videos.

According to the two observations, our approach exploits video temporal consistency to address the uncertainty caused by the lack of 3D joint annotations on training data. Specifically, our approach, called *temporal-aware self-supervised network (TASSN)*, can learn and infer 3D hand poses without using annotated 3D training data. Figure 1 shows the motivation and core idea of the proposed TASSN. TASSN explores video information by embedding a temporal structure to extract spatial-temporal features. We design a novel temporal self-consistency loss, which helps training the hand pose estimator without requiring annotated 3D training data. In addition to poses, we estimate hand meshes since meshes bring salient evidences for pose inference. With meshes, we can infer silhouettes to further regularize our model. The main contributions of this work are given below:

1. We develop a temporal consistency loss and a reversed temporal information technique for extracting spatio-temporal features. To the best of our knowledge, this work makes the first attempt to estimate 3D hand poses and meshes without using 3D annotations.
2. An end-to-end trainable framework, named temporal-aware self-supervised networks (TASSN), is proposed to learn an estimator without using annotated 3D training data. The learned estimator can jointly infer the 3D hand poses and meshes from video.
3. Our model achieves high accuracy with 3D prediction performance on par with state-of-the-art models trained with 3D ground truth.

2. Related Work

2.1. 3D Hand Pose Estimation from Depth Images

Since depth images contain surface geometry information of hands, they are widely used for hand pose estimation in the literature [40, 44, 11, 41, 14, 16, 27, 8, 9]. Most existing work adopts regression to fit the parameters of a deformed hand model [30, 22, 24, 40]. Recent work [14, 16] extracts depth image features and regress the joints through PointNet [34]. Wu *et al.* [41] leverage the depth image as

the intermediate guidance and conduct an end-to-end training framework. Despite the effectiveness, the aforementioned methods highly rely on accurate depth maps, and are less practical in the daily life since depth sensors are not available in many cases due to the high cost.

2.2. 3D Hand Pose Estimation from RGB Images

Owing to the wide accessibility of RGB cameras, estimating 3D hand poses from monocular images becomes an active research topic [5, 20, 31, 38, 43, 48] and significant improvement has been witnessed. These methods use convolutional neural networks (CNN) to extract features from RGB images. Zimmermann and Brox [48] feed these features to the 3D lift network and camera parameter estimation network for depth regression. Building on Zimmermann and Brox’s work, Iqbal *et al.* [20] add depth maps as intermediate guidance while Cai *et al.* [5] propose a weakly supervised approach to utilize depth maps for regularization. However, these methods suffer from limited training data since 3D hand annotations are hard to acquired. Also, they all dismiss the temporal information.

2.3. 3D Hand Mesh Estimation

3D hand mesh estimation is an active research topic [15, 3, 2, 21, 47]. Methods in [3, 2, 47] estimate hand meshes by using a pre-defined hand model, named MANO [35]. Due to the high degree of freedom of hand gestures, hand meshes lie in a high dimensional space. The MANO model serves as a kinematic and shape prior of meshes and can help reduce the dimension. However, since MANO is a linear model, it is not able to capture the nonlinear transformation for hand meshes [15]. Thus, mesh estimators based on MANO suffer from this issue. On the other hand, Ge *et al.* [15] regress 3D mesh vertices through graphical convolutional neural network (GCN) with down-sampling and up-sampling. Their work achieves the state-of-the-art performance, but it is trained on a dataset with 3D mesh ground truth which is even more difficult to label than 3D joint annotations. This drawback limits its applicability in practice.

2.4. Self-supervised Learning

Self-supervised learning [12, 33, 13] is a type of training methodologies, where training data are automatically labeled by exploiting existing information within the data. With this training scheme, manual annotations are not required for a given training set. This scheme is especially beneficial when data labeling is difficult or the data size is exceedingly large. Self-supervised learning has been applied to hand pose estimation. Similar to ours, the method in [13] adopts temporal cycle consistency for self-supervised learning. However, this method uses soft nearest neighbors to solve the video alignment problem, which is not applicable to 3D pose and mesh estimation. Simon *et*

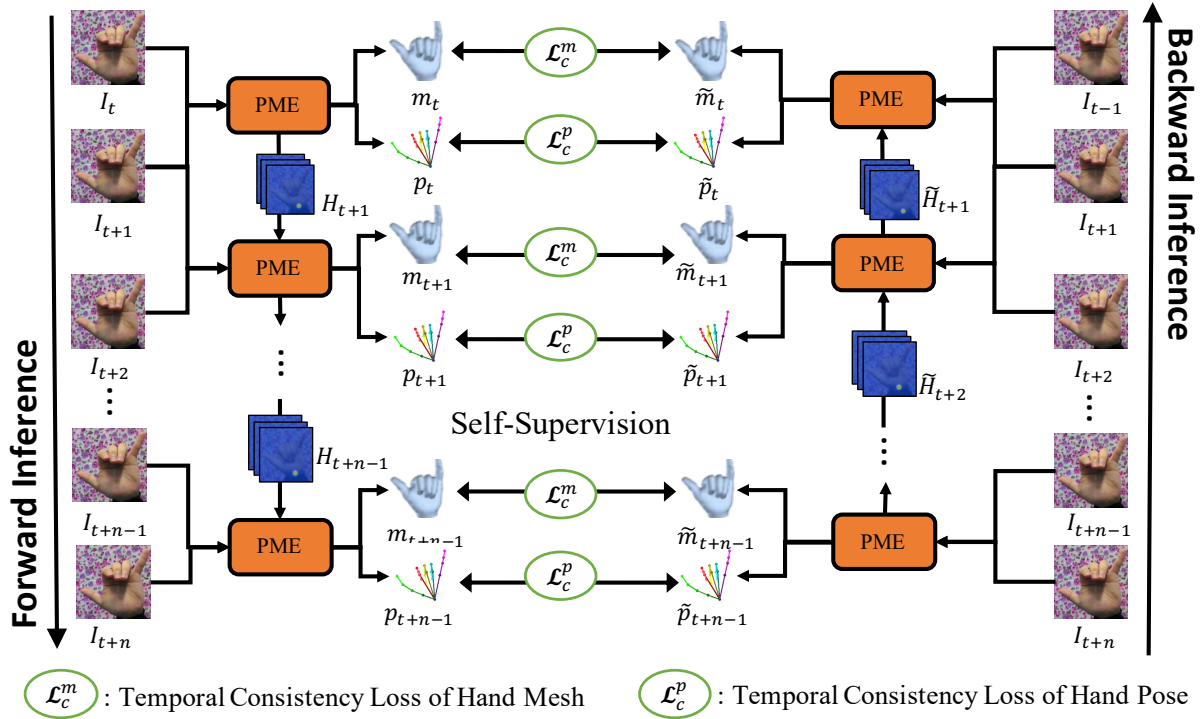


Figure 2. Overview of the proposed TASSN. TASSN involves both forward and backward inference to utilize temporal information. Namely, the hand poses estimated by forward and backward inference should be consistent. Our hand pose estimator leverages this observation and can be trained by using self-supervised learning without the need of 3D hand joint labels. Moreover, with the constraints of temporal consistency, either forward or backward inference can gain more accurate hand pose estimation results.

al. [37] adopt multi-view supervisory signals to regress 3D hand joint locations. While their approach resolves the hand self-occlusion issue using multi-view images, it in the training stage requires 3D joint annotations, which are difficult and expensive to get in this task. Another attempt of using self-supervised learning for hand pose estimation is presented in [39], where an approach leveraging a massive amount of unlabeled depth images is proposed. However, this approach may be limited due to the high variations of depth maps in diverse poses, scales, and sensing devices. Instead of leveraging multi-view consistency or depth consistency, the proposed self-supervised scheme relies on temporal consistency, which is inexpensive to get and does not require 3D keypoint annotations.

3. Proposed Method

We aim to train a 3D hand pose estimator from videos without 3D hand joint labels. To tackle the absence of 3D annotations, we adopt the temporal information from hand motion videos, and address the ambiguity caused by the lack of 3D joint ground truth. Specifically, we present a novel deep neural network, named temporal-aware self-supervised networks (TASSN). By developing the temporal consistency loss on the estimated hand gestures in a video,

TASSN can learn and infer 3D hand poses through self-supervised learning without using any 3D annotations.

3.1. Overview

Given an RGB hand motion video \mathbf{x} with N frames, $\mathbf{x} = \{I_1, \dots, I_N\}$, we aim at estimating 3D hand poses in this video, where $I_t \in \mathbb{R}^{3 \times W \times H}$ is the t -th frame, and W and H are the frame width and height, respectively. The 3D hand pose at frame t , $\mathbf{p}_t \in \mathbb{R}^{3 \times K}$, is represented by a set of K 3D keypoint coordinates of the hand. Figure 2 illustrates the network architecture of TASSN.

Leveraging the temporal consistency properties of videos, the hand poses and meshes predicted in the forward and backward inference orders can perform mutual supervision. Our model can be fine-tuned on any target dataset using this self-supervised learning and the temporal consistency is a good substitute for the hard-to-get 3D ground truth. TASSN alleviates the burden of annotating 3D ground-truth of a dataset without significantly sacrificing model performance.

Recent studies [15, 47] show that training pose estimators with hand meshes improves the performance because hand meshes can act as intermediate guidance for hand pose prediction. To this end, we propose a hand pose and mesh

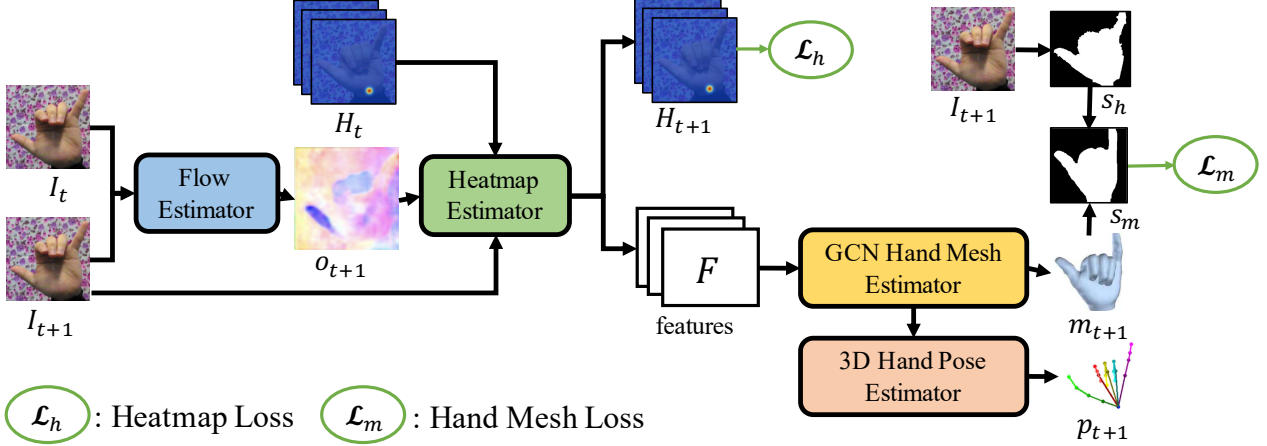


Figure 3. Network architecture of the pose and mesh estimation (PME) module. PME module consists of four sub-modules, including the flow, 2D keypoint heatmap, 3D hand mesh, and 3D hand pose estimators. The flow estimator computes the optical flow \mathbf{o}_{t+1} from two consecutive frames \mathbf{I}_t and \mathbf{I}_{t+1} . With \mathbf{I}_{t+1} , \mathbf{o}_{t+1} , and \mathbf{H}_t , the 2D heatmap estimator computes the keypoint heatmap \mathbf{H}_{t+1} at timestamp $t + 1$, as well as extract the image features. Based on the extracted image features, the 3D hand pose and mesh estimators predict the 3D hand pose \mathbf{p}_{t+1} and mesh \mathbf{m}_{t+1} at timestamp $t + 1$. Two loss terms, the heatmap loss \mathcal{L}_h and the hand mesh loss \mathcal{L}_m , are used for optimization.

estimation (PME) module, which jointly estimates the 2D hand keypoint heatmaps, 3D hand poses and meshes from every two adjacent frames \mathbf{I}_i and \mathbf{I}_{i+1} .

3.2. Pose and Mesh Estimation Module

The proposed PME module consists of four estimator sub-modules, including flow estimator, 2D keypoint heatmap estimator, 3D hand mesh estimator, and 3D hand pose estimator. Given two consecutive frames as input, it estimates the 3D hand pose and mesh. Figure 3 shows its network architecture.

Flow Estimator: To capture temporal clues from a hand gesture video, we adopt FlowNet [19] to estimate the optical flow $\mathbf{o}_{t+1} \in \mathbb{R}^{2 \times W \times H}$ between two consecutive frames \mathbf{I}_t and \mathbf{I}_{t+1} . In forward inference, FlowNet computes \mathbf{o}_{t+1} , the motion from frame \mathbf{I}_t to frame \mathbf{I}_{t+1} . In backward inference, FlowNet computes the reverse motion.

Heatmap Estimator: Our heatmap estimator computes 2D hand keypoints and generates the features for the 3D hand pose and mesh estimators. The estimated 2D keypoint heatmaps are denoted by $\mathbf{H} \in \mathbb{R}^{K \times W \times H}$, where K represents the number of keypoints. We adopt a two stacked hourglass network [32] to infer the hand keypoint heatmaps \mathbf{H} and compute the features \mathbf{F} . We concatenate \mathbf{I}_{t+1} , \mathbf{o}_{t+1} , and \mathbf{H}_t as input to the stacked hourglass network, which produces heatmaps \mathbf{H}_{t+1} , as shown in Figure 3. The estimated \mathbf{H}_{t+1} includes K heatmaps $\{\mathbf{H}_{t+1}^k \in \mathbb{R}^{W \times H}\}_{k=1}^K$, where \mathbf{H}_{t+1}^k expresses the confidence map of the location of the k th keypoint. The ground truth heatmap $\bar{\mathbf{H}}_{t+1}^k$ is the Gaussian blur of the Dirac- δ distribution centered at the ground truth location of the k th keypoint. The heatmap loss

\mathcal{L}_h at frame t is defined by

$$\mathcal{L}_h = \frac{1}{K} \sum_{k=1}^K \|\mathbf{H}_t^k - \bar{\mathbf{H}}_t^k\|_F^2. \quad (1)$$

3D Hand Mesh Estimator: Our 3D hand mesh estimator is developed based on Chebyshev spectral graph convolution network (GCN) [15], and it takes hand features \mathbf{F} as input and infers the 3D hand mesh. The output hand mesh $\mathbf{m}_t \in \mathbb{R}^{3 \times C}$ is represented by a set of 3D mesh vertices, where C is the number of vertices in a hand mesh.

To model hand mesh, we use an undirected graph $\mathcal{G}(\mathbf{V}, \mathbf{E})$, where \mathbf{V} and \mathbf{E} are the vertex and edge sets, respectively. The edge set \mathbf{E} can be represented by an adjacent matrix \mathbf{A} , where $\mathbf{A}_{i,j} = 1$ if edge $e(i, j) \in \mathbf{E}$, otherwise $\mathbf{A}_{i,j} = 0$. The normalized Laplacian normal matrix of \mathcal{G} is obtained via $\mathbf{L} = \mathbf{I} - \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$, where \mathbf{D} is the degree matrix and \mathbf{I} is the identity matrix. Since \mathbf{L} is a positive semi-definite matrix [4], it can be decomposed as $\mathbf{L} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T$, where $\mathbf{\Lambda} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_C)$, and C is the number of vertices in \mathcal{G} .

We follow the setting in [10], and set the convolution kernel to $\hat{\mathbf{A}} = \text{diag}(\sum_{i=0}^S \alpha_i \lambda_1^i, \dots, \sum_{i=0}^S \alpha_i \lambda_C^i)$, where α is the kernel parameter. The convolutional operations in \mathcal{G} can be calculated by $\mathbf{F}' = \mathbf{U} \hat{\mathbf{A}} \mathbf{U}^T \mathbf{F} \boldsymbol{\theta}_i = \sum_{i=0}^S \alpha_i \mathbf{L}^i \mathbf{F} \boldsymbol{\theta}_i$, where $\mathbf{F} \in \mathbb{R}^{N \times F_{in}}$ and $\mathbf{F}' \in \mathbb{R}^{N \times F_{out}}$ indicate the input and output features respectively, S is a preset hyperparameter used to control the receptive field, and $\boldsymbol{\theta}_i \in \mathbb{R}^{F_{in} \times F_{out}}$ is trainable parameter set used to control the number of output channels.

The Chebyshev polynomial is used to reduce the model complexity by approximating convolution operations, leading to the output features $\mathbf{F}' = \sum_{i=0}^S \alpha_i T_i(\hat{\mathbf{L}}) \boldsymbol{\theta}_i$ where




Figure 4. Examples of our estimated silhouettes. The first and third rows show the training images in STB and MHP datasets, respectively. The second and the fourth rows show the estimated silhouettes by our method.

$T_k(x)$ is the k -th Chebyshev polynomial and $\hat{\mathbf{L}} = 2\mathbf{L}/\lambda_{max} - \mathbf{I}$ is used to normalize the input features.

We adopt the scheme in [10, 15] to construct the hand mesh in a coarse-to-fine manner. We use the multi-level clustering algorithm for coarsening the graph, and then store the graph at each level and the mapping between graph nodes in every two consecutive levels. In forward inference, the GCN first up-samples the node features according to the stored mappings and graphs and then performs the graph convolutional operations.

Mesh Silhouette Constraint: In our model, without 3D mesh ground truth, the model tends to collapse to any kind of mesh as long as it is temporally consistent. To avoid this issue, we introduce the mesh loss \mathcal{L}_m to calculate the difference between the silhouette of the predicted hand mesh s_t and the ground-truth silhouette \bar{s}_t at frame t . The silhouette loss is defined by

$$\mathcal{L}_m = \|s_t - \bar{s}_t\|_F^2. \quad (2)$$

To obtain \bar{s}_t , we use GrabCut [36] to estimate the hand silhouettes from the training images. Some silhouettes estimated from training images are shown in Figure 4. The silhouette of our predicted hand mesh s_t is obtained by using the neural rendering approach in [23].

3D Hand Pose Estimator: The proposed 3D pose estimator directly infers 3D hand keypoints \mathbf{p}_t from the predicted hand mesh \mathbf{m}_t . Taking the mesh as the input, we adopt a network of two stacked GCNs, which has a similar structure to that used in 3D hand mesh estimator. We add a pooling layer to each GCN to extract the pose features from the mesh. Those pose features are then fed to two fully connected layers to regress the 3D hand pose \mathbf{p}_t .

3.3. Temporal Consistency Loss

Due to the lack of 3D keypoint annotations, conventional supervised learning schemes no longer work in model training. We propose a temporal consistency loss \mathcal{L}_c to solve this problem. Figure 2 shows the idea of our approach. Given a video clip with n frames, we feed every two adjacent frames $\{\mathbf{I}_i, \mathbf{I}_{i+1}\}_{i=t}^{t+n}$ to PME module for hand mesh and pose estimation, *i.e.*, $\{\mathbf{p}_i, \mathbf{m}_i\}_{i=t}^{t+n}$. TASSN analyzes the temporal information according to their relative input orders. Thus, we can reverse the input order from $\{\mathbf{I}_i, \mathbf{I}_{i+1}\}$ to $\{\mathbf{I}_{i+1}, \mathbf{I}_i\}$ to infer the pose and mesh in \mathbf{I}_i from \mathbf{I}_{i+1} . With this reversed temporal measurement (RTM) technique, we can infer the hand pose and mesh from the reversed temporal order. We denote the estimated pose and mesh in the reversed order as $\{\tilde{\mathbf{p}}_i, \tilde{\mathbf{m}}_i\}_{i=t}^{t+n}$. As shown in Figure 2, the prediction results estimated by the PME module in both forward and backward inference must be consistent with each other since the same mesh and pose are estimated at any frame. The temporal consistency loss on hand pose \mathcal{L}_c^p and mesh \mathcal{L}_c^m can be computed by

$$\mathcal{L}_c^p = \frac{1}{n} \sum_{i=t}^{t+n} \|\mathbf{p}_i - \tilde{\mathbf{p}}_i\|_F^2, \quad (3)$$

$$\mathcal{L}_c^m = \frac{1}{n} \sum_{i=t}^{t+n} \|\mathbf{m}_i - \tilde{\mathbf{m}}_i\|_F^2. \quad (4)$$

The temporal consistency loss \mathcal{L}_c is defined as the summation of \mathcal{L}_c^m and \mathcal{L}_c^p , *i.e.*,

$$\mathcal{L}_c = \lambda^m \mathcal{L}_c^m + \lambda^p \mathcal{L}_c^p, \quad (5)$$

where λ^m and λ^p are the weights of the corresponding losses.

3.4. TASSN Training

Suppose we are given an unlabeled hand pose dataset \mathbf{X} for training, which contains M hand gesture videos, $\mathbf{X} = \{\mathbf{x}^{(i)}\}_{i=1}^M$, where video $\mathbf{x}^{(i)} = \{I_1, \dots, I_N\}$ consists of N frames. We divide each training video into several video clips. Each training video clip \mathbf{v} is with n frames, *i.e.*, $\mathbf{v} = \{I_t, I_{t+1}, \dots, I_{t+n}\}$. With the losses defined in Eq. (1), Eq. (2), and Eq. (5), the objective for training the proposed TASSN is

$$\mathcal{L} = \lambda_s \mathcal{L}_m + \lambda_h \mathcal{L}_h + \mathcal{L}_c, \quad (6)$$

where λ^s and λ^h denote the weights of the loss \mathcal{L}_m and the loss \mathcal{L}_h , respectively. The details of parameter setting are given in the experiments.

4. Experiments Setting

4.1. Datasets for Evaluation

We evaluate our approach on two hand pose datasets, Stereo Tracking Benchmark Dataset (STB) [46] and Multi-



Figure 5. Some examples of the two hand pose datasets used for evaluation. The first row shows examples of the STB dataset [46] and the second row gives examples of the MHP dataset [17]. Both datasets include real hand video sequences performed by different subjects and have 3D hand keypoint annotations

view 3D Hand Pose dataset (MHP) [17]. These two datasets include real hand video sequences performed by different subjects and 3D hand keypoint annotations are provided for the hand video sequences.

For the STB dataset, we adopt its SK subset for training and evaluation. This subset contains 6 hand videos, each of which has 1,500 frames. Following the train-validation split setting used in [15], we use the first hand video as the validation set and the rest videos for training.

The MHP dataset includes 21 hand motion videos. Each video provides hand color images and different kinds of annotations for each sample, including the bounding box and the 2D and 3D location on the hand keypoints.

The following scheme of data pre-processing is applied to both STB and MHP datasets. We crop the hand from the original image by using the center of hand and the scale of the hand. Thus, the center of the hand is located at the center of the cropped images, and the cropped image covers the whole hand. We then resize the cropped image to 256×256 . As mentioned in [5, 48], the STB and MHP datasets use the palm center as the center of the hand. We use the mechanism introduced by [5] to change the center of hand from the palm center to the joint of wrist.

4.2. Metric

We follow the setting adopted in previous work [48, 15] and use *average End-Point-Error* (EPE) and *Area Under the Curve* (AUC) on the *Percentage of Correct Keypoints* (PCK) between threshold 20 millimeter (mm) and 50mm (AUC_{20-50}) as the two metrics. Beside, we adopt AUC on PCK between threshold 0mm and 50mm (AUC_{0-50}) as the third metrics for evaluating 3D hand pose estimation performance. The measuring unit of EPE is millimeter (mm).

4.3. Implementation Details

We implement our TASSN by using PyTorch. In training phase, we set the batch size to 24 and the initial learning rate to 10^{-5} . We train and evaluate our TASSN by using a machine with four GeForce GTX 1080Ti GPUs.

Table 1. 3D hand pose estimation results on the STB and MHP datasets. \uparrow : higher is better; \downarrow : lower is better; The measuring unit of EPE is millimeter (mm).

	$AUC_{0-50} \uparrow$	$AUC_{20-50} \uparrow$	EPE \downarrow
STB Dataset			
TASSN w/o \mathcal{L}_c	0.541	0.735	24.2
TASSN w/o \mathcal{L}_c^m	0.754	0.936	13.6
TASSN	0.773	0.972	11.3
MHP Dataset			
TASSN w/o \mathcal{L}_c	0.492	0.677	28.2
TASSN w/o \mathcal{L}_c^m	0.665	0.870	17.5
TASSN	0.689	0.892	16.2

Since end-to-end training a network from scratch with multiple modules is very difficult, we train our TASSN by using a three-stage procedure. In the first stage, we train the heatmap estimator with the loss \mathcal{L}_h . In the second stage, the GCN hand mesh estimator is initialized by using the pre-trained model provided by [15]. We jointly fine-tune heatmap and hand mesh estimator with the losses \mathcal{L}_h and \mathcal{L}_m on the target dataset without 3D supervision. In the final stage, we conduct an end-to-end training for our TASSN and fine-tune the weights of each sub-module. The model weights of heatmap, GCN hand mesh estimator, and 3D pose estimators are fine-tuned end-to-end. In this stage, we set $\lambda_s = 0.1$, $\lambda_h = 1$, and $\lambda_c^p = \lambda_c^m = 10$.

5. Experimental Results

5.1. Ablation Study of Temporal Consistency Losses

To study the impact of the proposed temporal consistency constraint, we train and evaluate TASSN under the following three settings: 1) TASSN is trained without using temporal consistency loss \mathcal{L}_c , i.e., without any temporal consistency constraint; 2) TASSN is trained without using temporal consistency loss of hand mesh \mathcal{L}_c^m , i.e., with temporal 3D pose constraint but not 3D mesh constraint; 3) TASSN is trained with all the proposed loss functions.

Table 1 shows the evaluation results on two 3D hand pose estimation tasks under the three different settings described above. The PCK curves corresponding to different settings are shown in Figure 6.

We note the following two observations from the ablation study. First, the temporal consistency constraint is critical for 3D pose estimation accuracy. This is clearly illustrated by comparing the results between settings 1 and 3. As shown in Figure 6, TASSN trained with the temporal consistency loss \mathcal{L}_c (red curve, setting 3) outperforms the TASSN trained without using temporal consistency loss (blue curve, setting 1) by a large margin on both the STB and MHP datasets. The quantitative results in Table 1 show that AUC_{0-50} , AUC_{20-50} and EPE, are improved by 0.232,




Figure 6. Performance in PCK on the (a) STB and (b) MHP datasets. TCL and TMCL denote the losses \mathcal{L}_c and \mathcal{L}_c^m , respectively.

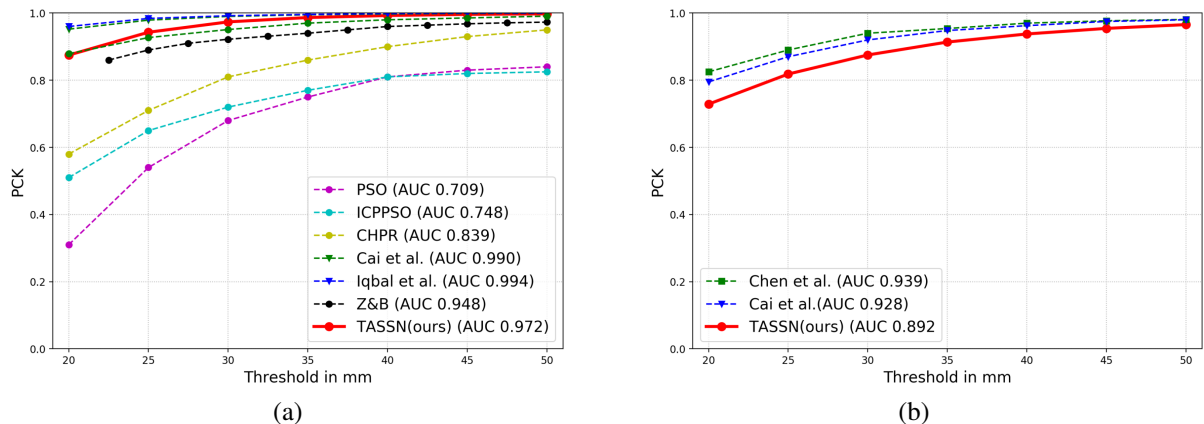


Figure 7. Comparison with the state-of-the-arts. Results in AUC_{20-50} on (a) the STB dataset and (b) the MHP dataset.

0.237, 12.9 on the STB dataset, respectively. A similar trend is also observed on the MHP dataset.

Second, imposing temporal mesh consistency constraints is beneficial for 3D pose estimation. This is illustrated by comparing the results between settings 2 and setting 3. By using the temporal mesh consistency loss \mathcal{L}_c^m , AUC_{0-50} , AUC_{20-50} , EPE improves by 0.024, 0.022, 1.3, respectively, on the STB dataset (Table 1). Results on MHP dataset share a same trend: Test AUCs are boosted by including the temporal mesh consistency loss \mathcal{L}_c^m . It points out that the temporal mesh consistency loss, as an intermediate constraint, facilitates 3D hand pose estimator learning.

In addition to the quantitative analysis, Figure 8 and Figure 9 display some estimated 3D hand poses for visual comparison among these settings on the STB and MHP datasets, respectively. We can see that TASSN, when trained with temporal consistency loss, can produce 3D hand pose estimations highly similar to the ground truth in diverse poses. It is worth noting that our GCN model is initialized with model [15] pretrained on the STB dataset. Our results on STB demonstrate that the temporal consistency is critical to enforce the 3D constraints, without which 3D predic-

tion accuracy drops substantially (Table 1). Moreover, our method generalizes well on other target datasets, e.g., the MHP dataset, where 3D annotations are not used in either model initialization or training. The pose categories and capturing environments are quite different between the two datasets (Figure 5). The effectiveness of our method on the MHP dataset can only be attributed to the temporal consistency constraint (Figure 6).

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

The state-of-the-art methods on both STB and MHP datasets are trained with the 3D annotations, while our method is not. Therefore, we take these methods as the upper bound of our method, and evaluate the performance gaps between these methods and ours.

For the STB dataset, we select six the-state-of-the-art methods for comparison. The selected methods include PSO [3], ICPPSO [8], CHPR [46], the method by Iqbal *et al.* [20], Cai *et al.* [5] and the approach by Zimmermann and Brox [48]. For the MHP dataset, we select two the-state-of-the-art methods for comparison including the approach by Cai *et al.* [5] and the method by Chen *et al.* [6].




Figure 8. Comparison among three different settings on the STB dataset. Columns 1 and 6 are RGB images. Columns 2 and 7 are the result by TASSN trained without temporal consistency loss. Columns 3 and 8 are the result by TASSN trained without temporal mesh consistency loss. Columns 4 and 9 are the result by TASSN. Columns 5 and 10 are the ground truth.




Figure 9. Comparison among three different settings on the MHP dataset. Columns 1 and 6 are RGB images. Columns 2 and 7 are the result by TASSN trained without temporal consistency loss. Columns 3 and 8 are the result by TASSN trained without temporal mesh consistency loss. Columns 4 and 9 are the result by TASSN. Columns 5 and 10 are the ground truth.

Figure 7(a) and Figure 7(b) show the comparison results on STB and MHP datasets, respectively. As expected, TASSN has a performance gap with current state-of-the-art methods on both datasets due to the lack of 3D annotation. However, the performance gaps are relative small. In STB dataset, as shown in Figure 7(a), our methods could even beat some of the methods trained with full 3D annotations.

All together, these results illustrate that 3D pose estimator can be trained without using 3D annotations. Estimating hand pose and mesh from single frames is challenging due to the ambiguities caused by the missing depth information and high flexibility of joints. These challenges can be partly mitigated by utilizing information from video, in which pose and the mesh are highly constrained by the adjacent frames. Temporal information offers an alternative way of enforcing constraints on 3D models for pose and mesh estimation.

6. Conclusions

We propose a video-based hand pose estimation model, temporal-aware self-supervised network (TASSN), to learn

and infer 3D hand pose and mesh from RGB videos. By leveraging temporal consistency between forward and reverse measurements, TASSN can be trained through self-supervised learning without explicit 3D annotations. The experimental results show that TASSN achieves reasonably good results with performance comparable to state-of-the-art models trained with 3D ground truth.

The temporal consistency constraint proposed here offers a convenient and yet effective mechanism for training 3D pose prediction models. Although we illustrate the efficacy of the model without using 3D annotations, it can be used in conjunction with direct supervision with a small number of 3D labeled samples to improve accuracy.

Acknowledgement. This work was supported in part by the Ministry of Science and Technology (MOST) under grants MOST 107-2628-E-009-007-MY3, MOST 109-2634-F-007-013, and MOST 109-2221-E-009-113-MY3, and by Qualcomm through a Taiwan University Research Collaboration Project.

References

- [1] Svitlana Antoshchuk, Mykyta Kovalenko, and Jürgen Sieck. Gesture recognition-based human-computer interaction interface for multimedia applications. In *Digitisation of Culture: Namibian and International Perspectives*. 2018.
- [2] Seungryul Baek, Kwang In Kim, and Tae-Kyun Kim. Pushing the envelope for rgb-based dense 3d hand pose estimation via neural rendering. In *CVPR*, 2019.
- [3] Adnane Boukhayma, Rodrigo de Bem, and Philip HS Torr. 3d hand shape and pose from images in the wild. In *CVPR*, 2019.
- [4] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- [5] Yujun Cai, Lihao Ge, Jianfei Cai, and Junsong Yuan. Weakly-supervised 3d hand pose estimation from monocular rgb images. In *ECCV*, 2018.
- [6] Liangjian Chen, Shih-Yao Lin, Yusheng Xie, Yen-Yu Lin, Wei Fan, and Xiaohui Xie. Drgan: Depth-image guided generative adversarial networks for disentangling rgb and depth images in 3d hand pose estimation. In *WACV*, 2020.
- [7] Liangjian Chen, Shih-Yao Lin, Yusheng Xie, Yen-Yu Lin, and Xiaohui Xie. Mvbm: A large-scale multi-view hand mesh benchmark for accurate 3d hand pose estimation. In *WACV*, 2021.
- [8] Liangjian Chen, Shih-Yao Lin, Yusheng Xie, Hui Tang, Yufan Xue, Yen-Yu Lin, Xiaohui Xie, and Wei Fan. Tagan: Tonality-alignment generative adversarial networks for realistic hand pose synthesis. In *BMVC*, 2019.
- [9] Liangjian Chen, Shih-Yao Lin, Yusheng Xie, Hui Tang, Yufan Xue, Xiaohui Xie, Yen-Yu Lin, and Wei Fan. Generating realistic training images based on tonality-alignment generative adversarial networks for hand pose estimation. *arXiv preprint arXiv:1811.09916*, 2018.
- [10] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *NeurIPS*, 2016.
- [11] Xiaoming Deng, Shuo Yang, Yinda Zhang, Ping Tan, Liang Chang, and Hongan Wang. Hand3d: Hand pose estimation using 3d neural network. *arXiv preprint arXiv:1704.02224*, 2017.
- [12] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *CVPR*, 2015.
- [13] Debidatta Dwivedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. Temporal cycle-consistency learning. 2019.
- [14] Lihao Ge, Yujun Cai, Junwu Weng, and Junsong Yuan. Hand pointnet: 3d hand pose estimation using point sets. In *CVPR*, 2018.
- [15] 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. In *CVPR*, 2019.
- [16] Lihao Ge, Zhou Ren, and Junsong Yuan. Point-to-point regression pointnet for 3d hand pose estimation. In *ECCV*, 2018.
- [17] Francisco Gomez-Donoso, Sergio Orts-Escolano, and Miguel Cazorla. Large-scale multiview 3d hand pose dataset. *arXiv preprint arXiv:1707.03742*, 2017.
- [18] Yi-Ping Hung and Shih-Yao Lin. Re-anchorable virtual panel in three-dimensional space, Dec. 27 2016. US Patent 9,529,446.
- [19] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017.
- [20] Umar Iqbal, Pavlo Molchanov, Thomas Breuel Juergen Gall, and Jan Kautz. Hand pose estimation via latent 2.5 d heatmap regression. In *ECCV*, 2018.
- [21] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *CVPR*, 2018.
- [22] David Joseph Tan, Thomas Cashman, Jonathan Taylor, Andrew Fitzgibbon, Daniel Tarlow, Sameh Khamis, Shahram Izadi, and Jamie Shotton. Fits like a glove: Rapid and reliable hand shape personalization. In *CVPR*, 2016.
- [23] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *CVPR*, 2018.
- [24] Sameh Khamis, Jonathan Taylor, Jamie Shotton, Cem Keskin, Shahram Izadi, and Andrew Fitzgibbon. Learning an efficient model of hand shape variation from depth images. In *CVPR*, 2015.
- [25] Deying Kong, Yifei Chen, Haoyu Ma, Xiangyi Yan, and Xiaohui Xie. Adaptive graphical model network for 2d hand-pose estimation. *arXiv preprint arXiv:1909.08205*, 2019.
- [26] Deying Kong, Haoyu Ma, and Xiaohui Xie. Sia-gcn: A spatial information aware graph neural network with 2d convolutions for hand pose estimation. *arXiv preprint arXiv:2009.12473*, 2020.
- [27] Shile Li and Dongheui Lee. Point-to-pose voting based hand pose estimation using residual permutation equivariant layer. *arXiv preprint arXiv:1812.02050*, 2018.
- [28] Shih-Yao Lin, Yun-Chien Lai, Li-Wei Chan, and Yi-Ping Hung. Real-time 3d model-based gesture tracking for multimedia control. In *ICPR*, 2010.
- [29] Shih-Yao Lin, Chuen-Kai Shie, Shen-Chi Chen, and Yi-Ping Hung. Airtouch panel: a re-anchorable virtual touch panel. In *MM*, 2013.
- [30] Alexandros Makris and A Argyros. Model-based 3d hand tracking with on-line hand shape adaptation. In *BMVC*, 2015.
- [31] 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 *CVPR*, 2018.
- [32] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016.
- [33] Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning features by watching objects move. In *CVPR*, 2017.
- [34] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017.

- [35] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics*, 36(6):245, 2017.
- [36] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. In *ACM Transactions on Graphics*, volume 23, pages 309–314, 2004.
- [37] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *CVPR*, 2017.
- [38] Bugra Tekin, Federica Bogo, and Marc Pollefeys. H+ o: Unified egocentric recognition of 3d hand-object poses and interactions. *CVPR*, 2019.
- [39] Chengde Wan, Thomas Probst, Luc Van Gool, and Angela Yao. Self-supervised 3d hand pose estimation through training by fitting. In *CVPR*, 2019.
- [40] Chengde Wan, Thomas Probst, Luc Van Gool, and Angela Yao. Dense 3d regression for hand pose estimation. In *CVPR*, 2018.
- [41] Xiaokun Wu, Daniel Finnegan, Eamonn O’Neill, and Yong-Liang Yang. Handmap: Robust hand pose estimation via intermediate dense guidance map supervision. In *ECCV*, 2018.
- [42] Zhenyu Wu, Duc Hoang, Shih-Yao Lin, Yusheng Xie, Liangjian Chen, Yen-Yu Lin, Zhangyang Wang, and Wei Fan. Mm-hand: 3d-aware multi-modal guided hand generative network for 3d hand pose synthesis. In *MM*, 2020.
- [43] Linlin Yang and Angela Yao. Disentangling latent hands for image synthesis and pose estimation. *CVPR*, 2019.
- [44] Shanxin Yuan, Guillermo Garcia-Hernando, Björn Stenger, Gyeongsik Moon, Ju Yong Chang, Kyoung Mu Lee, Pavlo Molchanov, Jan Kautz, Sina Honari, Lihao Ge, Junsong Yuan, Xinghao Chen, Guijin Wang, Fan Yang, Kai Akiyama, Yang Wu, Qingfu Wan, Meysam Madadi, Sergio Escalera, Shile Li, Dongheui Lee, Iason Oikonomidis, Antonis Argyros, and Tae-Kyun Kim. Depth-based 3d hand pose estimation: From current achievements to future goals. In *CVPR*, 2018.
- [45] Zahoor Zafrulla, Helene Brashear, Thad Starner, Harley Hamilton, and Peter Presti. American sign language recognition with the kinect. In *ICMI*, 2011.
- [46] 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.
- [47] Xiong Zhang, Qiang Li, Wenbo Zhang, and Wen Zheng. End-to-end hand mesh recovery from a monocular rgb image. *arXiv preprint arXiv:1902.09305*, 2019.
- [48] Christian Zimmermann and Thomas Brox. Learning to estimate 3d hand pose from single rgb images. In *CVPR*, 2017.