

# Hierarchical Temporal Transformer for 3D Hand Pose Estimation and Action Recognition from Egocentric RGB Videos

Yilin Wen<sup>1</sup> Hao Pan<sup>2</sup> Lei Yang<sup>3,1</sup> Jia Pan<sup>1</sup> Taku Komura<sup>1</sup> Wenping Wang<sup>4</sup>

<sup>1</sup>The University of Hong Kong <sup>2</sup>Microsoft Research Asia <sup>3</sup>TransGP <sup>4</sup>Texas A&M University  
 {ylwen, jpan, taku}@cs.hku.hk haopan@microsoft.com l.yang@transgp.hk wenping@tamu.edu

## Abstract

Understanding dynamic hand motions and actions from egocentric RGB videos is a fundamental yet challenging task due to self-occlusion and ambiguity. To address occlusion and ambiguity, we develop a transformer-based framework to exploit temporal information for robust estimation. Noticing the different temporal granularity of and the semantic correlation between hand pose estimation and action recognition, we build a network hierarchy with two cascaded transformer encoders, where the first one exploits the short-term temporal cue for hand pose estimation, and the latter aggregates per-frame pose and object information over a longer time span to recognize the action. Our approach achieves competitive results on two first-person hand action benchmarks, namely FPHA and H2O. Extensive ablation studies verify our design choices.

## 1. Introduction

Perceiving dynamic interacting human hands is fundamental in fields such as human-robot collaboration, imitation learning, and VR/AR applications. Viewing through the egocentric RGB video is especially challenging, as there are frequent self-occlusions between hands and objects, as well as severe ambiguity of action types judged from individual frames (*e.g.* see Fig. 1 where the actions of *pour milk* and *place milk* can only be discerned at complete sequences).

Recent years have witnessed tremendous improvement in 3D hand pose estimation and action recognition. While many works focus on only one of these tasks [10, 12, 13, 18, 23, 39, 51, 56], unified frameworks [27, 47, 52] have also been proposed to address both tasks simultaneously, based on the critical observation that the temporal context of hand poses



Figure 1. Image sequences for *pour milk* and *place milk* under egocentric view from H2O [27], with frequently occluded hand joints and ambiguous action type judged by individual frames. Using temporal information can benefit both tasks of 3D hand pose estimation and action recognition.

helps resolve action ambiguity, implemented via models like LSTM, graph convolutional network or temporal convolutional network. However, we note that temporal information can also benefit hand pose estimation: while interacting hands are usually under partial occlusion and truncation, especially in the egocentric view, they can be inferred more reliably from neighboring frames with different views by temporal motion continuity. Indeed, this idea has not been fully utilized yet among the existing works [27, 47, 52]: *e.g.*, [27, 47] perform hand pose estimation at each frame, leaving the temporal dimension unexplored, and [52] jointly refines action and hand pose through hand-crafted multiple-order motion features and a complex iterative scheme.

We build a simple end-to-end trainable framework to exploit the temporal dimension and achieve effective hand pose estimation and action recognition with a single feed-forward pass. To exploit the relationship among frames, we adopt the transformer architecture [48] which has demonstrated superior performance in sequence modeling. However, action and pose have different temporal granularity: while the action is related to longer time spans lasting for several seconds, the hand pose depicts instantaneous motions. Correspondingly, we use two transformer encoders with different window sizes to respectively leverage the short-term and long-term temporal cues for the per-frame

Work is partially done during the internship of Y. Wen with Microsoft Research Asia. Code and data are available at <https://github.com/fylwen/HTT>.

pose estimation and the action recognition of a whole sequence. Moreover, we notice that the action has a higher semantic level, which is usually defined in the form of “*verb + noun*” [14,27], where *verb* can be derived from the hand motion and the *noun* is the object being manipulated. We thus follow this pattern to build a hierarchy by cascading the pose and action blocks, where the pose block outputs the per-frame hand pose and object label, which are then aggregated by the action block for action recognition.

We evaluate our approach on FPHA [14] and H2O [27], and achieve state-of-the-art performances for 3D hand pose estimation and action recognition from egocentric RGB videos. Our contribution is summarized as follows:

- We propose a simple but efficient end-to-end trainable framework to leverage the temporal information for 3D hand pose estimation and action recognition from egocentric RGB videos.
- We build a hierarchical temporal transformer with two cascaded blocks, to leverage different time spans for pose and action estimation, and model their semantic correlation by deriving the high-level action from the low-level hand motion and manipulated object label.
- We show state-of-the-art performance on two public datasets FPHA [14] and H2O [27].

## 2. Related Work

**3D hand pose estimation from monocular RGB image/video** Massive literature focuses only on the task of 3D pose estimation for hand (or two hands) recorded in the monocular RGB input, to output the 3D skeleton [3, 10, 23, 25, 38–40, 46, 56], or reconstruct also the hand mesh [2, 5, 6, 15–17, 28, 30, 31, 41, 49, 54, 55]. While most of these articles exploit only the spatial dimensions, [3, 10, 17, 40, 49] work on the video-based setting and leverage also temporal consistency for improved robustness, using models such as LSTM in [10] or graph convolutional network in [3].

Same as ours, [27, 47, 52] simultaneously tackle tasks of hand pose estimation and dynamic action recognition from video. However, H+O [47] and H2O [27] perform image-based hand pose estimation by outputting the 3D joint position based on a 3D grid representation, leaving the temporal consistency unexplored. Collaborative [52] initializes the estimation by deriving per-joint 2D heatmaps and depth maps, and iteratively refines the hand pose with temporal action prior. In comparison, we use hierarchical transformers to exploit short-term temporal cues for hand pose estimation and long-term temporal cues for action recognition, with all computations done in a simple and efficient feed-forward pass.

**Action recognition from monocular RGB video** In addition to hand pose, interests grow in recognizing semantic

hand actions where objects are manipulated [14,27]. To recognize the performed hand action (or more general human actions), temporal information is usually necessary. For example, [4, 13, 43] build on a two-stream convolutional network to exploit the spatial-temporal features from the input frames, while [11, 12] leverage on different frame rates for improved flexibility and efficiency. However, they neither stress the pose estimation nor utilize the semantic relationship between pose and action. On the other hand, follow-up research has shown the benefits of explicitly modeling the pose-action relationship and deriving the action type from motion features. For example, [29,37,44] and [7,34] leverage the 2D hand masks or locations to better understand the hand-object interaction, which further facilitates action recognition or anticipation. Another group of works refers to the 3D skeleton features for action recognition, based on temporal models such as temporal convolutional network in [24,36,45,52], LSTM in [32,33,47], and graph convolutional network in [27,42,51].

Among works that simultaneously process both 3D hand pose estimation and action recognition from the given input video, Collaborative [52] aggregates the slow-fast features for multiple orders in the temporal dimension and multi-scale relations on the topology of hand skeleton, and outputs the action category with the temporal convolutional network. H+O [47] first outputs the per-frame pose for the hand and manipulated object (if applicable), as well as the verb category and noun category for action, and integrate the per-frame estimations across the video based on an LSTM framework for improved action recognition. H2O [27] improves H+O [47] with a topology-aware graph convolutional network (TA-GCN) to simultaneously model both the inter-frame and inner-frame relationship for hand object interaction. We follow the semantic relationship to derive action from hand motion and object label, but use a more powerful hierarchical transformer framework to adaptively attend to frames simultaneously, for robust pose and action estimation (see Fig. 7).

**Transformers** Transformers [48] prevail in natural language processing and computer vision. With multi-head self-attention, this framework effectively models relationships among different tokens of an input sequence. For hand pose estimation, the transformer framework has been adopted mainly to capture interactions of image features and key joints in the spatial domain [21,22,28,30,31,38,41]. Meanwhile, [1, 53] adopt clip-level transformers with video-level fusion, to leverage the coherence of clips and the whole sequence for robust action recognition.

We adopt the transformer framework to simultaneously achieve tasks at two levels (pose estimation and action recognition), by respecting their different temporal granularities and modeling their semantic dependency. Moreover, inspired by [35] that builds a hierarchy of shifted windows

to exploit multi-scale features within the image plane, we extend the hierarchical shifting window strategy in the context of multitasking, to learn temporal models of different granularities and model the dependant semantic relationship with efficient computation.

### 3. Methodology

Given a first-person RGB video recording the dynamic hand (or two-hands) performing an action, we estimate the per-frame 3D hand pose and the action category from a given taxonomy [14, 27, 47, 52].

Our network is visualized in Fig. 2. To encode the spatial information of each frame, we first feed each frame to a ResNet-18 [19] and obtain its feature vector from the last layer of the ResNet before softmax. To leverage temporal information, the sequence of frame features is passed to our hierarchical temporal transformer (Sec. 3.1), which is a cascaded framework that uses different time spans for per-frame 3D hand pose estimation (Sec. 3.2) and action recognition (Sec. 3.3). To implement the different time spans efficiently, we adopt the shifting window strategy (see Fig. 3) to split the video into sub-sequences for pose estimation and action recognition.

#### 3.1. Hierarchical Temporal Transformer

Our core design is the hierarchical temporal transformer (HTT, see Fig. 2), which exploits temporal cues from the input video clip  $S = \{I_{S,i} \in \mathbb{R}^{3 \times H \times W} | i = 1, \dots, T\}$  consisting of  $T$  frames. The key ideas are two-fold: On one hand, noticing that the high-level task of action recognition (*e.g.*, pour milk) is defined as a combination of two low-level tasks, namely the hand motion (*e.g.*, the movement of pouring) and the object in manipulation (*e.g.*, the milk bottle) [14, 27], we follow this semantic hierarchy and divide HTT into two cascaded parts, namely the pose block  $\mathbf{P}$  and the action block  $\mathbf{A}$ . The pose block  $\mathbf{P}$  first estimates the per-frame 3D hand pose and the interacting object category, and subsequently the action block  $\mathbf{A}$  aggregates the predicted hand motion and object label over  $S$  for action recognition.

On the other hand, to cope with the different temporal granularities associated with the long-term action and the instantaneous pose, although we adopt transformer architecture for both  $\mathbf{P}$  and  $\mathbf{A}$ , we focus  $\mathbf{P}$  on a narrower temporal receptive field with only  $t$  consecutive frames ( $t < T$ ), while applying  $\mathbf{A}$  over all the  $T$  frames.

#### 3.2. Hand Pose Estimation with Short-Term Temporal Cue

To overcome the frequent occlusion and truncation of hands in interaction in the egocentric view, we leverage temporal cues to improve the robustness of hand pose estimation. Hand poses represent instant motions in time, so

referring to a long time span can overemphasize the temporally distant frames, which could sacrifice the accuracy of local motion (see Sec. 4.4 and the supplementary video for an ablation). Therefore, we localize the time span for pose estimation by dividing the video clip  $S$  into  $m$  consecutive segments  $\text{seg}_t(S) = (\bar{S}_1, \bar{S}_2, \dots, \bar{S}_m)$ , where  $m = \lceil T/t \rceil$ ,  $\bar{S}_i = \{I_{\bar{S}_i,j} = I_{S,k} \in S | k = (i-1)t + j, j = 1, \dots, t\}$  (see Fig. 2); tokens beyond the length  $T$  are padded but masked out from self-attention computation. This scheme can be regarded as a shifting window strategy with window size  $t$ . The module  $\mathbf{P}$  processes each segment  $\bar{S} \in \text{seg}_t(S)$  in parallel to capture the temporal cue for hand pose estimation.

For each local segment  $\bar{S} \in \text{seg}_t(S)$ ,  $\mathbf{P}$  takes the sequence of per-frame ResNet features  $(f(I_{\bar{S},1}), \dots, f(I_{\bar{S},t}))$  as input, and outputs a sequence  $(g_{\bar{S}}(I_{\bar{S},1}), \dots, g_{\bar{S}}(I_{\bar{S},t}))$ . The  $j$ -th token  $g_{\bar{S}}(I_{\bar{S},j}) \in \mathbb{R}^d$  ( $j = 1, \dots, t$ ) corresponds to the frame  $I_{\bar{S},j}$  and encodes also the temporal cue from  $\bar{S}$ . We then decode the hand pose for  $I \in \bar{S}$  from its temporal-dependent feature  $g_{\bar{S}}(I)$ :

$$P_I = (P_I^{2D}, P_I^{dep}) = \text{MLP}_1(g_{\bar{S}}(I)) \quad (1)$$

where for hand (or two-hands) with  $J$  hand joints,  $P$  is the concatenation of the joint coordinates in the 2D image plane  $P_I^{2D} \in \mathbb{R}^{J \times 2}$  and the joint depth to the camera  $P_I^{dep} \in \mathbb{R}^{J \times 1}$ , while  $\text{MLP}_1$  has three layers of width  $[d, d, 3J]$  with LeakyReLU as the activation functions for the hidden layers. For supervision we compare the prediction with the groundtruth hand pose  $(P_{I,gt}^{2D}, P_{I,gt}^{dep})$  and minimize the  $L1$ -loss:

$$L_H(I) = \frac{1}{J} (\|P_I^{2D} - P_{I,gt}^{2D}\|_1 + \lambda_1 \|P_I^{dep} - P_{I,gt}^{dep}\|_1) \quad (2)$$

where  $\lambda_1$  is a hyper-parameter to balance the different magnitudes of the 2D loss and the depth loss. The 3D positions of the hand joints in the camera space  $P_I^{3D} \in \mathbb{R}^{J \times 3}$  for  $I$  can then be recovered given the camera intrinsics.

As the category of the manipulated object supplies the noun of an action [47], we also predict from  $g_{\bar{S}}(I)$  the object category. Denoting the number of object classes as  $n_o$ , we obtain an  $n_o$ -dim classification probability vector from  $g_{\bar{S}}(I)$  with another  $\text{MLP}_2$  with two layers of width  $[d, n_o]$ :

$$O_I = [p(o_1|I), \dots, p(o_{n_o}|I)] = \text{softmax}(\text{MLP}_2(g_{\bar{S}}(I))). \quad (3)$$

Given the groundtruth object label  $o_{gt}$ , the target probability is defined as a one-hot vector  $w^o(I) = (w_{I,1}^o, \dots, w_{I,n_o}^o)$  with only  $w_{I,o_{gt}}^o = 1$  for  $o_{gt}$ . The object classification task is supervised to minimize the cross entropy between the predicted  $O_I$  and the groundtruth  $w^o(I)$ :

$$L_O(I) = - \sum_{i=1}^{n_o} w_{I,i}^o \log p(o_i|I). \quad (4)$$

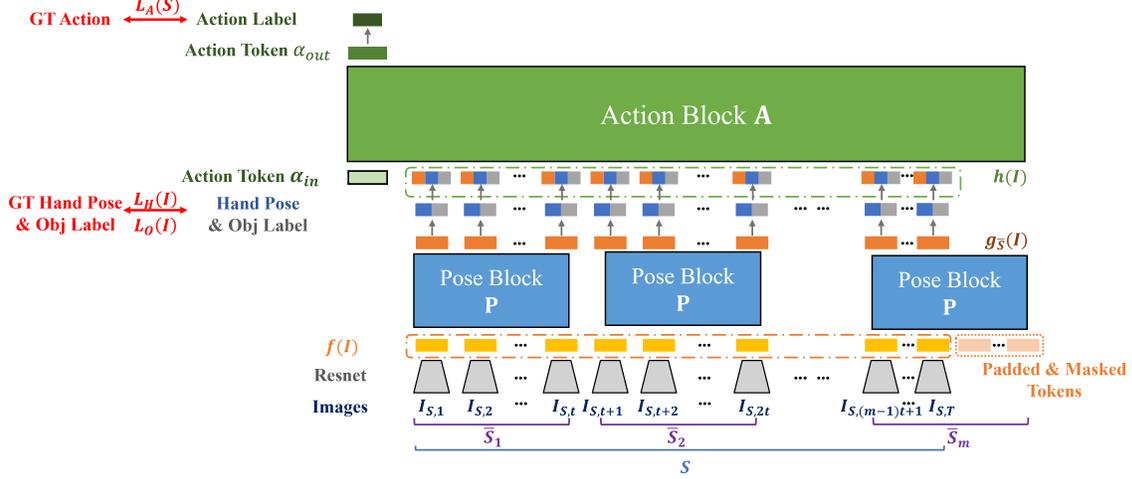


Figure 2. Overview of our framework. Given input video  $S$ , we first feed each image to a ResNet feature extractor, and then leverage short-term temporal cue via  $\mathbf{P}$  applied to shifted windowed frames, to estimate per-frame 3D hand pose and object label. We finally aggregate the long-term temporal cue with  $\mathbf{A}$ , to predict the performed action label for  $S$  from the hand motion and manipulated object label. We supervise the learning with GT labels.

### 3.3. Action Recognition with Long-Term Temporal Cue

Perceiving a long time span clarifies the action performed (Fig. 1). Therefore we have the action module  $\mathbf{A}$  of **HTT** that leverages the full input sequence  $S$  to predict the action label.

Specifically, the input of  $\mathbf{A}$  is a sequence of  $T + 1$  tokens  $(\alpha_{in}, h(I_{S,1}), \dots, h(I_{S,T}))$ . We follow previous works [8, 9] to introduce a trainable token  $\alpha_{in} \in \mathbb{R}^d$  that aggregates the global information across  $S$  for action classification. The remaining tokens encode the per-frame information of the hand pose and object label. For a frame  $I \in \bar{S}$  with  $\bar{S} \in \text{seg}_t(S)$ ,  $h(I) \in \mathbb{R}^d$  mixes its 2D hand pose  $P_I^{2D}$ , the probability distribution of object classification  $O_I$  and the image feature  $g_{\bar{S}}(I)$  computed by  $\mathbf{P}$ , i.e.:

$$h(I) = \text{FC}_1[\text{FC}_2(P_I^{2D}), \text{FC}_3(O_I), g_{\bar{S}}(I)] \quad (5)$$

where  $\text{FC}_2(\cdot)$  and  $\text{FC}_3(\cdot)$  respectively output  $d$ -dim features, and  $\text{FC}_1[\cdot, \cdot, \cdot]$  reduces the concatenation of three input features into  $d$ -dim to fit in the token dimension of  $\mathbf{A}$ . Hence,  $\mathbf{A}$  is expected to recognize the performed action based on the hand motion, the label of the manipulated object, and also the per-frame image features which may encode other useful cues like object appearance and hand-object contacts. We provide an ablation for the input feature leveraged by  $\mathbf{A}$  in Sec. 4.4 and Tab. 5.

To classify the action for  $S$ , we make use of the first token  $\alpha_{out} \in \mathbb{R}^d$  of the output sequence by  $\mathbf{A}$ , where with an  $\text{FC}_4$  we regress the probability distribution over the given action taxonomy with  $n_a$  pre-defined actions as:

$$\mathbf{A}(S) = [p(a_1|S), \dots, p(a_{n_a}|S)] = \text{softmax}(\text{FC}_4(\alpha_{out})). \quad (6)$$

Given the groundtruth action  $a_{gt}$  and the target probability as a one-hot vector  $w(S) = (w_{S,1}, \dots, w_{S,n_a})$  with only  $w_{S,a_{gt}} = 1$  for  $a_{gt}$ , we supervise by minimizing the cross entropy for classifying the action category as:

$$L_A(S) = - \sum_{i=1}^{n_a} w_{S,i} \log p(a_i|S). \quad (7)$$

To summarize, given an input video clip  $S$  to **HTT**, the total training loss is:

$$L = L_A(S) + \frac{1}{T} \sum_{\bar{S} \in \text{seg}_t(S)} \sum_{I \in \bar{S}} (\lambda_2 L_H(I) + \lambda_3 L_O(I)) \quad (8)$$

where  $\lambda_2, \lambda_3$  are hyperparameters to balance different loss terms.

### 3.4. Implementation Details

**Network parameters** We set  $T = 128$  and  $t = 16$  as the maximum input sequence length for  $\mathbf{A}$  and  $\mathbf{P}$ , respectively. All input images are resized to  $H = 270, W = 480$ . The ResNet feature, as well as the tokens of  $\mathbf{P}$  and  $\mathbf{A}$  are all of the dimension  $d = 512$ . Both  $\mathbf{P}$  and  $\mathbf{A}$  have two transformer encoder layers, use the fixed sine/cosine position encoding [48] and put layer normalization before the attention and feed-forward operations; each layer has 8 attention heads, and the feed-forward block has a dimension of 2048.

The value of  $T$  derives from the limitation of available computational resources, as simultaneously processing all frames for an input video  $V$  with more than  $T$  frames causes an out-of-memory error. To process a longer video, we split  $V$  into a clip set  $\text{seg}_T(V)$  where each clip  $S \in \text{seg}_T(V)$  can be processed by **HTT**. The construction of  $\text{seg}_T(V)$  has two

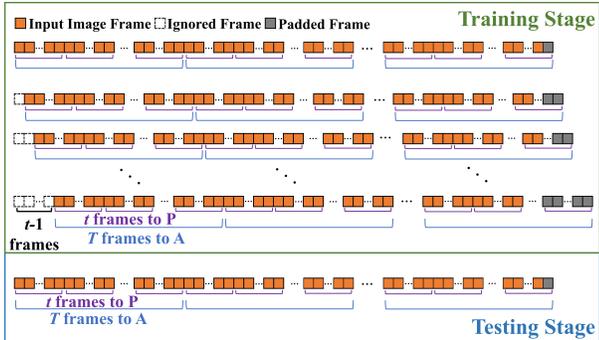


Figure 3. Segmentation strategy for dividing a long video into inputs of **HTT**. In the testing stage, we start from the first frame, while in the training stage, we offset the starting frame within  $t$  frames to augment the training data diversity.

steps: we first downsample  $V$  into two sub-sequences with a sampling ratio of 2, and then divide both of the two sub-sequences into consecutive clips by adopting the shifting window strategy with a window size  $T$ .

**Training stage details** We set  $\lambda_1 = 200$ ,  $\lambda_2 = 0.5$ ,  $\lambda_3 = 1$  to balance different loss terms, and train with the Adam optimizer [26] with an initial learning rate of  $3 \times 10^{-5}$ , where we halve the learning rate every 15 epochs. Same as [18], the ResNet-18 is initialized with weights pretrained on ImageNet, which are updated during backpropagation, except that the batch normalization layers are frozen. The whole network was trained on 2 GPUs with a total batch size of 2, and we observe convergence after 45 epochs. We also follow [18] to add online augmentation to input image sequences, by randomly adjusting the hue, saturation, contrast, and brightness, as well as adding random Gaussian blur and 2D translation.

To augment sampling variations of training data, we offset the starting frame to each of the first  $t$  frames, as illustrated in Fig. 3. Offsets within the  $t$  segment ensure that both **P** and **A** consume different augmented data generated from the same sequence; in contrast, an offset larger than  $t$  will cause duplicated segments fed to **P**.

**Testing stage computation** For a video with more than  $T$  frames, we first obtain the set  $\text{seg}_T(V)$  for the original test video  $V$ , then feed each  $S \in \text{seg}_T(V)$  to **HTT**, and obtain the per-image 3D hand pose from the output of **P** and the action category for  $V$  by voting from the output category among  $S \in \text{seg}_T(V)$  (Fig. 3). Note that in this way we achieve efficient computation as each image is processed only once by both **P** and **A**.

## 4. Experiments

### 4.1. Dataset

We train and test the proposed method on two public datasets FPFA [14] and H2O [27] for 3D hand pose

estimation and action recognition from first-person views. Both datasets are recorded in multiple indoor scenarios and have a frame rate of 30 fps. We use the groundtruth labels for hand pose, action, and object category provided by the datasets for supervision and evaluation.

**FPFA [14]** This dataset records 6 subjects performing 45 actions, interacting with 26 rigid or non-rigid objects. It annotates  $J = 21$  joints for only the subject’s right hand, with the hand pose data collected from the wearable magnetic sensors. Following the *action split* [14, 47, 52] we train our method on 600 videos and test it on 575 videos. All subjects and actions are seen at both the training and testing stages.

**H2O [27]** This dataset has 4 subjects performing 36 actions related to manipulating 8 objects. Markerless 3D annotations for both hands with a total number of  $J = 21 \times 2$  joints are provided. We follow [27] to use egocentric view sequences with annotated actions for training and testing, where the training split has 569 videos including all actions for the first 3 subjects, and the testing split has 242 videos of the remaining subject unseen in training.

### 4.2. Metrics

**Action recognition** We report the *Classification Accuracy* over the test split, by comparing the predicted and groundtruth action categories for each test video.

**3D hand pose estimation** We evaluate the proposed method by comparing its estimated 3D joint positions with the groundtruth. On FPFA [14], we follow the baseline methods [10, 47, 52] to report the *Percentage of Correct Keypoints (PCK)* under different error thresholds for joints [56] and the corresponding *Area Under the Curve (AUC)*. We respectively report *3D PCK* and *3D PCK-RA* for evaluation in the camera space and the root-aligned (RA) space where for each frame the estimated wrist is aligned with its groundtruth position.

On H2O [27], we follow its benchmark [27] to evaluate in the camera space, by reporting the *3D PCK* metric and the *Mean End-Point Error (MEPE)* for hands [56]. We also establish a baseline for the root-aligned space, where for each frame we align the estimated wrist with its groundtruth position and report the corresponding *3D PCK-RA* and *MEPE-RA* for the non-wrist joints.

### 4.3. Comparison with Related Works

**Comparison on FPFA [14]** Our closest state-of-the-art baseline method is Collaborative [52] which leverages the temporal cue for both 3D hand pose estimation and action recognition. Another baseline method is H+O [47], which first conducts single image-based pose estimation and then exploits the temporal domain for action recognition. Moreover, for hand pose estimation, we also compare our method with ACE-Net [10] which focuses only on video-based 3D hand pose estimation and reports on FPFA [14]. Among

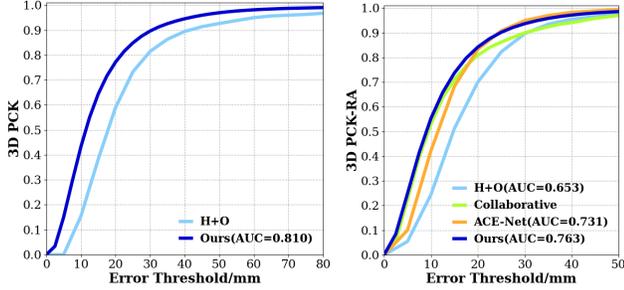


Figure 4. 3D PCK(-RA) of hand pose estimation on FPFA [14]. We report the 3D PCK(-RA) versus different error thresholds by respectively evaluating in the camera space (Left figure) and the root-aligned space (Right figure).

	Joule-color [20]	Two-stream [13]	H+O [47]	Collaborative [52]	Ours
Accuracy	66.78	75.30	82.43	85.22	<b>94.09</b>

Table 1. Classification accuracy of action recognition for RGB-based methods on FPFA [14].

RGB-based methods that tackle only the action recognition, we add baseline methods [13, 20] from FPFA [14] to our discussion.

We demonstrate the effectiveness of our method with competitive results for both action recognition and hand pose estimation, as shown in Tab. 1 and Fig. 4. For accuracy of action recognition, we outperform baseline methods by respectively improving for more than 8% over Collaborative [52] and 11% over H+O [47]. For hand pose estimation, Fig. 4 shows the 3D PCK and the 3D PCK-RA, where we note that only H+O [47] reports results in the camera space. Compared with our closest related work Collaborative [52], we improve the performance with a notable margin for error thresholds > 15mm. We also show better accuracy than H+O [47] at all reported error thresholds, and outperform ACE-Net [10] for error thresholds less than 20mm and the AUC metrics. Unfortunately, the comparing methods have not yet released their code or data for inspecting the specific cases, but from our qualitative cases for 3D pose estimation in Fig. 6 and the supplementary video we can see that our results are robust to self-occlusions and truncations that are common in egocentric view (more discussions in Sec. 4.4), which can be the primary reason for the differences.

**Comparison on H2O [27]** We focus on comparison with H+O [47] and H2O [27], as these two methods output both the per-frame 3D hand pose and the action category for the input video. For each task, we additionally compare our method with other RGB-based methods that solve only the given task: for action recognition, C2D [50], I3D [4] and SlowFast [12] are compared; for hand pose estimation, we add LPC [18] into the comparison. Results of these baseline methods are copied from the benchmark of H2O [27]. We note that LPC [18], H+O [47] and H2O [27] con-

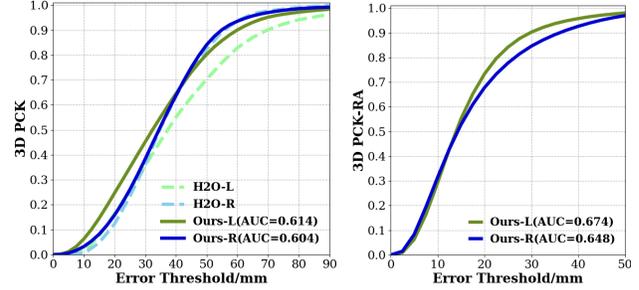


Figure 5. 3D PCK(-RA) of hand pose estimation on the test split of H2O [27]. We report the 3D PCK(-RA) versus different error thresholds by respectively evaluating in the camera space (Left figure) and the root-aligned space (Right figure).

	MEPE in Camera Space			MEPE-RA	
	H+O [47]	LPC [18]	H2O [27]	Ours	Ours
Left	41.42	39.56	41.45	<b>35.02</b>	16.59
Right	38.86	41.87	37.21	<b>35.63</b>	17.91

Table 2. MEPE and MEPE-RA of hand pose estimation on the test split of H2O [27], the unit is mm.

duct image-based pose estimation. Furthermore, H+O [47] and LPC [18] train separate networks for each hand, while H2O [27] and ours estimate for both hands within a unified network.

We first report the accuracy of action recognition in Tab. 3, where we outperform baseline methods and improve the state-of-the-art setup proposed in H2O [27] with a margin over 7%. For 3D hand pose estimation, we first follow the H2O benchmark [27] to evaluate the results in the camera space, reporting the 3D PCK and MEPE in Fig. 5 and Tab. 2, respectively. Our better performance demonstrates the importance and benefits of leveraging the temporal coherence for robust hand pose estimation under frequent occlusions in the egocentric view. We further establish a baseline by evaluating the results in the root-aligned space, and report our MEPE-RA and 3D PCK-RA in Tab. 2 and Fig. 5, respectively. We show our qualitative cases in Fig. 6 and the supplementary video.

#### 4.4. Ablation Study

We verify our key design choices with ablation study on both the FPFA [14] and H2O [27] datasets. For pose estimation, we report MEPE and AUC for 3D PCK versus different error thresholds, with regard to both the camera space and per-frame root-aligned space. For action, we report the classification accuracy.

**Short-term temporal cue for pose** We first demonstrate the benefits brought by leveraging temporal cue for pose estimation, where we remove the pose block **P** and conduct image-based pose estimation by regressing the hand pose

	C2D [50]	I3D [4]	SlowFast [12]	H+O [47]	H2O w/ ST-GCN [51]	H2O w/ TA-GCN [27]	Ours
Accuracy	70.66	75.21	77.69	68.88	73.86	79.25	<b>86.36</b>

Table 3. Classification accuracy of action recognition for RGB-based methods on H2O [27].

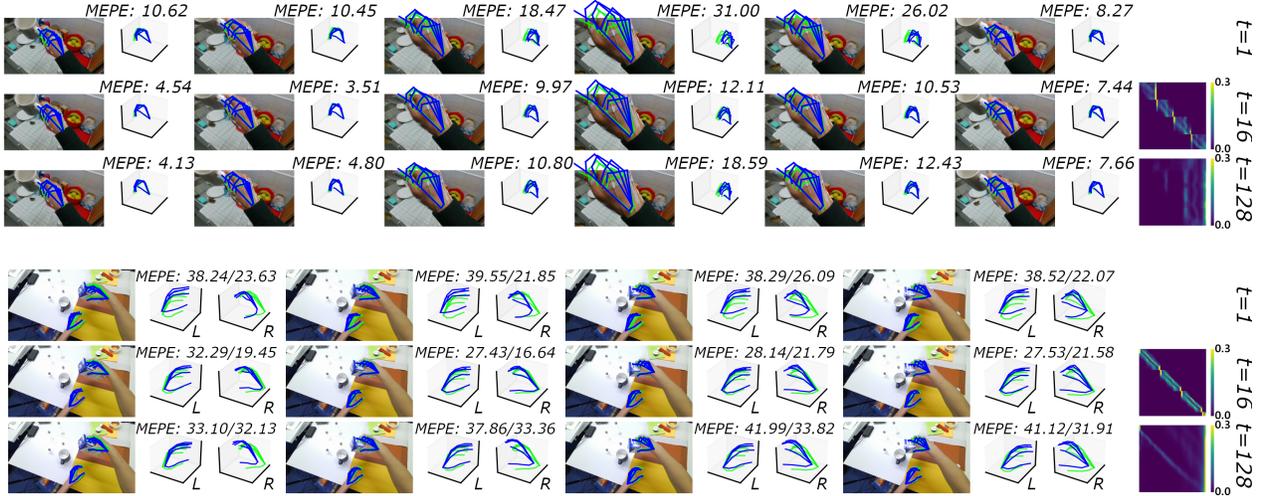


Figure 6. Qualitative cases of 3D hand pose estimation in the camera space and its 2D projection for sequential images from FPHA [14](Upper) and H2O [27](Lower). We show our choice of leveraging the short-term temporal cue with  $t = 16$ , and compare it with the image-based baseline of  $t = 1$ , and the setting using a long-term temporal cue with  $t = 128$ . The estimation and GT are respectively denoted in blue and green, where the MEPE with unit *mm* is also attached. For  $t = 16, 128$ , the attention weights in the final layer of  $\mathbf{P}$  are visualized. Our  $t = 16$  shows enhanced robustness under invisible joints compared with  $t = 1$ , while avoiding over-attending to distant frames and ensuring sharp local motion compared with a long-term  $t = 128$ .

and object label from the ResNet feature. Thus, we have the temporal window size  $t = 1$  for pose and keep  $T = 128$  for action. We report the evaluation of the 3D hand pose for this setup in Tab. 4, whose degraded performance in both camera space and root-aligned space for both datasets validates the benefits brought by exploiting short-term temporal cues with  $t > 1$ . Moreover, we qualitatively compare the two setups with  $t = 1$  and our  $t = 16$  in Fig. 6 and the supplementary video, where with temporal cue we see improved robustness under occlusion and truncation.

We then move on to discuss the impact of various time spans  $t$ , and report in Tab. 4 for hand pose estimation. In the camera space,  $t = 16$  shows the best performance while longer time spans degrade accuracy, which reveals that attending to temporally distant frames in a long time span can override the local sharpness for motion. For further validation, we present qualitative cases in Fig. 6 and the supplementary video, where we compare the two setups of  $t = T = 128$  and our  $t = 16$  by visualizing the hand pose estimation and the attention map for the final layer of  $\mathbf{P}$ , which confirms that the long time-spans cause diffused attention maps that cannot capture localized motions. Moreover, in the per-frame root-aligned space, leveraging a local time span with  $t$  as 32 or 64 report the best results, while our choice of using a smaller  $t = 16$  shows comparable per-

formance. These results validate our choice of  $t = 16$  for robust and accurate pose estimation.

**Long-term temporal cue for action** Given the necessity of temporal information for resolving action ambiguity (see Fig. 1), we first verify the benefits of using a long time span for action, by varying  $T$  from  $T = t = 16$  to  $T = 128$ . We report the accuracy for both datasets in the upper part of Tab. 5. Unlike per-frame hand pose estimation which benefits from a short-term temporal cue, here longer time spans improve the accuracy of action recognition, which supports our design of leveraging different time spans for tasks with different temporal granularity.

For our  $T = 128$ , we also visualize in Fig. 7 the attention weights from the action token to the per-frame tokens in the final layer of  $\mathbf{A}$ , obtained on a video for *take out espresso* whose action can only be judged by the last few frames depicting the process of taking the capsule out of the box. We observe that these keyframes for action recognition are given the most attention. More cases for various actions are provided in the supplementary video, where the distribution of attention weights shows clear correspondence patterns with respect to different actions.

**Cascaded hierarchy for action recognition** To echo the “*verb+noun*” pattern described by action labels, we cascade the two blocks  $\mathbf{P}$  and  $\mathbf{A}$  to classify the action category, and

$t$	FPHA [14]				H2O [27]							
	In Camera Space		In Root-Aligned Space		In Camera Space				In Root-Aligned Space			
	AUC(0-80)	MEPE	AUC-RA(0-50)	MEPE-RA	AUC(0-90)		MEPE		AUC-RA(0-50)		MEPE-RA	
					Left	Right	Left	Right	Left	Right	Left	Right
1	0.776	18.78	0.707	15.01	0.563	0.551	40.12	40.62	0.648	0.590	17.97	21.22
8	0.802	16.54	0.750	12.77	0.607	0.600	35.91	36.00	0.674	0.629	16.58	18.89
16	<b>0.810</b>	<b>15.81</b>	0.763	12.13	<b>0.614</b>	<b>0.604</b>	<b>35.02</b>	<b>35.63</b>	0.674	0.648	16.59	17.91
32	0.805	16.14	0.766	11.96	0.594	0.573	36.78	38.48	0.676	<b>0.650</b>	16.56	<b>17.85</b>
64	0.805	16.19	<b>0.767</b>	<b>11.92</b>	0.599	0.567	36.47	39.02	<b>0.684</b>	0.643	<b>16.23</b>	18.10
128	0.800	16.73	0.761	12.20	0.599	0.564	36.41	39.36	0.673	0.636	16.71	18.66

Table 4. Ablation on time span  $t$  for pose block  $\mathbf{P}$  on the FPHA [14] and H2O [27] dataset. Reported are the AUC(0- $E$ ) with  $E \in \{50, 80, 90\}$  for 3D PCK(-RA) at error thresholds ranging from 0 to  $E$  mm and the MEPE(-RA) in the unit of mm. All setups have  $T = 128$  for the action block  $\mathbf{A}$ .



Figure 7. Visualization for weights of attention in the final layer of  $\mathbf{A}$ , from the action token to the frames. Presented is a video of *take out espresso*, whose down-sampled image sequence is shown in the top row. The last few frames are the key to recognizing the action; in response, our network pays the most attention to these frames. The full image sequence of this demo can be found in the supplementary video.

$T$	Cascaded $\mathbf{P}, \mathbf{A}$	Input Feature for $\mathbf{A}$			Classification Accuracy	
		Image Feature	Hand Pose	Object Label	FPHA [14]	H2O [27]
16	✓	✓	✓	✓	90.96	74.38
32	✓	✓	✓	✓	91.65	79.34
64	✓	✓	✓	✓	92.35	78.51
128	✓	✓	✓	✓	<b>94.09</b>	<b>86.36</b>
128		✓			93.22	80.17
128	✓	✓	✓	✓	<b>94.09</b>	<b>86.36</b>
128	✓	✓	✓		93.57	82.65
128	✓	✓		✓	91.65	85.12
128	✓		✓	✓	90.26	75.21
128	✓	✓	✓	✓	<b>94.09</b>	<b>86.36</b>

Table 5. Ablation on key designs for action recognition. Reported is the accuracy on the FPHA [14] and H2O [27] dataset. All setups have  $t = 16$  for the pose block  $\mathbf{P}$ .

feed the per-frame hand pose, object label, and image feature obtained by  $\mathbf{P}$  as input for  $\mathbf{A}$  (see Eq. 5). We first verify the necessity of the cascaded design. For comparison, we set  $\mathbf{P}$  and  $\mathbf{A}$  in a parallel structure, where the ResNet image feature is the per-frame input token for both  $\mathbf{P}$  and  $\mathbf{A}$ . We observe degraded performance with this parallel design, as reported in the middle part of Tab. 5.

Then based on our cascaded design, we ablate the components of the per-frame input to  $\mathbf{A}$ . As reported in the lower part in Tab. 5, our full version shows the best performance, which demonstrates the benefits of modeling the semantic correlation among hand pose, object label, and action, as well as the gains of letting  $\mathbf{A}$  exploit also the other miscellaneous cues encoded in the image feature.

## 5. Conclusion

In this paper, we have proposed a unified framework to simultaneously handle the tasks of 3D hand pose estimation

and action recognition for an egocentric RGB video. Our core framework is a hierarchical temporal transformer that has two cascaded parts, where the first one refers to a relatively short time span to output per-frame 3D hand pose and object label with enhanced robustness under occlusion and truncation, and the second one exploits the long-term temporal cue by aggregating the per-frame information for action recognition. In this way, we model the correlation between the two tasks with different semantic levels, and leverage different time spans according to their temporal granularity, both contributing to improved performances as verified by ablation studies. Extensive evaluations on two first-person hand action benchmarks demonstrate the effectiveness of our method.

**Limitation and future work** We mainly focus on learning to exploit the temporal dimension for pose estimation and action recognition with a transformer-based framework, but adaptively modeling the spatial interactions of hand joints and objects (*e.g.*, via a transformer module rather than a ResNet feature extractor) may further improve the performance, which we leave as future work. Another interesting direction is that our hierarchical sequential framework can potentially be extended to model motion prediction and generation, which are fundamental to tasks such as early action detection and forecasting for human-robot collaboration.

**Acknowledgments** This work is partially funded by the Research Grant Council of Hong Kong (GRF 17210222). This work is also partially supported by the Innovation and Technology Commission of the HKSAR Government under the InnoHK initiative.

## References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6836–6846, 2021. 2
- [2] Adnane Boukhayma, Rodrigo de Bem, and Philip HS Torr. 3d hand shape and pose from images in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10843–10852, 2019. 2
- [3] Yujun Cai, Lihao Ge, Jun Liu, Jianfei Cai, Tat-Jen Cham, Junsong Yuan, and Nadia Magnenat Thalmann. Exploiting spatial-temporal relationships for 3d pose estimation via graph convolutional networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2272–2281, 2019. 2
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 2, 6, 7
- [5] Xingyu Chen, Yufeng Liu, Chongyang Ma, Jianlong Chang, Huayan Wang, Tian Chen, Xiaoyan Guo, Pengfei Wan, and Wen Zheng. Camera-space hand mesh recovery via semantic aggregation and adaptive 2d-1d registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13274–13283, 2021. 2
- [6] Hongsuk Choi, Gyeongsik Moon, and Kyoung Mu Lee. Pose2mesh: Graph convolutional network for 3d human pose and mesh recovery from a 2d human pose. In *European Conference on Computer Vision*, pages 769–787. Springer, 2020. 2
- [7] Eadom Dessalene, Chinmaya Devaraj, Michael Maynard, Cornelia Fermuller, and Yiannis Aloimonos. Forecasting action through contact representations from first person video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *Association for Computational Linguistics*, 2019. 4
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations*, 2021. 4
- [10] Zhipeng Fan, Jun Liu, and Yao Wang. Adaptive computationally efficient network for monocular 3d hand pose estimation. In *European Conference on Computer Vision*, pages 127–144. Springer, 2020. 1, 2, 5, 6
- [11] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 203–213, 2020. 2
- [12] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019. 1, 2, 6, 7
- [13] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1933–1941, 2016. 1, 2, 6
- [14] Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, and Tae-Kyun Kim. First-person hand action benchmark with rgb-d videos and 3d hand pose annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 409–419, 2018. 2, 3, 5, 6, 7, 8
- [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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10833–10842, 2019. 2
- [16] Shreyas Hampali, Sayan Deb Sarkar, Mahdi Rad, and Vincent Lepetit. Keypoint transformer: Solving joint identification in challenging hands and object interactions for accurate 3d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11090–11100, 2022. 2
- [17] Shangchen Han, Beibei Liu, Randi Cabezas, Christopher D Twigg, Peizhao Zhang, Jeff Petkau, Tsz-Ho Yu, Chun-Jung Tai, Muzaffer Akbay, Zheng Wang, et al. Megatrack: monochrome egocentric articulated hand-tracking for virtual reality. *ACM Transactions on Graphics (ToG)*, 39(4):87–1, 2020. 2
- [18] Yana Hasson, Bugra Tekin, Federica Bogo, Ivan Laptev, Marc Pollefeys, and Cordelia Schmid. Leveraging photometric consistency over time for sparsely supervised hand-object reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 571–580, 2020. 1, 5, 6
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition.

- In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3
- [20] Jian-Fang Hu, Wei-Shi Zheng, Jianhuang Lai, and Jianguo Zhang. Jointly learning heterogeneous features for rgb-d activity recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5344–5352, 2015. 6
- [21] Lin Huang, Jianchao Tan, Ji Liu, and Junsong Yuan. Hand-transformer: Non-autoregressive structured modeling for 3d hand pose estimation. In *European Conference on Computer Vision*, pages 17–33. Springer, 2020. 2
- [22] Lin Huang, Jianchao Tan, Jingjing Meng, Ji Liu, and Junsong Yuan. Hot-net: Non-autoregressive transformer for 3d hand-object pose estimation. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 3136–3145, 2020. 2
- [23] Umar Iqbal, Pavlo Molchanov, Thomas Breuel, Juer-gen Gall, and Jan Kautz. Hand pose estimation via latent 2.5 d heatmap regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 118–134, 2018. 1, 2
- [24] Qihong Ke, Mohammed Bennamoun, Senjian An, Ferdous Sohel, and Farid Boussaid. A new representation of skeleton sequences for 3d action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3288–3297, 2017. 2
- [25] Dong Uk Kim, Kwang In Kim, and Seungryul Baek. End-to-end detection and pose estimation of two interacting hands. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11189–11198, 2021. 2
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015. 5
- [27] Taein Kwon, Bugra Tekin, Jan Stühmer, Federica Bogo, and Marc Pollefeys. H2o: Two hands manipulating objects for first person interaction recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10138–10148, 2021. 1, 2, 3, 5, 6, 7, 8
- [28] Mengcheng Li, Liang An, Hongwen Zhang, Lianpeng Wu, Feng Chen, Tao Yu, and Yebin Liu. Interacting attention graph for single image two-hand reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2761–2770, 2022. 2
- [29] Yin Li, Zhefan Ye, and James M Rehg. Delving into egocentric actions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 287–295, 2015. 2
- [30] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1954–1963, 2021. 2
- [31] Kevin Lin, Lijuan Wang, and Zicheng Liu. Mesh graphormer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12939–12948, 2021. 2
- [32] Jun Liu, Amir Shahroudy, Dong Xu, and Gang Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In *European conference on computer vision*, pages 816–833. Springer, 2016. 2
- [33] Jun Liu, Gang Wang, Ping Hu, Ling-Yu Duan, and Alex C Kot. Global context-aware attention lstm networks for 3d action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1647–1656, 2017. 2
- [34] Miao Liu, Siyu Tang, Yin Li, and James M Rehg. Forecasting human-object interaction: joint prediction of motor attention and actions in first person video. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 704–721. Springer, 2020. 2
- [35] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 2
- [36] Diogo C Luvizon, David Picard, and Hedi Tabia. 2d/3d pose estimation and action recognition using multitask deep learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5137–5146, 2018. 2
- [37] Minghuang Ma, Haoqi Fan, and Kris M Kitani. Going deeper into first-person activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1894–1903, 2016. 2
- [38] Hao Meng, Sheng Jin, Wentao Liu, Chen Qian, Mengxiang Lin, Wanli Ouyang, and Ping Luo. 3d interacting hand pose estimation by hand de-occlusion and removal. *European Conference on Computer Vision*, 2022. 2
- [39] Gyeongsik Moon, Shou-I Yu, He Wen, Takaaki Shiratori, and Kyoung Mu Lee. Interhand2.6m: A dataset and baseline for 3d interacting hand pose estimation from a single rgb image. In *European Conference on*

- Computer Vision*, pages 548–564. Springer, 2020. 1, 2
- [40] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 49–59, 2018. 2
- [41] JoonKyu Park, Yeonguk Oh, Gyeongsik Moon, Hong-suk Choi, and Kyoung Mu Lee. Handocnet: Occlusion-robust 3d hand mesh estimation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1496–1505, 2022. 2
- [42] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12026–12035, 2019. 2
- [43] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *Advances in neural information processing systems*, 27, 2014. 2
- [44] Suriya Singh, Chetan Arora, and CV Jawahar. First person action recognition using deep learned descriptors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2620–2628, 2016. 2
- [45] Tae Soo Kim and Austin Reiter. Interpretable 3d human action analysis with temporal convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 20–28, 2017. 2
- [46] Adrian Spurr, Umar Iqbal, Pavlo Molchanov, Otmar Hilliges, and Jan Kautz. Weakly supervised 3d hand pose estimation via biomechanical constraints. In *European Conference on Computer Vision*, pages 211–228. Springer, 2020. 2
- [47] Bugra Tekin, Federica Bogo, and Marc Pollefeys. H+o: Unified egocentric recognition of 3d hand-object poses and interactions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4511–4520, 2019. 1, 2, 3, 5, 6, 7
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 1, 2, 4
- [49] Jiayi Wang, Franziska Mueller, Florian Bernard, Suzanne Sorli, Oleksandr Sotnychenko, Neng Qian, Miguel A Otaduy, Dan Casas, and Christian Theobalt. Rgb2hands: real-time tracking of 3d hand interactions from monocular rgb video. *ACM Transactions on Graphics (ToG)*, 39(6):1–16, 2020. 2
- [50] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. 6, 7
- [51] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Thirty-second AAAI conference on artificial intelligence*, 2018. 1, 2, 7
- [52] Siyuan Yang, Jun Liu, Shijian Lu, Meng Hwa Er, and Alex C Kot. Collaborative learning of gesture recognition and 3d hand pose estimation with multi-order feature analysis. In *European Conference on Computer Vision*, pages 769–786. Springer, 2020. 1, 2, 3, 5, 6
- [53] Xitong Yang, Haoqi Fan, Lorenzo Torresani, Larry S Davis, and Heng Wang. Beyond short clips: End-to-end video-level learning with collaborative memories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7567–7576, 2021. 2
- [54] Baowen Zhang, Yangang Wang, Xiaoming Deng, Yinda Zhang, Ping Tan, Cuixia Ma, and Hongan Wang. Interacting two-hand 3d pose and shape reconstruction from single color image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11354–11363, 2021. 2
- [55] Xiong Zhang, Qiang Li, Hong Mo, Wenbo Zhang, and Wen Zheng. End-to-end hand mesh recovery from a monocular rgb image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2354–2364, 2019. 2
- [56] Christian Zimmermann and Thomas Brox. Learning to estimate 3d hand pose from single rgb images. In *Proceedings of the IEEE international conference on computer vision*, pages 4903–4911, 2017. 1, 2, 5