

# SpatioTemporal Learning for Human Pose Estimation in Sparsely-Labeled Videos

Yingying Jiao<sup>1,2\*</sup>, Zhigang Wang<sup>3\*†</sup>, Sifan Wu<sup>1,2</sup>, Shaojing Fan<sup>4†</sup>,  
Zhenguang Liu<sup>5,6†</sup>, Zhuoyue Xu<sup>3</sup>, Zheqi Wu<sup>3</sup>

<sup>1</sup>College of Computer Science and Technology, Jilin University

<sup>2</sup>Key Laboratory of Symbolic Computation and Knowledge Engineering of Ministry of Education, Jilin University

<sup>3</sup>College of Computer Science and Technology, Zhejiang Gongshang University

<sup>4</sup>School of Computing, National University of Singapore

<sup>5</sup>The State Key Laboratory of Blockchain and Data Security, Zhejiang University

<sup>6</sup>Hangzhou High-Tech Zone (Binjiang) Institute of Blockchain and Data Security

jiaoyy21@mails.jlu.edu.cn, wangzhigang2024@gmail.com, wusifan2021@gmail.com, fanshaojing@gmail.com,  
liuzhenguang2008@gmail.com, 1069516849xzyy@gmail.com, chasewoo17@gmail.com

## Abstract

Human pose estimation in videos remains a challenge, largely due to the reliance on extensive manual annotation of large datasets, which is expensive and labor-intensive. Furthermore, existing approaches often struggle to capture long-range temporal dependencies and overlook the complementary relationship between temporal pose heatmaps and visual features. To address these limitations, we introduce STDPose, a novel framework that enhances human pose estimation by learning spatiotemporal dynamics in sparsely-labeled videos. STDPose incorporates two key innovations: 1) A novel Dynamic-Aware Mask to capture long-range motion context, allowing for a nuanced understanding of pose changes. 2) A system for encoding and aggregating spatiotemporal representations and motion dynamics to effectively model spatiotemporal relationships, improving the accuracy and robustness of pose estimation. STDPose establishes a new performance benchmark for both video pose propagation (*i.e.*, propagating pose annotations from labeled frames to unlabeled frames) and pose estimation tasks, across three large-scale evaluation datasets. Additionally, utilizing pseudo-labels generated by pose propagation, STDPose achieves competitive performance with only 26.7% labeled data.

## 1 Introduction

In recent years, visual perception tasks (Dosovitskiy et al. 2020; Kirillov et al. 2023; Shuai et al. 2023) have achieved significant research breakthroughs, largely owing to the continuous advancement of model architectures (Vaswani et al. 2017; Liu et al. 2022b) and the rollout of large-scale datasets (Deng et al. 2009; Lin et al. 2014). Correspondingly, *human pose estimation* (Liu et al. 2021a; Sun et al. 2019), as a foundational task in computer vision (Chen et al. 2023; Wang et al. 2022), has flourished over the past few years and is particularly valuable in a wide range of applications, including *sports analytics*, *surveillance*, *augmented reality*, and *human-computer interaction* (Schmidtke et al.

\*These authors contributed equally.

†Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

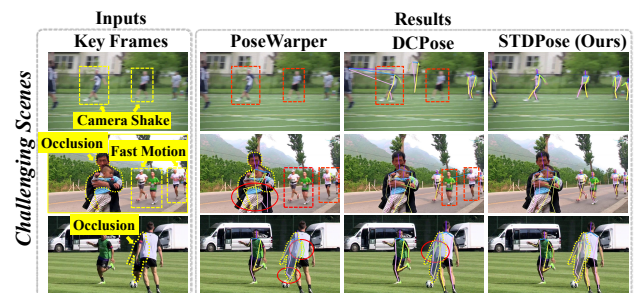


Figure 1: Our model, STDPose, consistently demonstrates high accuracy and robustness in human pose estimation, even in challenging video scenes with **blur** and **occlusion**, thanks to its innovative approach to capturing spatiotemporal information and long-range motion cues. However, state-of-the-art methods like PoseWarper (Bertasius et al. 2019) and DCPose (Liu et al. 2021a) struggle in such scenarios. Red rectangles in our visual comparisons indicate where these methods completely failed to detect **blurred** individuals, while red ellipses highlight their incorrect detections of wrist and ankle joints caused by severe **occlusion**. All results are from models trained on sparsely labeled (*i.e.*, every 7 frames) videos.

2021; Tse et al. 2022; Yang et al. 2023; Liu et al. 2022c; Wu et al. 2024b; Su et al. 2021).

However, the development of multi-person video pose estimation has plateaued, facing numerous challenges that demand innovative solutions. **First**, traditional model architectures, originally designed for static images, continue to struggle with common video issues such as pose occlusion and blur due to fast motion or camera shake. **Second**, most leading methods heavily rely on large-scale benchmark datasets that require intensive time-consuming and labor-intensive pose annotations. Notably, the temporal information in video sequences often shows substantial redundancy, and the spatial changes in pose from one frame to the next are typically minor (Bertasius et al. 2019), resulting in substantial repetitive effort for manual annotation.

Efforts to overcome the inherent limitations of traditional static image-based methods for human pose estimation have propelled substantial progress. Video-focused methods (Bertasius et al. 2019; Wang, Tighe, and Modolo 2020) are designed to capture the temporal dependencies and correlations within video data, which are often neglected by conventional static image methods (Newell, Yang, and Deng 2016; Sun et al. 2019). For instance, PoseWarper (Bertasius et al. 2019) and DCPose (Liu et al. 2021a) model pose residuals for aggregating temporal contexts, utilizing deformable convolutions (Zhu et al. 2019).

Empirically, we observe that in challenging scenes such as pose occlusion and blur, existing methods (Bertasius et al. 2019; Liu et al. 2021a) frequently exhibit poor performance, as illustrated in Figure 1. Our experiments and analyses suggest multifaceted reasons for these shortcomings: **(1)** Current state-of-the-art methods directly aggregate spatiotemporal representations, lacking mathematical constraints guaranteeing that label-relevant clues are extracted and redundant information is reduced, which leads to suboptimal pose estimation outcomes. Furthermore, in scenarios of pose occlusion and blur, temporally distant frames may carry more pertinent supplementary evidence than adjacent frames due to the temporal similarity between consecutive frames, yet current methods (Liu et al. 2021a; Bertasius et al. 2019) fail to effectively harness these long-range spatiotemporal contexts. **(2)** Despite these approaches (Liu et al. 2021a, 2022a) employing deformable convolutions to simulate various receptive fields, they excessively emphasize local pose variations and neglect global spatial correlations, potentially limiting a comprehensive understanding. **(3)** Recent methods (Liu et al. 2022a; Feng et al. 2023a; Wu et al. 2024a) focus solely on temporal features and overlook the integration of complementary information from pose heatmaps, consequently restricting the incorporation of pose annotations in pose propagation.

To address the limitations of existing methods, we present a novel framework, termed **STDPose**, which encodes and aggregates spatiotemporal representations and motion dynamics to learn **S**patio**T**emporal **D**ynamics for Human **P**ose Estimation. The motivation behind STDPose is to propagate pose annotations from labeled (auxiliary) frames to unlabeled (key) frames within sparsely-labeled videos at intervals (*i.e.*, every  $T$  frames), to reduce manual labor. In particular, our STDPose embraces two key components: **(i)** A SpatioTemporal Representation Encoder (STRE) is proposed to address limitations (1) and (3) by collaboratively integrating multi-frame visual features and pose heatmaps to comprehensively capture spatiotemporal dependencies through its two designed specialized submodules. Additionally, a mutual information objective is utilized to supervise cross-frame task-relevant knowledge extraction. **(ii)** To address limitations (1) and (2), we further introduce a novel Dynamic-Aware Mask (DAM) that dedicated to effectively capture long-range motion contexts through a modified sigmoid function. Finally, the SpatioTemporal Dynamics Aggregation module (STDA) aggregates spatiotemporal representations and motion dynamics, enhancing spatial and temporal coherence. Our method achieves state-of-the-art re-

sults for both the video pose propagation and video pose estimation tasks across three benchmark datasets. To summarize, the main contributions of this paper are as follows:

- We develop a novel framework that effectively learns SpatioTemporal Dynamics for Human Pose Estimation in sparsely-labeled videos by encoding and aggregating spatiotemporal representations and motion details, enhancing pose estimation accuracy and robustness through leveraging the temporal continuity between frames.
- We introduce a pioneering technique called the Dynamic-Aware Mask (DAM), which dynamically captures long-range motion clues. This allows for a more nuanced understanding of pose offsets, especially in areas prone to occlusion or blur.
- Our method STDPose significantly advances the field by establishing new state-of-the-art benchmarks for both the pose propagation task and standard pose estimation task across three challenging benchmark datasets: PoseTrack2017, PoseTrack2018, and PoseTrack2021. By automatically generating accurate pose annotations throughout the entire video with only a few manually labeled frames, STDPose reduces the reliance on labor-intensive manual annotation and provides the research community with new insights for video pose estimation.

## 2 Related Work

**Image-based human pose estimation.** Traditional human pose estimation methods rely on pictorial structures (Zhang et al. 2009; Sapp, Toshev, and Taskar 2010) but are limited by handcrafted features and poor generalization. Inspired by advances in deep learning and benchmarks like PoseTrack (Doering et al. 2022) and COCO (Lin et al. 2014), recent deep learning based approaches (Artacho and Savakis 2020; Wei et al. 2016; Xiao, Wu, and Wei 2018) have emerged. HRNet (Sun et al. 2019) excels with high-resolution features, and vision transformers (Dosovitskiy et al. 2020; Xu et al. 2022) show promising performance. However, image-based methods struggle with video-specific challenges like blur and occlusion due to a lack of temporal context. This work advances by focusing on video-based pose estimation to address these dynamic issues effectively.

**Video-based human pose estimation.** To address the shortcomings of image-based approaches in videos, video-based methods (Bertasius et al. 2019; Liu et al. 2021a, 2022a; Feng et al. 2023a; He and Yang 2024; Jin, Lee, and Lee 2022) focus on learning temporal information from neighboring frames. Techniques such as calculating dense optical flow (Pfister, Charles, and Zisserman 2015; Ilg et al. 2017; Xiu et al. 2018) between frames use enriched flow representations to enhance prediction accuracy. DefTrack (Wang, Tighe, and Modolo 2020) integrates 3D-HRNet and 3D convolutions for slight improvements in pose sequence estimation. However, existing methods struggle with three main issues: i) dependence on extensively labeled data; ii) difficulty capturing long-range temporal dependencies; and iii) separate handling of pose heatmaps and visual features. Our approach aims to capture long-range motion

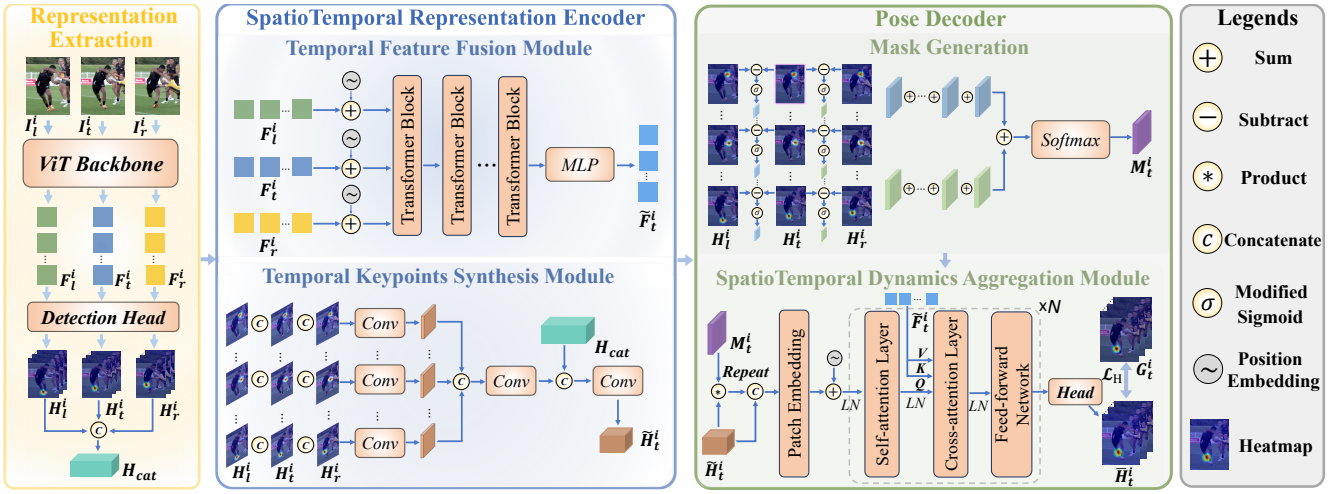


Figure 2: The overall pipeline of our STDPose framework. Given an input sequence  $\langle I_l^i, I_t^i, I_r^i \rangle$ , our goal is to estimate the human pose of the key frame  $I_t^i$ .

dynamics in sparsely-labeled videos by modeling and aggregating temporal representations and motion context simultaneously. This enhances pose detection accuracy and reduces the need for dense annotations.

### 3 Proposed Method

**Problem formulation.** Following the top-down pose estimation paradigm, we initially utilize an object detector to extract the bounding box for each individual person in a video frame  $I_t$ . Each bounding box is then enlarged by 25% to crop the same individual person in key frame  $I_t$  and two auxiliary frames  $I_l$  and  $I_r$ .  $I_l$  and  $I_r$  represent the key frame’s left and right frames in the video sequence, respectively. Note that the time interval  $T$  is set to 7, meaning there are 6 frames between the two auxiliary frames. In this manner, we obtain a cropped image sequence  $\mathcal{S}_t^i = \langle I_l^i, I_t^i, I_r^i \rangle$  for person  $i$ . Given  $\mathcal{S}_t^i$ , our goal is to estimate the pose in  $I_t^i$ .

**Method overview.** As shown in Figure 2, our framework consists of three main components: a Representation Extraction module, a SpatioTemporal Representation Encoder (STRE) (Sec 3.1), and a Pose Decoder (Sec 3.2). (1) Specifically, we first perform spatial feature extraction on the image sequence  $\mathcal{S}_t^i$  to obtain a feature sequence  $\mathcal{F}_t^i = \langle F_l^i, F_t^i, F_r^i \rangle$ , and then utilize a head to convert  $\mathcal{F}_t^i$  into a keypoint heatmap sequence  $\mathcal{H}_t^i = \langle H_l^i, H_t^i, H_r^i \rangle$ . (2) We further feed  $\mathcal{F}_t^i$  and  $\mathcal{H}_t^i$  into the TFF and TKS of the STRE, respectively, to implement feature fusion and keypoint merging, resulting in a fused feature  $\tilde{F}_t^i$  and a merged keypoint heatmaps  $\tilde{H}_t^i$ . (3) Following this, we apply a novel modified sigmoid function to highlight the motion regions and perform a channel compression operation to obtain a dynamic-aware mask  $M_t^i$ . Finally, STDA utilizes a cross-attention algorithm to aggregate  $\tilde{F}_t^i$ ,  $\tilde{H}_t^i$ , and  $M_t^i$  into the final representation, which is then passed to a detection head to output the pose estimation  $\tilde{H}_t^i$ . In the following sections,

we will elaborate on these components in detail.

#### 3.1 SpatioTemporal Representation Encoder

Leveraging the established excellence of Vision Transformers (Dosovitskiy et al. 2020; Liu et al. 2021b) (ViT) in extracting spatial features, we choose ViT as the backbone of our model to extract pose features from the input sequence  $\mathcal{S}_t^i$ . However, we encounter two significant challenges:

(1) While ViT is adept at capturing global spatial dependencies, it struggles with visual tasks such as video pose estimation, which require precise localization of keypoints and often involve subtle local pose variations within videos. To overcome this limitation, we introduce a SpatioTemporal Representation Encoder (STRE) that models both temporal features and temporal pose heatmaps concurrently. The pose heatmaps provide essential local spatial information to accurately pinpoint the locations of individual joints, complemented by global semantic information from the features to effectively handle complex scenes, such as pose occlusion and blur.

(2) Additionally, directly fusing these temporal representations can introduce a substantial amount of task-irrelevant information, particularly because the temporal distance between auxiliary frames and key frames in sparsely-labeled videos can be significant. Drawing inspiration from prior works (Liu et al. 2022a; Hjelm et al. 2018; Tian et al. 2021), we employ a mutual information objective to refine the process of temporal information extraction, ensuring that only relevant data contributes to pose estimation.

These innovations mark our contribution in enhancing the accuracy of video pose estimation, addressing both the challenge of precise localization and the effective integration of temporal dynamics.

**Temporal Feature Fusion (TFF) module.** We feed the  $\mathcal{F}_t^i = \langle F_l^i, F_t^i, F_r^i \rangle$  obtained from the backbone into the Temporal Feature Fusion module to output  $\tilde{F}_t^i$ . Specifically,

we first add a new learnable position embedding  $PE$  to each visual feature and concatenate all the features together. We then feed them into cascaded Transformer blocks. Each block contains a multi-head self-attention layer and a feed-forward neural network. Finally, a Multilayer Perceptron is applied to aggregate all the encoded features to obtain  $\tilde{F}_t^i$ . In summary, the Temporal Feature Fusion module is capable of effectively aggregating multi-frame temporal features, which markedly enhances the extraction of spatiotemporal information, making STDPose superior to existing methods that perform feature alignment (Liu et al. 2022a) and representation difference learning (Feng et al. 2023a).

**Temporal Keypoints Synthesis (TKS) module.** While there is an inherent spatial correlation among the adjacent joints of the human body, the temporal trajectories of each individual joint possess a degree of independence (He and Yang 2024). Thus, given a keypoint heatmap sequence  $\mathcal{H}_t^i = \langle H_l^i, H_t^i, H_r^i \rangle$ , we first merge each keypoint temporally using convolutional blocks, then merge all keypoints in the spatial dimension, and finally concatenate them with the sequence  $\mathcal{H}_t^i$  for the final synthesis to obtain the merged heatmaps  $\tilde{H}_t^i$ . By such, our framework is able to learn spatiotemporal poses by synthesizing temporal keypoint heatmaps stepwise from different dimensions, which is a significant advancement over the previous approaches that involved the simplistic aggregation of pose information (Liu et al. 2021a; Bertasius et al. 2019).

**Mutual Information (MI) objective.** Directly modeling temporal pose and feature yields a significant amount of task-irrelevant clues in sparsely-labeled videos. To effectively extract temporal information, we introduce a mutual information (MI) objective. Within this framework, our main objective for learning effective temporal features and temporal poses can be formulated as:

$$\max \left[ \mathcal{I} \left( y_t^i; \tilde{F}_t^i \mid F_t^i \right) + \mathcal{I} \left( y_t^i; \tilde{H}_t^i \mid H_t^i \right) \right], \quad (1)$$

where  $y_t^i$  denotes the pose label. The terms  $\mathcal{I} \left( y_t^i; \tilde{F}_t^i \mid F_t^i \right)$  and  $\mathcal{I} \left( y_t^i; \tilde{H}_t^i \mid H_t^i \right)$  each represent the measure of task-relevant information contained within the fused feature  $\tilde{F}_t^i$  and the merged heatmaps  $\tilde{H}_t^i$ , respectively, that is in addition to the information already present in  $F_t^i$  and  $H_t^i$ . Due to the difficulty in calculating mutual information (Liu et al. 2022a), we have simplified the formula, as detailed in Appendix. The proposed MI objective  $\mathcal{L}_{MI}$  is as follows:

$$\mathcal{L}_{MI} = -\alpha \cdot \mathcal{I} \left( y_t^i; \tilde{F}_t^i \mid F_t^i \right) - \beta \cdot \mathcal{I} \left( y_t^i; \tilde{H}_t^i \mid H_t^i \right), \quad (2)$$

where  $\alpha$  and  $\beta$  are hyperparameters to balance the ratio of different MI loss terms.

### 3.2 Pose Decoder

Instead of merely employing convolutional operations to combine keypoint heatmaps with features, which often leads to suboptimal outcomes due to a narrow focus on local interactions, we propose a more innovative strategy. Drawing on

insights from recent multimodal research (Rombach et al. 2022), we formulate the integration of keypoint heatmaps and visual features as a quasi-multimodal task using a cross-attention algorithm. This algorithm excels at both local positioning and global searching through its robust similarity computation mechanism, allowing the framework to capture broader spatial interdependencies that are crucial for accurate pose estimation in complex video scenarios.

Moreover, temporal inconsistencies across frames are frequently encountered in sparsely-labeled videos. To tackle this challenge, we introduce a novel concept: the Dynamic-Aware Mask (DAM). This mask is designed to learn motion context effectively, thereby highlighting areas of movement and enhancing the reliability and precision of our framework. This innovative approach allows the framework to capture long-range motion clues, thus enabling a more nuanced understanding of pose dynamics, especially in challenging scenarios prone to occlusion or blur.

**Dynamic-Aware Mask (DAM) generation.** We are the first to propose the Dynamic-Aware Mask (DAM), designed to capture subtle motion clues in complex spatiotemporal interaction scenarios and offer a nuanced comprehension of pose dynamics. We first obtain forward and backward pose residuals through heatmap subtraction, then activate motion areas using a modified sigmoid function on each channel. We then compress the channels into a single channel and add the forward and backward feature maps with weighting. Finally, we utilize a softmax function to derive the mask  $M_t^i$ . The proposed modified sigmoid function as follows:

$$\mathbf{Sig}_m(x, k, \theta) = \frac{1}{1 + e^{-k \cdot (|x| - \theta)}}, \quad (3)$$

where  $\mathbf{Sig}_m(\cdot)$  is the modified sigmoid function.  $x$  is the input of the function.  $k$  is a positive slope parameter that controls the steepness of the function.  $\theta$  is a threshold parameter that determines the value of the function when  $x = 0$ . We aim to keep the value at  $x = 0$  sufficiently small but not zero.  $|x|$  represents the absolute value of  $x$ , ensuring that the function responds identically to both positive and negative values of  $x$ . This is because negative pose residual values also represent significant local spatial changes, which a regular sigmoid function would suppress.  $e$  represents the mathematical constant. We empirically set  $k$  and  $\theta$  to 1.5 and 0.5, respectively. The mask generated by utilizing this sigmoid function can learn local spatial pose differences, thereby capturing subtle motion dynamics and enhancing the robustness of our framework.

**SpatioTemporal Dynamics Aggregation (STDA) module.** Given  $\tilde{F}_t^i$ ,  $\tilde{H}_t^i$ ,  $M_t^i$ , the goal of STDA is to aggregate them and output the final heatmaps  $\tilde{H}_t^i$ . The process starts by performing a dot product between  $\tilde{H}_t^i$  and  $M_t^i$  to obtain the masked heatmaps. These masked heatmaps are then concatenated with  $\tilde{F}_t^i$  and fed into a patch embedding layer that embeds the heatmaps into tokens, also adding position embedding to get  $\hat{H}_t^i$ . Both  $\tilde{F}_t^i$  and  $\hat{H}_t^i$  are then fed into Pose Aggregation (PA) blocks, each consisting of a self-attention layer, a cross-attention layer, and a Feed-forward neural network. We also insert a LayerNorm operation before each

Dataset	Method	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Mean
PoseTrack17 Val Set	Farneback (Farneback 2003)	76.5	82.3	74.3	69.2	80.8	74.8	70.1	75.5
	SimpleBaseline (Xiao, Wu, and Wei 2018)	87.3	88.4	83.6	77.6	83.2	78.4	73.7	82.0
	FlowNet2 (Ilg et al. 2017)	82.7	91.0	83.8	78.4	89.7	83.6	78.1	83.8
	DCPose (Liu et al. 2021a)	91.2	90.8	88.4	84.3	87.7	86.7	83.2	87.7
	PoseWarper (Bertasius et al. 2019)	86.0	92.7	89.5	86.0	91.5	89.1	86.6	88.7
	<b>STDPose (Ours)</b>	<b>92.7</b>	<b>93.1</b>	<b>91.4</b>	<b>88.1</b>	<b>91.9</b>	<b>90.3</b>	<b>88.1</b>	<b>90.9</b>
PoseTrack18 Val Set	SimpleBaseline (Xiao, Wu, and Wei 2018)	82.7	80.1	72.7	66.1	72.1	69.7	65.6	73.4
	PoseWarper (Bertasius et al. 2019)	87.0	88.5	84.8	80.4	81.4	82.1	79.9	83.7
	DCPose (Liu et al. 2021a)	88.9	89.0	85.6	81.8	84.8	82.7	80.3	85.0
	<b>STDPose (Ours)</b>	<b>90.4</b>	<b>91.6</b>	<b>87.6</b>	<b>85.8</b>	<b>86.9</b>	<b>87.0</b>	<b>85.8</b>	<b>88.0</b>
PoseTrack21 Val Set	SimpleBaseline (Xiao, Wu, and Wei 2018)	82.2	77.6	71.1	64.1	68.6	65.0	59.6	70.6
	PoseWarper (Bertasius et al. 2019)	88.6	87.0	83.3	79.2	80.7	80.4	77.3	82.8
	DCPose (Liu et al. 2021a)	88.1	87.3	82.8	79.7	83.0	79.6	78.2	83.1
	<b>STDPose (Ours)</b>	<b>91.5</b>	<b>90.4</b>	<b>87.2</b>	<b>83.3</b>	<b>85.9</b>	<b>85.2</b>	<b>83.2</b>	<b>86.9</b>

Table 1: The results of video pose propagation on PoseTrack2017 (Iqbal, Milan, and Gall 2017), PoseTrack2018 (Andriluka et al. 2018), and PoseTrack2021 (Doering et al. 2022) datasets. Same as PoseWarper (Bertasius et al. 2019), all time intervals  $T$  are set to 7, *i.e.*, pose annotations are given every 7 frames. The evaluation metric is mean Average Precision (mAP).

layer. Ultimately, a classic pose detection head is engaged to upsample the features emanating from the last block, yielding the final predicted heatmaps  $\bar{H}_t^i$ . Overall, the SpatioTemporal Dynamics Aggregation module integrates spatiotemporal representations and performs global-local learning of spatiotemporal dynamics to enhance the performance of pose detection using a cross-attention algorithm instead of deformable convolutions (Zhu et al. 2019; Dai et al. 2017).

### 3.3 Loss Functions

Our loss functions consist of two portions. (1) We employ the common pose heatmap loss  $\mathcal{L}_H$  to supervise the learning of the final pose heatmaps  $\bar{H}_t^i$ :

$$\mathcal{L}_H = \left\| \bar{H}_t^i - G_t^i \right\|_2^2, \quad (4)$$

where  $G_t^i$  denotes the ground-truth heatmaps. (2) Furthermore, We adopt the MI loss  $\mathcal{L}_{MI}$  from Eq. 2 to supervise the extraction of temporal information. The total loss function  $\mathcal{L}_{total}$  is given by:

$$\mathcal{L}_{total} = \mathcal{L}_H + \mathcal{L}_{MI}. \quad (5)$$

## 4 Experiments

### 4.1 Experimental Settings

We carried out thorough evaluations for video pose propagation and video pose estimation tasks on three popular benchmarks: PoseTrack2017 (Iqbal, Milan, and Gall 2017), PoseTrack2018 (Andriluka et al. 2018), and PoseTrack21 (Doering et al. 2022). The videos in these datasets feature diverse challenges, such as crowded scenes and rapid movements. The input image size is  $256 \times 192$ . We utilize a standard Vision Transformer (Dosovitskiy et al. 2020) pretrained on the COCO dataset (Lin et al. 2014) as the backbone network of our STDPose framework. We set the parameters  $\alpha$  to 0.1 and  $\beta$  to 0.01 in Eq. 2, and have not densely tuned them. We evaluate our model using the standard pose estimation metric, average precision (AP), by initially calculating the AP for each joint and subsequently deriving the model’s overall performance through the mean average precision (mAP) across all joints.

### 4.2 Comparison on Pose Propagation

We apply our model to the video pose propagation task, *i.e.*, propagating poses across time from a few labeled frames. Specifically, during training, every 7th frame of the training videos acts as an auxiliary frame (*i.e.*, there are 6 key frames between each pair of auxiliary frames) with no pose annotation. Subsequently, during inference, we provide pose annotations for the auxiliary frames to facilitate pose propagation from labeled (auxiliary) frames to all unlabeled frames. We compare our model with several state-of-the-art methods, including DCPose (Liu et al. 2021a), and PoseWarper (Bertasius et al. 2019), among others. We test the performance of applying DCPose to sparsely-labeled videos based on its open-source release. Regrettably, a direct quantitative comparison with recent studies (Liu et al. 2022a; Feng et al. 2023a) is not possible due to their inherent limitations in aggregating pose labels and propagating poses within sparsely-labeled videos. The experimental results on three benchmark datasets are in Table 1.

**PoseTrack2017.** The proposed STDPose consistently surpasses existing methods, achieving an mAP of 90.9. Our model achieves a 2.2 mAP gain over the previous state-of-the-art approach PoseWarper (Bertasius et al. 2019). Especially, we obtain promising improvements for the more challenging joints (*i.e.*, wrist, ankle), with an mAP of 88.1 ( $\uparrow 2.1$ ) for wrists and an mAP of 88.1 ( $\uparrow 1.5$ ) for ankles.

**PoseTrack2018.** STDPose surpasses the previous state-of-the-art method DCPose (Liu et al. 2021a), attaining an mAP of 88.0 ( $\uparrow 3.0$ ), with an mAP of 85.8, 87.0, and 85.8 for the wrist, knee, and ankle, respectively. Our model outperforms DCPose and PoseWarper (Bertasius et al. 2019), which only process heatmaps, demonstrating the contribution of concurrently integrating both heatmaps and features in our approach.

**PoseTrack2021.** STDPose achieves a 86.9 mAP—3.8 higher than DCPose (Liu et al. 2021a). It greatly surpasses other methods on wrists and ankles, showcasing its effectiveness in tackling challenging scenarios where these joints are often blurred or occluded due to pose occlusion and rapid movement.

Dataset	Method	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Mean
PoseTrack17 Val Set	PoseFlow (Xiu et al. 2018)	66.7	73.3	68.3	61.1	67.5	67.0	61.3	66.5
	PoseWarper (Bertasius et al. 2019)	81.4	88.3	83.9	78.0	82.4	80.5	73.6	81.2
	DCPose (Liu et al. 2021a)	88.0	88.7	84.1	78.4	83.0	81.4	74.2	82.8
	FAMI-Pose (Liu et al. 2022a)	89.6	90.1	86.3	80.0	84.6	83.4	77.0	84.8
	DSTA (He and Yang 2024)	89.3	90.6	87.3	82.6	84.5	85.1	77.8	85.6
	TDMI (Feng et al. 2023a)	<b>90.6</b>	91.0	87.2	81.5	85.2	84.5	78.7	85.9
	DiffPose (Feng et al. 2023b)	89.0	91.2	87.4	83.5	85.5	87.2	80.2	86.4
	<b>STDPose (Ours)</b>	89.7	<b>91.3</b>	<b>88.5</b>	<b>84.7</b>	<b>88.7</b>	<b>87.9</b>	<b>80.6</b>	<b>87.4</b>
PoseTrack18 Val Set	PoseWarper (Bertasius et al. 2019)	79.9	86.3	82.4	77.5	79.8	78.8	73.2	79.7
	DCPose (Liu et al. 2021a)	84.0	86.6	82.7	78.0	80.4	79.3	73.8	80.9
	FAMI-Pose (Liu et al. 2022a)	85.5	87.7	84.2	79.2	81.4	81.1	74.9	82.2
	DiffPose (Feng et al. 2023b)	85.0	87.7	84.3	81.5	81.4	82.9	77.6	83.0
	DSTA (He and Yang 2024)	85.9	88.8	85.0	81.1	81.5	83.0	77.4	83.4
	TDMI (Feng et al. 2023a)	<b>86.7</b>	<b>88.9</b>	85.4	80.6	82.4	82.1	77.6	83.6
	<b>STDPose (Ours)</b>	84.9	88.3	<b>85.5</b>	<b>82.3</b>	<b>85.9</b>	<b>84.9</b>	<b>79.6</b>	<b>84.5</b>
PoseTrack21 Val Set	DCPose (Liu et al. 2021a)	83.2	84.7	82.3	78.1	80.3	79.2	73.5	80.5
	FAMI-Pose (Liu et al. 2022a)	83.3	85.4	82.9	78.6	81.3	80.5	75.3	81.2
	DiffPose (Feng et al. 2023b)	84.7	85.6	83.6	80.8	81.4	83.5	<b>80.0</b>	82.9
	DSTA (He and Yang 2024)	<b>87.5</b>	87.0	84.2	81.4	82.3	82.5	77.7	83.5
	TDMI (Feng et al. 2023a)	86.8	87.4	85.1	81.4	83.8	82.7	78.0	83.8
	<b>STDPose (Ours)</b>	84.4	<b>88.5</b>	<b>85.3</b>	<b>82.3</b>	<b>85.5</b>	<b>84.1</b>	<b>80.0</b>	<b>84.3</b>

Table 2: Comparisons with the state-of-the-art methods for video pose estimation on the validation sets of the PoseTrack2017 (Iqbal, Milan, and Gall 2017), PoseTrack2018 (Andriluka et al. 2018), and PoseTrack2021 (Doering et al. 2022) datasets. Note that during training, we aggregate temporal information from neighboring frames (*i.e.*, one frame to the left and one to the right), and during inference, the pose labels of neighboring frames are not provided.

### 4.3 Comparison on Pose Estimation

**Pose estimation trained on full video annotation.** Unlike pose propagation, during pose estimation training, the auxiliary frames are provided from neighboring frames, and no auxiliary frame pose labels are required during inference. Notably, similar to DCPose (Liu et al. 2021a), we only require two auxiliary frames in training, whereas recent methods like (He and Yang 2024; Feng et al. 2023a) require four.

As shown in Table 2, STDPose outperforms all comparing methods across three datasets, achieving a 87.4 mAP on PoseTrack2017 (1.0 higher than DiffPose (Feng et al. 2023b)) and a 0.9 mAP improvement on PoseTrack2018 over TDMI (Feng et al. 2023a) (with a significant 2.0 mAP increase for the ankle joint). STDPose also tops on the PoseTrack2021 dataset with a 84.3 mAP. Our method excels in detecting challenging joints like wrists and ankles. This is attributed to the combination of the global receptive field provided by the attention mechanism and the dynamic awareness facilitated by the DAM. The superior performance in complex interaction scenarios demonstrates the robustness of our model.

**Pose estimation trained with pseudo-labels generated by pose propagation.** We further demonstrate the effectiveness of our pose propagation model in enhancing pose estimation on sparsely-labeled videos. Specifically, we train our model using a combination of manual annotations and pseudo-labels generated by the pose propagation model on the PoseTrack2017 training set. By varying parameter  $T$ , we control the proportion of manually-labeled frames, with  $T=2$  indicating a 50/50 split. We then evaluate the pose estimation performance on PoseTrack2017 validation set.

As shown in Table 3, pseudo-labels generated from pose

propagation significantly improves pose estimation when dealing with sparsely-labeled videos. Our model achieves 84.3 mAP at  $T=4$ , close to FAMI-Pose (Liu et al. 2022a). Notably, at  $T=2$ , **our model excels over FAMI-Pose (Liu et al. 2022a), achieving 85.2 mAP with only 50% of the manually-labeled frames, demonstrating superior performance with only half the labeled data.** Our model also outperforms FAMI-Pose in ankle joint detection accuracy at  $T=4$ . These results demonstrate that high accuracy can be achieved with minimal labeled data, reducing the need for large, densely annotated datasets. For brevity, we compare only with the classic method FAMI-Pose (Liu et al. 2022a) in Table 3, while comparisons with other state-of-the-art methods are available in Table 2.

### 4.4 Ablation Study

We first perform ablation experiments to examine the influence of each component in our proposed method STDPose for pose propagation task. We also examine the performance changes in pose propagation as the time interval  $T$  increases or decreases. All the ablation studies are conducted on the PoseTrack2017 validation set.

**Study on components of STDPose.** We conduct a comprehensive evaluation of each component in our proposed STDPose framework, presenting the quantitative results in Table 4. **(a)** As a baseline, we incorporate a pose detection head after the ViT backbone. **(b)** We add only the Temporal Feature Fusion (TFF) module. **(c)** We include only the Temporal Keypoints Synthesis (TKS) module. **(d)** Both the TFF and TKS modules are added, utilizing the SpatioTemporal Dynamics Aggregation (STDA) module to aggregate spatiotemporal representations. **(e)** The Dynamic-Aware Mask

Model	$T$	Labeled Frame Ratio	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Mean
FAMI-Pose (Liu et al. 2022a)	-	<b>100%</b>	89.6	90.1	86.3	80.0	84.6	83.4	77.0	<b>84.8</b>
STDTPose (Ours)	$T = 7$	16.7%	87.1	85.9	82.8	79.6	83.1	82.0	76.3	82.4 ( $\downarrow$ 2.4)
	$T = 4$	<b>26.7%</b>	88.3	87.8	84.1	81.5	85.3	84.2	79.0	<b>84.3</b> ( $\downarrow$ 0.5)
	$T = 2$	<b>50.0%</b>	88.9	89.0	85.4	82.6	86.6	84.9	79.6	<b>85.2</b> ( $\uparrow$ 0.4)
	-	100.0%	89.7	91.3	88.5	84.7	88.7	87.9	80.6	87.4 ( $\uparrow$ 2.6)

Table 3: The results of the pose estimation model trained with pseudo-labels generated by a pose propagation model at different  $T$  on the PoseTrack2017 validation set.

Methods	$T$	Labeled Frame Ratio	Components					mAP (%)
			TFF	TKS	DAM	STDA	MI	
(a)	-	0.0%						86.8
(b)	7	16.7%	✓					88.4 ( $\uparrow$ 1.6)
(c)	7	16.7%		✓				88.6 ( $\uparrow$ 1.8)
(d)	7	16.7%	✓	✓		✓		89.6 ( $\uparrow$ 2.8)
(e)	7	16.7%	✓	✓	✓	✓		90.3 ( $\uparrow$ 3.5)
(f)	7	16.7%	✓	✓	✓	✓	✓	90.9 ( $\uparrow$ 4.1)
STDTPose (Ours)	15	10.0%	✓	✓	✓	✓	✓	83.6
	9	13.3%	✓	✓	✓	✓	✓	90.1
	5	23.3%	✓	✓	✓	✓	✓	92.5
	3	36.7%	✓	✓	✓	✓	✓	96.8
	2	50.0%	✓	✓	✓	✓	✓	<b>98.1</b>

Table 4: The upper half of the table presents an ablation study on various components within our STDTPose method, conducted on the PoseTrack2017 (Iqbal, Milan, and Gall 2017) validation set for the pose propagation task. The lower half illustrates the impact of different time intervals  $T$  on performance.

(DAM) is introduced to effectively learn motion context. (f) We incorporate the Mutual Information objective into the loss function to form the complete STDTPose. Results reveal that combining TFF and TKS offers superior performance compared to using each module individually. Additionally, incorporating DAM results in a performance improvement of 0.7 mAP. Lastly, the integration of MI allows our framework to achieve the highest performance, improving by 4.1 mAP over the baseline.

**Study on different time intervals.** Additionally, we assess the performance of our method across varying time intervals  $T$ , as detailed in Table 4. Considering the average video contains about 30 frames, setting  $T$  to 7 results in roughly 5 evenly distributed labeled frames per video. Increasing  $T$  reduces the number of required labeled frames: at  $T=15$  a video contains only 3 labeled frames, representing just 10% of the total. However, this configuration leads to a significant drop in performance. Conversely, when  $T$  is reduced to 2, equating the number of labeled and unlabeled frames, our method reaches an mAP of 98.1, nearly matching a 100% ideal precision of manual annotation. This promising result highlights our method’s ability to efficiently automate video annotation, achieving near-perfect accuracy with just 50% of frames labeled manually.

#### 4.5 Discussion

In summary, STDTPose consistently excels in both pose propagation and pose estimation across all datasets, particularly

in challenging scenarios with blurred or occluded joints like wrists and ankles. Our pose propagation model significantly boosts pose estimation on sparsely-labeled videos, achieving competitive performance with minimal labeled data. These results showcase STDTPose’s effectiveness in learning video spatiotemporal clues, not only enhancing performance but also tremendously reducing manual labeling efforts.

## 5 Conclusion

We introduce STDTPose, a novel architecture for video pose propagation and pose estimation. STDTPose innovatively models temporal features and pose heatmaps simultaneously, setting a new standard in the field. We also introduce a theoretical advancement, the Dynamic-Aware Mask, which is specifically designed to learn and interpret motion dynamics effectively, especially in challenging occasions such as occlusion and blur. STDTPose pushes the edge in performance on three benchmark datasets, while also reducing the need for extensive manual video annotations.

## Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. 62372402), and the Key R&D Program of Zhejiang Province (No. 2023C01217).

## References

- Andriluka, M.; Iqbal, U.; Insafutdinov, E.; Pishchulin, L.; Milan, A.; Gall, J.; and Schiele, B. 2018. Posetrack: A benchmark for human pose estimation and tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5167–5176.
- Artacho, B.; and Savakis, A. 2020. Unipose: Unified human pose estimation in single images and videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7035–7044.
- Bertasius, G.; Feichtenhofer, C.; Tran, D.; Shi, J.; and Torresani, L. 2019. Learning temporal pose estimation from sparsely-labeled videos. *Advances in neural information processing systems*, 32.
- Chen, X.; Wang, X.; Zhou, J.; Qiao, Y.; and Dong, C. 2023. Activating more pixels in image super-resolution transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22367–22377.
- Dai, J.; Qi, H.; Xiong, Y.; Li, Y.; Zhang, G.; Hu, H.; and Wei, Y. 2017. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 764–773.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Doering, A.; Chen, D.; Zhang, S.; Schiele, B.; and Gall, J. 2022. Posetrack21: A dataset for person search, multi-object tracking and multi-person pose tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20963–20972.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Farnebäck, G. 2003. Two-frame motion estimation based on polynomial expansion. In *Image Analysis: 13th Scandinavian Conference, SCIA 2003 Halmstad, Sweden, June 29–July 2, 2003 Proceedings 13*, 363–370. Springer.
- Feng, R.; Gao, Y.; Ma, X.; Tse, T. H. E.; and Chang, H. J. 2023a. Mutual information-based temporal difference learning for human pose estimation in video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17131–17141.
- Feng, R.; Gao, Y.; Tse, T. H. E.; Ma, X.; and Chang, H. J. 2023b. DiffPose: SpatioTemporal diffusion model for video-based human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14861–14872.
- He, J.; and Yang, W. 2024. Video-Based Human Pose Regression via Decoupled Space-Time Aggregation. *arXiv preprint arXiv:2403.19926*.
- Hjelm, R. D.; Fedorov, A.; Lavoie-Marchildon, S.; Grewal, K.; Bachman, P.; Trischler, A.; and Bengio, Y. 2018. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*.
- Ilg, E.; Mayer, N.; Saikia, T.; Keuper, M.; Dosovitskiy, A.; and Brox, T. 2017. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2462–2470.
- Iqbal, U.; Milan, A.; and Gall, J. 2017. Posetrack: Joint multi-person pose estimation and tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2011–2020.
- Jin, K.-M.; Lee, G.-H.; and Lee, S.-W. 2022. OTPose: occlusion-aware transformer for pose estimation in sparsely-labeled videos. In *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 3255–3260. IEEE.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Liu, Z.; Chen, H.; Feng, R.; Wu, S.; Ji, S.; Yang, B.; and Wang, X. 2021a. Deep dual consecutive network for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 525–534.
- Liu, Z.; Feng, R.; Chen, H.; Wu, S.; Gao, Y.; Gao, Y.; and Wang, X. 2022a. Temporal feature alignment and mutual information maximization for video-based human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11006–11016.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021b. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022b. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11976–11986.
- Liu, Z.; Wu, S.; Xu, C.; Wang, X.; Zhu, L.; Wu, S.; and Feng, F. 2022c. Copy motion from one to another: Fake motion video generation. *arXiv preprint arXiv:2205.01373*.
- Newell, A.; Yang, K.; and Deng, J. 2016. Stacked hourglass networks for human pose estimation. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, 483–499. Springer.
- Pfister, T.; Charles, J.; and Zisserman, A. 2015. Flowing convnets for human pose estimation in videos. In *Proceedings of the IEEE international conference on computer vision*, 1913–1921.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent

- diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Sapp, B.; Toshev, A.; and Taskar, B. 2010. Cascaded models for articulated pose estimation. In *Computer Vision–ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part II 11*, 406–420. Springer.
- Schmidtke, L.; Vlontzos, A.; Ellershaw, S.; Lukens, A.; Arichi, T.; and Kainz, B. 2021. Unsupervised human pose estimation through transforming shape templates. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2484–2494.
- Shuai, C.; Zhong, J.; Wu, S.; Lin, F.; Wang, Z.; Ba, Z.; Liu, Z.; Cavallaro, L.; and Ren, K. 2023. Locate and verify: A two-stream network for improved deepfake detection. In *Proceedings of the 31st ACM International Conference on Multimedia*, 7131–7142.
- Su, P.; Liu, Z.; Wu, S.; Zhu, L.; Yin, Y.; and Shen, X. 2021. Motion prediction via joint dependency modeling in phase space. In *Proceedings of the 29th ACM international conference on multimedia*, 713–721.
- Sun, K.; Xiao, B.; Liu, D.; and Wang, J. 2019. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5693–5703.
- Tian, X.; Zhang, Z.; Lin, S.; Qu, Y.; Xie, Y.; and Ma, L. 2021. Farewell to mutual information: Variational distillation for cross-modal person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1522–1531.
- Tse, T. H. E.; Kim, K. I.; Leonardis, A.; and Chang, H. J. 2022. Collaborative learning for hand and object reconstruction with attention-guided graph convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1664–1674.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, M.; Tighe, J.; and Modolo, D. 2020. Combining detection and tracking for human pose estimation in videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11088–11096.
- Wang, Y.; Li, K.; Li, Y.; He, Y.; Huang, B.; Zhao, Z.; Zhang, H.; Xu, J.; Liu, Y.; Wang, Z.; et al. 2022. Internvideo: General video foundation models via generative and discriminative learning. *arXiv preprint arXiv:2212.03191*.
- Wei, S.-E.; Ramakrishna, V.; Kanade, T.; and Sheikh, Y. 2016. Convolutional pose machines. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 4724–4732.
- Wu, S.; Chen, H.; Yin, Y.; Hu, S.; Feng, R.; Jiao, Y.; Yang, Z.; and Liu, Z. 2024a. Joint-Motion Mutual Learning for Pose Estimation in Video. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 8962–8971.
- Wu, S.; Liu, Z.; Zhang, B.; Zimmermann, R.; Ba, Z.; Zhang, X.; and Ren, K. 2024b. Do as I Do: Pose Guided Human Motion Copy. *IEEE Transactions on Dependable and Secure Computing*.
- Xiao, B.; Wu, H.; and Wei, Y. 2018. Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision (ECCV)*, 466–481.
- Xiu, Y.; Li, J.; Wang, H.; Fang, Y.; and Lu, C. 2018. Pose Flow: Efficient online pose tracking. *arXiv preprint arXiv:1802.00977*.
- Xu, Y.; Zhang, J.; Zhang, Q.; and Tao, D. 2022. Vitpose: Simple vision transformer baselines for human pose estimation. *Advances in Neural Information Processing Systems*, 35: 38571–38584.
- Yang, Y.; Chen, H.; Liu, Z.; Lyu, Y.; Zhang, B.; Wu, S.; Wang, Z.; and Ren, K. 2023. Action recognition with multi-stream motion modeling and mutual information maximization. *arXiv preprint arXiv:2306.07576*.
- Zhang, X.; Li, C.; Tong, X.; Hu, W.; Maybank, S.; and Zhang, Y. 2009. Efficient human pose estimation via parsing a tree structure based human model. In *2009 IEEE 12th International Conference on Computer Vision*, 1349–1356. IEEE.
- Zhu, X.; Hu, H.; Lin, S.; and Dai, J. 2019. Deformable convnets v2: More deformable, better results. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9308–9316.