

Causal-Inspired Multitask Learning for Video-Based Human Pose Estimation

Haipeng Chen^{1,2}, Sifan Wu^{1,2*}, Zhigang Wang^{3*},
Yifang Yin⁴, Yingying Jiao^{1,2*}, Yingda Lyu^{1,5*}, Zhenguang Liu^{6,7}

¹College of Computer Science and Technology, Jilin University,

²Key Laboratory of Symbolic Computation and Knowledge Engineering of Ministry of Education, Jilin University,

³College of Computer Science and Technology, Zhejiang Gongshang University,

⁴Institute for Infocomm Research (*I²R*), A*STAR,

⁵Public Computer Education and Research Center, Jilin University,

⁶The State Key Laboratory of Blockchain and Data Security, Zhejiang University,

⁷Hangzhou High-Tech Zone (Binjiang) Institute of Blockchain and Data Security,

chenhp@jlu.edu.cn, wusifan2021@gmail.com, wangzhigang2024@gmail.com,

yin_yifang@i2r.a-star.edu.sg, jiaoyy21@mails.jlu.edu.cn, ydlv@jlu.edu.cn, liuzhenguang2008@gmail.com

Abstract

Video-based human pose estimation has long been a fundamental yet challenging problem in computer vision. Previous studies focus on spatio-temporal modeling through the enhancement of architecture design and optimization strategies. However, they overlook the causal relationships in the joints, leading to models that may be overly tailored and thus estimate poorly to challenging scenes. Therefore, adequate causal reasoning capability, coupled with good interpretability of model, are both indispensable and prerequisite for achieving reliable results. In this paper, we pioneer a causal perspective on pose estimation and introduce a causal-inspired multitask learning framework, consisting of two stages. *In the first stage*, we try to endow the model with causal spatio-temporal modeling ability by introducing two self-supervision auxiliary tasks. Specifically, these auxiliary tasks enable the network to infer challenging keypoints based on observed keypoint information, thereby imbuing causal reasoning capabilities into the model and making it robust to challenging scenes. *In the second stage*, we argue that not all feature tokens contribute equally to pose estimation. Prioritizing causal (keypoint-relevant) tokens is crucial to achieve reliable results, which could improve the interpretability of the model. To this end, we propose a Token Causal Importance Selection module to identify the causal tokens and non-causal tokens (*e.g.*, background and objects). Additionally, non-causal tokens could provide potentially beneficial cues but may be redundant. We further introduce a non-causal tokens clustering module to merge the similar non-causal tokens. Extensive experiments show that our method outperforms state-of-the-art methods on three large-scale benchmark datasets.

Introduction

Estimating human poses from videos is a fundamental topic in artificial intelligence, which aims to identify and localize anatomical keypoints on the human body. In recent years,

*Corresponding to: Yingying Jiao, Sifan Wu, Zhigang Wang, Yingda Lyu.

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

this task has garnered an increasing interest from researchers and industry, accelerating advances in human-centric applications in diverse scenes *such as* action recognition (Yang et al. 2023), motion prediction (Su et al. 2021) and motion transfer (Liu et al. 2022b; Wu et al. 2024b).

With the continual breakthroughs of deep learning algorithms and models (Yao, Li, and Xiao 2024; Hu et al. 2024b), AI has achieved success in various fields (Shuai et al. 2023; Chen et al. 2024c; Guan et al. 2024b). For pose estimation, one line of work focuses on designing different network structures. (Artacho and Savakis 2020) adopts CNNs and LSTM to extract the intrinsic motion dynamic of persons. (Feng et al. 2023) proposes feature difference method to capture spatio-temporal dependencies in videos. Another line of work introduces specific loss functions to supervise the network. (Liu et al. 2022a; Feng et al. 2023) propose a mutual information objective to align the features of sequences and obtain an informative representation.

Though promising, there are a few clouds on the horizon for video-based pose estimation. 1) **Robustness**. Over-tailored network structures (Tang et al. 2024; Xu et al. 2024) and the lack of causal perceptual capabilities compromise the robustness of models (Zhang, Ji, and Liu 2023). Challenging scenes such as pose occlusion and video defocus often appear in videos, where the model fails to accurately estimate the human poses. For instance, when multiple people play football, the legs of the persons in the video occlude each other. Existing methods primarily concentrate on the occluded parts and attempt to identify the positions of the occluded keypoints. However, they usually sacrifice an explicit understanding of observed (non-entangled region) visual cues. In other words, they lack an ability to causal reasoning, utilizing observed variables to speculate about unknown (entangled region) variables. 2) **Interpretability**. For pose estimation tasks, each frame is a mixture of causal (keypoint-relevant) and non-causal (*e.g.*, background, objects) factors (Liu et al. 2024). Obviously, the contribution of causal factors is much greater than that of non-causal factors. Existing methods struggle to capture the causal features

that are crucial for identifying the human pose, which deprives the network of model interpretability.

In this work, adopting a causal look at pose estimation, we present a **Causal-inspired Multitask learning Pose estimation framework (CM-Pose)** in a two-stage process. *In the first stage*, we aim to foster the network with causal reasoning ability, enabling it to infer the locations of challenging (occluded or blurred) keypoints based on the observed keypoint information. Specifically, we present a novel multitask learning framework, which introduces different *auxiliary tasks* on top of the *primary pose estimation task*. Central to the idea of auxiliary tasks is to randomly corrupt partial feature tokens and then recover them utilizing normal tokens based on the causal spatio-temporal modeling capability of network. Our framework complements existing pose estimation approaches by emphasizing an effective multitask learning paradigm in a causal view.

Technically, for the challenging scenes of pose occlusion and video defocus, we propose two self-supervision auxiliary tasks: a masked token reconstruction task and a denoising token task. The masked token reconstruction task randomly masks the feature tokens, with the objective of reconstructing the masked tokens. This task simulates the pose occlusion scene where some keypoints are occluded. Similarly, the token denoising task randomly adds Gaussian noise to the feature tokens, with the goal of recovering the corrupted tokens. The task aims to simulate the video defocus scenario where some areas of the image become blurred. Compared with the existing masked/denoising learning paradigms, they always deal with data types such as images (He et al. 2022) and motion sequences (Xu et al. 2023). To the best of our knowledge, we are the first work to propose masked/denoising tokens reconstruction learning to promote the primary pose estimation task. To better cultivate the spatio-temporal causal reasoning ability of multitask network, we introduce a simple yet effective network, a criss-cross spatio-temporal attention network. These tasks share the same network, and the auxiliary tasks are only employed during the training process, incurring no additional computational costs.

In the second stage, we inject interpretability into our model. Not all feature tokens contribute equally to pose estimation. We argue that revealing “which features of images is important for keypoint position” is the key to accurately and reliably estimate human poses. Specifically, we propose a Token Causal Importance Selection module and a Non-causal Tokens Clustering module. Firstly, given coarse tokens from multitask coarse learning, the token causal importance selection module differentiates these tokens into causal (keypoint region) and non-causal (background or objects) tokens based on keypoint tokens attention. While non-causal tokens are not directly related to human pose, they could supply the potential clues and improve the expressivity of the model (Long et al. 2023). Therefore, it is inappropriate to directly discard non-causal tokens. However, non-causal tokens may contain redundant information, such as multiple background tokens with the same semantics. To this end, the non-causal tokens clustering module applies a density peaks clustering algorithm based on k nearest neighbor to effectively cluster similar non-causal tokens and merge

these tokens from the same group into a new token. Finally, causal tokens and merged non-causal tokens are aggregated to informative tokens, which are fed into a pose detection head to estimate the human pose. We believe that these insights will open avenues for future research on video-based human pose estimation.

Contributions. The key contributions of this work are summarized as follows:

- In this paper, we investigate the video-based pose estimation task from a causal perspective, aiming to inject good robustness and interpretability into our model in challenging scenes.
- We propose a causal-inspired multitask learning framework that enables the model to perform causal spatio-temporal modeling on challenging scenes. We further introduce a token causal importance selection module and a non-causal token clustering module to identify causal features and compact (less redundancy) non-causal features, which improves the interpretability of the model.
- Our method achieves state-of-the-art performance on three benchmark datasets, *i.e.*, PoseTrack2017, PoseTrack2018, PoseTrack2021.

Related Works

Vision Transformer on Image-based Pose Estimation. Vision transformer (ViT) (Dosovitskiy et al. 2020) provides an alternative to CNN-based methods for various visual tasks (Tang et al. 2023; Chen et al. 2023, 2024a). (Li et al. 2021) proposes Tokenpose, which first utilizes CNN to extract feature maps and then employs ViT to estimate the human pose in images. (Xu et al. 2022) first proposes to use ViT model without CNN to perform image-level pose estimation. Image-based pose estimation methods have gained perfect performance. However, they are not suitable for video-based pose estimation tasks because they focus only on intra-frame spatial relationships and ignore the abundant inter-frame temporal dependencies.

Video-Based Human Pose Estimation. Existing methods (Jin, Lee, and Lee 2022; Jiao et al. 2022; Wu et al. 2024a) resort to different frameworks or loss functions to model spatio-temporal dependencies of videos. (Song et al. 2017) employ dense optical flow between frames to polish the pose of keyframe. (Feng et al. 2023) introduces feature difference to capture spatial features and temporal dynamics. In addition to designing various network structures, (Liu et al. 2022a) proposes a mutual information loss function to supervise the network. In general, the performance of pose estimation is heavily dependent on the dedicated networks or loss functions. In contrast, we adopt a new training strategy that efficiently trains the network by leveraging causal-inspired multitask learning framework.

Auxiliary Task Learning. Auxiliary task learning is a novel learning paradigm that aims to simulate challenging scenes to make network more robust, imbuing the network with robust information. Auxiliary task learning has demonstrated extraordinary results in various tasks, such as video captioning (Gao et al. 2021) and motion prediction (Xu et al. 2023). To the best of our knowledge, auxiliary task learning

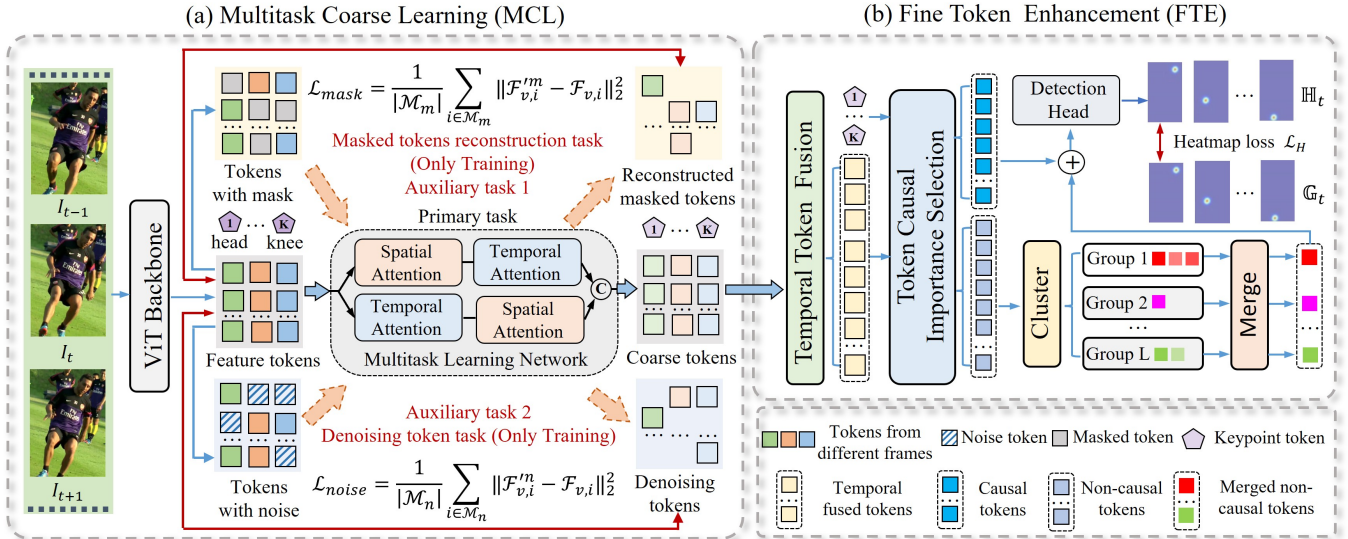


Figure 1: Overall pipeline of our CM-Pose framework. The goal is to identify the pose of keyframe I_t . CM-Pose includes two key components: Multitask Coarse Learning (MCL) and Fine Token Enhancement (FTE). (a) The MCL module consists of three tasks: the primary pose estimation task (middle branch), the masked token reconstruction task (upper branch), and the denoising token task (lower branch). To model the causal spatio-temporal dependencies for these tasks, we design a simple yet effective multitask learning network, a criss-cross spatio-temporal attention network. (b) The FTE module first performs temporal fusion on the coarse tokens. We then decouple the causal and non-causal tokens according to keypoint tokens attention. We further cluster non-causal tokens and merge the tokens from the same group into a new token. Finally, we concatenate the causal tokens and the merged non-causal tokens to estimate the pose \mathbb{H}_t through a detection head.

has not been attempted in pose estimation. In this paper, we propose CM-Pose, which explores the potential of auxiliary task learning in video-based human pose estimation.

Method

Problem Formulation

Given a video sequence, the video-based human pose estimation task aims to identify the pose of all person in every frame. Technically, we first utilize an object detector (Qiao, Chen, and Yuille 2021) to obtain the bounding boxes of each person in video frames. To crop the same person in the video sequence, we then enlarge all bounding boxes by 25% and obtain the video clip of person i , *i.e.*, $\mathcal{I}_t^i = \{I_{t-\omega}^i, \dots, I_t^i, \dots, I_{t+\omega}^i\}$ (where ω being a predefined time span) centered on the keyframe I_t^i . We are interested in modeling the spatio-temporal dependencies in \mathcal{I}_t^i to estimate the pose in I_t^i . To facilitate representation and understanding, we set $\omega = 1$ and omit the superscript i .

Method Overview

As shown in Figure 1, the proposed method CM-Pose comprises two key components: Multitask Coarse Learning (MCL) and Fine Token Enhancement (FTE). (1) We take a causal view to propose MCL, which learns the primary pose estimation task along with extra-designed auxiliary tasks. MCL enable the network to have more comprehensive causal spatio-temporal modeling capabilities and be

robust to challenging scenes. (2) The FTE module categorizes the coarse tokens (generated from primary task in the MCL) into causal and non-causal tokens based on keypoint token attention scores (Guan et al. 2024a), injecting good interpretability into the model. In order to reduce the information redundancy of non-causal tokens, we cluster and merge similar non-causal tokens. Finally, we concatenate the causal tokens and the merged non-causal tokens to jointly infer human pose. In the following, we introduce the two components in detail.

Multitask Coarse Learning

The key to accurately estimating pose is to model the spatial and temporal dependencies of joints based on feature tokens. Existing methods focus on spatio-temporal modeling through the enhancement of architecture design or optimization strategies. However, they ignore the causal correlation in the feature tokens, resulting in models that may be overly tailored and thus estimate poorly to challenging scenes. We consider a novel and interesting research line, a causal-inspired multitask learning framework, that imposes the network to capture more comprehensive causal spatio-temporal dependencies by introducing extra-designed auxiliary tasks. There are three important steps: feature extraction, auxiliary tasks, and multitask learning network.

Feature Extraction. Given a frame sequence $\mathcal{I}_t = \{I_{t-1}, I_t, I_{t+1}\}$, we first utilize a backbone network (Dosovitskiy et al. 2020) to generate the initial feature tokens $\{F_{t-1}, F_t, F_{t+1}\}$. We then concatenate them to obtain $\mathcal{F}_v =$

$F_{t-1} \oplus F_t \oplus F_{t+1}$, $\mathcal{F}_v \in \mathbb{R}^{3N \times D}$, where \oplus represents the concatenate operation, N denotes the number of tokens in a frame and D is the dimension of hidden embedding. Besides, following (Li et al. 2021), we also introduce K additional learnable keypoint tokens $\mathcal{F}_k \in \mathbb{R}^{K \times D}$ to represent K keypoints.

Auxiliary Tasks. To enable the network to perform robust pose estimation in challenging scenes (*i.e.*, pose occlusion and video defocus), we propose two auxiliary tasks, respectively: masked token reconstruction task and denoising token task. Specifically, given initial tokens \mathcal{F}_v , in the masked token reconstruction task, each token of \mathcal{F}_v is masked to zero value with a probability p_{tm} (*i.e.*, $3 \cdot N \cdot p_{tm}$ tokens are masked), obtaining masked tokens \mathcal{F}_v^m . Similarly, in the denoising token task, each token of \mathcal{F}_v has a probability p_{tn} of being added with a Gaussian noise (*i.e.*, $3 \cdot N \cdot p_{tn}$ tokens are added noise). We can obtain the tokens with noise \mathcal{F}_v^n . The auxiliary task requires the network to reconstruct corrupted tokens by utilizing the limited semantic context available from normal tokens, which cultivates the causal spatio-temporal modeling ability of network and perform better to challenging scenes. It is worth noting that introducing additional auxiliary tasks in the learning framework, which share the multitask learning network with primary task, does not increase the model size. During inference, we remove the auxiliary tasks. Besides, the flag of the masked tokens is formulated as:

$$M_m(i) = \begin{cases} 0, & \text{if token } i \text{ is masked,} \\ 1, & \text{otherwise.} \end{cases} \quad (1)$$

Similarly, the flag of token denoising task is M_n , which helps in auxiliary tasks to supervise the network training utilizing reconstruction losses. Mathematically, assuming the masked token set $\mathcal{M}_m = \{i | M_m(i) = 0\}$ and the noised token set $\mathcal{M}_n = \{i | M_n(i) = 0\}$, the loss functions of two auxiliary tasks are formulated as:

$$\mathcal{L}_{mask} = \frac{1}{|\mathcal{M}_m|} \sum_{i \in \mathcal{M}_m} \left\| \mathcal{F}'_{v,i} - \mathcal{F}_{v,i} \right\|_2^2, \quad (2)$$

$$\mathcal{L}_{denoise} = \frac{1}{|\mathcal{M}_n|} \sum_{i \in \mathcal{M}_n} \left\| \mathcal{F}'_{v,i} - \mathcal{F}_{v,i} \right\|_2^2, \quad (3)$$

where \mathcal{F}'_v^m , \mathcal{F}'_v^n , \mathcal{F}_v denote the reconstructed masked tokens, denoising tokens, and the initial tokens, respectively.

Multitask Learning Network. To more effectively model causal spatio-temporal dependencies between joints based on feature tokens, a naive method is to feed the feature tokens into a self-attention module (Chen et al. 2024b). However, there is a natural spatial association between adjacent joints in a human pose, but the temporal trajectory of each joint is independent (He and Yang 2024). This indicates the importance of decoupling the spatial and temporal dependencies between joints. Motivated by the observation and insight, we propose a simple but efficient multitask learning network: a criss-cross spatio-temporal attention network, including a spatio-temporal pathway and a temporal-spatio pathway. The primary task and auxiliary tasks share this network, where the auxiliary tasks promote

the network to infer unknown tokens based on normal tokens, giving the network causal spatio-temporal modeling capabilities. Concretely, the input of these tasks can be obtained by:

$$X_v = \mathcal{F}_v \oplus \mathcal{F}_k, X_v^m = \mathcal{F}_v^m \oplus \mathcal{F}_k, X_v^n = \mathcal{F}_v^n \oplus \mathcal{F}_k, \quad (4)$$

where \oplus denotes the concatenate operation, X_v , X_v^m , and X_v^n are the input tokens of three tasks. \mathcal{F}_v , \mathcal{F}_v^m , \mathcal{F}_v^n and \mathcal{F}_k represent the initial feature tokens, masked tokens, noised tokens, and the keypoint tokens, respectively. The output of multitask learning network is formulated as:

$$\{\mathcal{F}'_v; \mathcal{F}'_k\} = D_T(D_S(X_v)) \oplus D_S(D_T(X_v)), \quad (5)$$

$$\{\mathcal{F}'_v^m; \mathcal{F}'_k^m\} = D_T(D_S(X_v^m)) \oplus D_S(D_T(X_v^m)), \quad (6)$$

$$\{\mathcal{F}'_v^n; \mathcal{F}'_k^n\} = D_T(D_S(X_v^n)) \oplus D_S(D_T(X_v^n)), \quad (7)$$

where D_T and D_S represent the temporal and spatial self-attention module, respectively. \mathcal{F}'_v , \mathcal{F}'_v^m , and \mathcal{F}'_v^n denotes the refined coarse tokens, the reconstructed masked tokens, and the denoising tokens, respectively. \mathcal{F}'_k , \mathcal{F}'_k^m , and \mathcal{F}'_k^n represent the corresponding learned keypoint tokens.

Fine Token Enhancement

In human pose estimation tasks, causal feature tokens are more valuable than the non-causal tokens (such as background, objects, *etc.*). An intuitive approach is to utilize the causal tokens and discard the other tokens to estimate the human pose. However, motivated by the observation and insight in (Xiao et al. 2020; Long et al. 2023), image background could improve the accuracy of vision task due to their potential and implicit relations to the human.

Tokens Causal Importance Selection. To prioritize the causal tokens for pose estimation, we propose tokens causal importance selection module. We first employ a temporal fusion layer to aggregate the multi-frame feature tokens \mathcal{F}'_v to a tokens $\hat{\mathcal{F}}_v$. To endow the interpretability of our method, we decouple the tokens $\hat{\mathcal{F}}_v$ into two categories (*i.e.*, the causal and the non-causal tokens) by comparing the similarity with the learned keypoint tokens \mathcal{F}'_k . Mathematically, we compute the similarity score between the keypoint tokens and all feature tokens as:

$$S^i = \text{Softmax}\left(\frac{\mathbf{Q}_k^i \mathbf{K}_v^T}{\sqrt{D}}\right), \quad (8)$$

where \mathbf{Q}_k^i represents the query vector of the i^{th} keypoint token \mathcal{F}'_k^i , \mathbf{K}_v is the key matrix of feature tokens $\hat{\mathcal{F}}_v$, D is the latent dimension. For each keypoint token \mathcal{F}'_k^i , we select the top- n tokens as causal tokens for i^{th} keypoint token according to similarity scores S^i . For K keypoint tokens, we can obtain nK causal tokens $\hat{\mathcal{F}}_v^c$ and the remaining $N - nK$ tokens are set as non-causal tokens $\hat{\mathcal{F}}_v^{nc}$.

Non-Causal Tokens Clustering. Although non-causal tokens provide implicit and potential clues, several non-causal tokens often correspond to the identical region (*e.g.*, background, objects) and the semantic information may be

redundant. To this end, we propose a non-causal tokens clustering module, which cluster the non-causal tokens $\hat{\mathcal{F}}_v^{nc}$ and then merge the tokens from the identical group into new tokens through importance weight. Specifically, we employ a density peak clustering algorithm based on k-nearest neighbor (DPC-KNN) (Zeng et al. 2022) to cluster feature tokens. DPC-KNN follows two assumptions: (1) the density of the cluster centers is higher than other samples around them and (2) the distance between different cluster centers is far. This derives two concepts: local density ρ and relative distance δ for each non-causal token. Given non-causal tokens, ρ and δ are formulated as:

$$\rho_i = \exp\left(-\frac{1}{k} \sum_{f_j \in \text{KNN}(f_i)} \|f_i - f_j\|_2^2\right), \quad (9)$$

$$\delta_i = \begin{cases} \min_{j: \rho_j > \rho_i} (\|f_i - f_j\|_2) & \text{if } \exists j \text{ s.t. } \rho_j > \rho_i \\ \max_j (\|f_i - f_j\|_2) & \text{otherwise} \end{cases}, \quad (10)$$

where $f_i, f_j \in \hat{\mathcal{F}}_v^{nc}$, and $\hat{\mathcal{F}}_v^{nc}$ represents the non-causal tokens. ρ_i, δ_i denote the local density and relative distance of token f_i , $\text{KNN}(f_i)$ denotes a set of k-nearest neighbors of token f_i . We consider two factors (*i.e.*, ρ_i and δ_i) to determine the cluster center score of token f_i as $\rho_i \times \delta_i$. We choose the top-L tokens with the highest scores as cluster centers, and then assign the remaining tokens to the cluster center with the closest distance.

Each cluster may contain a different number of tokens, and different tokens may have different importance. Inspired by (Long et al. 2023), instead of mindlessly averaging the tokens in the same group, we merge the tokens by a weighted sum. Specifically, we introduce the keypoint token attention to denote importance (Hu et al. 2024a), we merge the same group of tokens into a new token as:

$$\bar{f}_i = \sum_{j \in C_i} s_j f_j, \quad (11)$$

where \bar{f}_i denotes the merged token from C_i , C_i denotes the i^{th} cluster, s_j represents the importance score of token f_j . We concatenate all \bar{f}_i to get merged non-causal tokens $\hat{\mathcal{F}}_v^{nc}$.

Heatmap Generation. Finally, we aggregate causal tokens $\hat{\mathcal{F}}_v^c$ and the merged non-causal tokens $\hat{\mathcal{F}}_v^{nc}$ to obtain a fine feature tokens \mathbb{F}_v as follows:

$$\mathbb{F}_v = \hat{\mathcal{F}}_v^c \oplus \hat{\mathcal{F}}_v^{nc}, \quad (12)$$

where \oplus represents the concatenate operation. \mathbb{F}_v is then fed into a detection head to obtain the pose heatmaps \mathbb{H}_t .

Loss Functions

We use a heatmap loss to supervise the pose estimation:

$$\mathcal{L}_H = \|\mathbb{H}_t - \mathbb{G}_t\|_2^2, \quad (13)$$

where \mathbb{H}_t and \mathbb{G}_t denote the estimated heatmap and the ground truth heatmap, respectively. For auxiliary task learning, we adopt the reconstruction loss \mathcal{L}_{mask} in the Eq. 2 and $\mathcal{L}_{denoise}$ in the Eq. 3 to supervise the *multitask learning network*, aiming to capture more causal spatio-temporal modeling ability. The total loss is given by:

$$\mathcal{L}_{total} = \mathcal{L}_H + \lambda(\mathcal{L}_{mask} + \mathcal{L}_{denoise}), \quad (14)$$

where λ is a weight hyper-parameter.

Experiments

Experimental Setup

Dataset. We evaluate the proposed CM-Pose for video-based human pose estimation in three widely used datasets: PoseTrack2017 (Iqbal, Milan, and Gall 2017), PoseTrack2018 (Andriluka et al. 2018), and PoseTrack2021 (Doring et al. 2022). **PosTrack2017** includes 80,144 pose annotations and has two subsets, *i.e.*, training (train) and validation (val) with 250 videos and 50 videos (split according to the official protocol), respectively. **PoseTrack2018** largely increases the number of video clips and pose annotations including 593 videos for training, 170 videos for validation, and the total number of pose annotations is 153,615. PoseTrack2018 also introduces an additional flag characterizing joint visibility. **PoseTrack2021** further increases the number of pose annotations for small or crowded persons, including 177,164 labels. All three datasets identify 15 keypoints and the training set is densely labeled in the center 30 frames, while the validation set contains additional pose annotations every 4 frames.

Implementation Details. We implement our method CM-Pose for human pose estimation with Pytorch, which is trained on 2 Nvidia Geforce RTX 4090 GPUs and terminated with 20 epochs. We use Vision Transformer (ViT) (Dosovitskiy et al. 2020), pre-trained on the COCO dataset, as the backbone network. For data augmentation, we adopt random scale with a factor of ± 0.35 , random rotation $[-45^\circ, 45^\circ]$, truncation, and flipping. We set the image size as 256×192 . The time span ω is set to 1. The number of keypoint tokens K is 15. We use AdamW optimizer to train the model with an initial learning rate of $2e - 4$ (decays to $2e - 5, 2e - 6, 2e - 7$ at the 5-th, 12-th, 18-th epochs, respectively).

Evaluation Metric. As in previous works (Liu et al. 2022a), we use average precision (AP) to evaluate the performance of our model. We calculate the AP for each keypoint and average them to get the final performance (mAP).

Comparison with State-of-the-art Methods

Results on the PoseTrack2017 Dataset. We evaluate our method CM-Pose with existing 11 pose estimation methods, and the results are tabulated in Table 1. Compared to previous methods, the proposed CM-Pose achieves a new state-of-the-art performance of 87.5 mAP and delivers a gain of 1.8 mAP over the best-performing previous work TDMI (Feng et al. 2023). Specifically, we also observe that the encouraging improvement for challenging joints (*i.e.*, wrist, hip): with an mAP of 85.6 ($\uparrow 3.0$) for wrists and an mAP of 88.6 ($\uparrow 3.4$) for hips. As would be expected, the proposed auxiliary task learning (*i.e.*, masked token reconstruction task, denoising token task) helps the network to infer challenging keypoints from known keypoints based on spatio-temporal dependencies and causal reasoning ability, which is especially important for challenging scenes such as pose occlusion and video defocus.

Results on the PoseTrack2018 Dataset. We further evaluate our model on the PoseTrack2018 dataset. Table 2 reports the empirical comparisons on validation set. We can

Method	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Mean
HRNet (Sun et al. 2019)	82.1	83.6	80.4	73.3	75.5	75.3	68.5	77.3
CorrTrack (Rafi, Leibe, and Gall 2020)	86.1	87.0	83.4	76.4	77.3	79.2	73.3	80.8
Dynamic-GNN (Yang et al. 2021)	88.4	88.4	82.0	74.5	79.1	78.3	73.1	81.1
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. 2021)	88.0	88.7	84.1	78.4	83.0	81.4	74.2	82.8
SLT-Pose (Gai et al. 2023)	88.9	89.7	85.6	79.5	84.2	83.1	75.8	84.2
HANet (Jin et al. 2023)	90.0	90.0	85.0	78.8	83.1	82.1	77.1	84.2
M-HANet (Jin et al. 2024)	90.3	90.7	85.3	79.2	83.4	82.6	77.8	84.8
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. 2023)	90.0	91.1	87.1	81.4	85.2	84.5	78.5	85.7
CM-Pose (Ours)	89.2	92.0	89.0	85.6	88.6	87.2	81.1	87.5

Table 1: Quantitative results on the **PoseTrack2017** dataset.

Method	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Mean
Dynamic-GNN (Yang et al. 2021)	80.6	84.5	80.6	74.4	75.0	76.7	71.8	77.9
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. 2021)	84.0	86.6	82.7	78.0	80.4	79.3	73.8	80.9
SLT-Pose (Gai et al. 2023)	84.3	87.5	83.5	78.5	80.9	80.2	74.4	81.5
FAMI-Pose (Liu et al. 2022a)	85.5	87.7	84.2	79.2	81.4	81.1	74.9	82.2
HANet (Jin et al. 2023)	86.1	88.5	84.1	78.7	79.0	80.3	77.4	82.3
M-HANet (Jin et al. 2024)	86.7	88.9	84.6	79.2	79.7	81.3	78.7	82.7
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. 2023)	86.2	88.7	85.4	80.6	82.4	82.1	77.5	83.5
CM-Pose (Ours)	85.7	88.9	85.8	81.0	84.4	84.2	80.1	84.4

Table 2: Quantitative results on the **PoseTrack2018** dataset.

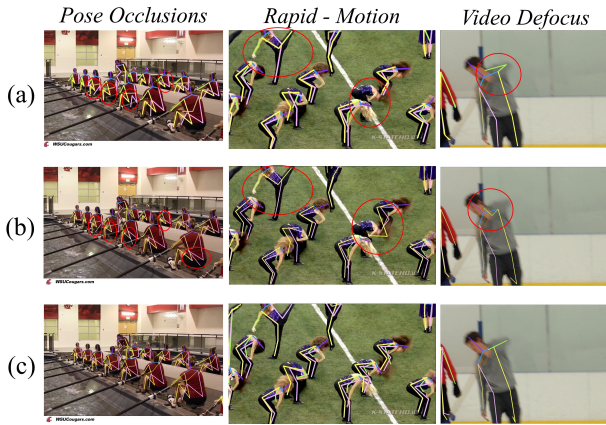


Figure 2: The keyframe (a) and visual comparisons of results obtained from FAMI (a), TDMI (b), and our CM-Pose (c) on challenging scenes in the PoseTrack dataset. Inaccurate predictions are highlighted with the red solid circles.

observe that our method one again surpasses the state-of-the-art methods, reaching an mAP of 84.4, with an mAP of 84.4, 84.2, 80.1 for the hip, knee, and ankle joints.

Results on the PoseTrack2021 Dataset. We present our results of PoseTrack2021 dataset in Table 3, comparing our model with previous state-of-the-art methods. We see that existing methods such as TDMI (Feng et al. 2023) and DSTA (He and Yang 2024) have already reached an impressive performance of 83.5 mAP. In contrast, the proposed method CM-Pose achieves a 84.3 mAP. We also obtain an 88.9 (\uparrow 3.1) mAP for the head, 83.7 (\uparrow 1.3) mAP for the knee, and 84.6 (\uparrow 1.1) mAP for hip joints.

Comparison of Visual Results. We further visualize the results with challenging scenarios such as pose occlusion and video defocus to examine the robustness of our method. We illustrate in Figure 2 the side-by-side comparisons of state-of-the-art approaches a) FAMI (Liu et al. 2022a), b) TDMI (Feng et al. 2023) and c) our CM-Pose. From the Figure 2, we see that our CM-Pose consistently estimates more robust and accurate human pose for challenging scenes. FAMI and TDMI design feature alignment or difference to model spatio-temporal dependencies between joints. Through the integration of auxiliary tasks, our CM-Pose capture more comprehensive spatio-temporal dependencies and grasp better causal reasoning capability. On the other hand, by prioritizing causal tokens and integrating potential information in non-causal tokens, CM-Pose have a better interpretability for reliable pose estimation.

Ablation Study

We perform ablative analysis to examine the contribution of each component in the proposed CM-Pose, including Multi-task Coarse Learning (MCL) and Fine Token Enhancement (FTE). We also investigate the effect of different ratio of corrupted tokens (*i.e.*, p_{tm} and p_{tn}). These experiments are performed on the PoseTrack2017 validation dataset.

Study on Components of MCL. We explore the effectiveness of the two auxiliary tasks in the MCL: masked token reconstruction task (“Mask”) and denoising token task (“De-noise”), and show the result in Table 4. “Primary” represents the primary task. From the second, third, and fourth lines of Table 4, we can see that introducing the masked token reconstruction task and denoising token task alone could improve the performance of pose estimation, proving the effectiveness of these auxiliary tasks. Additionally, as shown in the

Method	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Mean
CorrTrack (Rafi, Leibe, and Gall 2020)	-	-	-	-	-	-	-	72.3
CorrTrack* (Rafi, Leibe, and Gall 2020)	-	-	-	-	-	-	-	72.7
DCPose (Liu et al. 2021)	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
DSTA (He and Yang 2024)	87.5	87.0	84.2	81.4	82.3	82.5	77.7	83.5
TDMI (Feng et al. 2023)	85.8	87.5	85.1	81.2	83.5	82.4	77.9	83.5
CM-Pose (Ours)	88.9	88.3	84.4	81.9	84.6	83.7	78.8	84.3

Table 3: Quantitative results on the **Posetrack2021** dataset. CorrTrack* denotes CorrTrack without ReID.

Primary	Mask	Denoise	Mean	Declines
✓	✓	✓	87.5	-
✓	✓	-	86.9	0.6 (↓)
✓	-	✓	87.0	0.5 (↓)
✓	-	-	85.6	1.9 (↓)

Table 4: Ablation study on auxiliary tasks in **MCL**.

Ablation	$\hat{\mathcal{F}}_v^c$	$\bar{\mathcal{F}}_v^{ir}$	Mean	Declines
CM-Pose	✓	✓	87.5	-
(a)	✓	-	85.4	2.1 (↓)
(b)	-	-	86.0	1.5 (↓)

Table 5: Ablation of different designs in **FTE**.

second and fifth line of Table 4, we also find that jointly adopting two auxiliary tasks can further improve the pose estimation performance.

Study on the Masked and Noise Ratios. We perform experiments to explore the effectiveness of different masked ratio p_{tm} and noise ratio p_{tn} . As shown in Fig. 3, we can observe that if the masked ratio p_{tm} and noise ratio p_{tn} are set too low or too high, both lead to poor performance. We conjecture the possible reason is that the inappropriate ratios make the auxiliary task unsuitable for model learning. Furthermore, a moderate ratio makes the model work best, where the appropriate range of masked ratio p_{tm} is 0.4 to 0.5 and noise ratio p_{tn} is 0.3 to 0.6.

Study on Components of FTE. We further verify the impact of the fine token enhancement under different setting and tabulated in Table 5. (a) For the first experiment setting, we discard the non-causal tokens $\bar{\mathcal{F}}_v^{nc}$ and only utilize causal tokens $\hat{\mathcal{F}}_v^c$ to estimate the human pose. The mAP performance drops from 87.5 to 85.4. This results deterioration on top of the non-causal tokens also include potentially useful semantic information for pose estimation. (b) For the next experiment setting, we remove the FTE module from CM-Pose. It should be noted that after removing FTE, the tokens from temporal fusion layer are fed to detection head. We see that the performance drops from 87.5 to 86.0 mAP. This indicates the importance of our FTE module, which could pay more attention to causal tokens while also understanding the potential contribution of non-causal tokens to the pose estimation. Interestingly, the setting (b) works better than (a), which also demonstrates the necessity of retaining non-causal tokens.

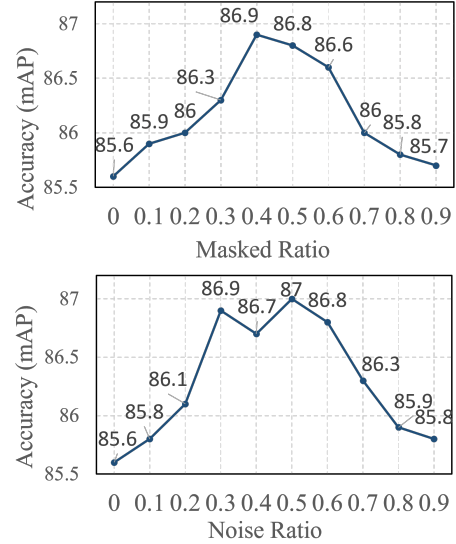


Figure 3: Impact of different masked ratios and noise ratios in training process on PoseTrack2017 validation set.

Conclusion

In this work, we explore the video-based human pose estimation task from a causal perspective. We propose a novel causal-inspired multitask learning framework, termed CM-Pose, which to the best of our knowledge is the first to leverage auxiliary tasks rather than elaborate network design to grasp more causal spatio-temporal modeling ability. To further identify the causal tokens and integrate non-causal tokens, we present a Token Causal Importance Selection module and a Non-Causal Tokens Clustering module. These modules enhance the interpretability of our model and leads to a more reliable pose estimation results. Empirical experiments on three datasets show our method delivers state-of-the-art performance and equips with higher robustness on challenging scenes. In the future, we will try different auxiliary tasks for human pose estimation and extend our effort to other tasks (Yao et al. 2024; Zhang et al. 2024).

Acknowledgments

This research is supported by the National Natural Science Foundation of China (No. 62276112, No. 62372402), Key Projects of Science and Technology Development Plan of Jilin Province (No. 20230201088GX), the Key R&D Program of Zhejiang Province (No. 2023C01217), and Graduate Innovation Fund of Jilin University (No. 2024CX089).

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. In *Advances in Neural Information Processing Systems*, 3027–3038.
- Chen, Y.; Huang, W.; Zhou, S.; Chen, Q.; and Xiong, Z. 2023. Self-supervised neuron segmentation with multi-agent reinforcement learning. In *IJCAI*, 609–617.
- Chen, Y.; Shi, H.; Liu, X.; Shi, T.; Zhang, R.; Liu, D.; Xiong, Z.; and Wu, F. 2024a. TokenUnify: Scalable Autoregressive Visual Pre-training with Mixture Token Prediction. *arXiv preprint arXiv:2405.16847*.
- Chen, Y.; Zheng, S.; Jin, M.; Chang, Y.; and Wang, N. 2024b. DualFluidNet: An attention-based dual-pipeline network for fluid simulation. *Neural Networks*, 177: 106401.
- Chen, Y.; Zheng, S.; Wang, N.; Jin, M.; and Chang, Y. 2024c. A Pioneering Neural Network Method for Efficient and Robust Fuel Sloshing Simulation in Aircraft. *arXiv:2412.10748*.
- 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*.
- Feng, R.; Gao, Y.; Ma, X.; Tse, T. H. E.; and Chang, H. J. 2023. 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.
- Gai, D.; Feng, R.; Min, W.; Yang, X.; Su, P.; Wang, Q.; and Han, Q. 2023. Spatiotemporal Learning Transformer for Video-Based Human Pose Estimation. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Gao, L.; Lei, Y.; Zeng, P.; Song, J.; Wang, M.; and Shen, H. T. 2021. Hierarchical representation network with auxiliary tasks for video captioning and video question answering. *IEEE Transactions on Image Processing*, 31: 202–215.
- Guan, R.; Li, Z.; Tu, W.; Wang, J.; Liu, Y.; Li, X.; Tang, C.; and Feng, R. 2024a. Contrastive Multiview Subspace Clustering of Hyperspectral Images Based on Graph Convolutional Networks. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–14.
- Guan, R.; Tu, W.; Li, Z.; Yu, H.; Hu, D.; Chen, Y.; Tang, C.; Yuan, Q.; and Liu, X. 2024b. Spatial-Spectral Graph Contrastive Clustering with Hard Sample Mining for Hyperspectral Images. *IEEE Transactions on Geoscience and Remote Sensing*, 1–16.
- He, J.; and Yang, W. 2024. Video-Based Human Pose Regression via Decoupled Space-Time Aggregation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1022–1031.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16000–16009.
- Hu, D.; Dong, Z.; Liang, K.; Yu, H.; Wang, S.; and Liu, X. 2024a. High-order Topology for Deep Single-cell Multi-view Fuzzy Clustering. *IEEE Transactions on Fuzzy Systems*.
- Hu, D.; Liu, S.; Wang, J.; Zhang, J.; Wang, S.; Hu, X.; Zhu, X.; Tang, C.; and Liu, X. 2024b. Reliable Attribute-missing Multi-view Clustering with Instance-level and feature-level Cooperative Imputation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 1456–1466.
- 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.
- Jiao, Y.; Chen, H.; Feng, R.; Chen, H.; Wu, S.; Yin, Y.; and Liu, Z. 2022. GLPose: Global-local representation learning for human pose estimation. *ACM transactions on multimedia computing, communications, and applications (TOMM)*, 18(2s): 1–16.
- 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.
- Jin, K.-M.; Lee, G.-H.; Nam, W.-J.; Kang, T.-K.; Kim, H.-W.; and Lee, S.-W. 2024. Masked Kinematic Continuity-aware Hierarchical Attention Network for pose estimation in videos. *Neural Networks*, 169: 282–292.
- Jin, K.-M.; Lim, B.-S.; Lee, G.-H.; Kang, T.-K.; and Lee, S.-W. 2023. Kinematic-aware hierarchical attention network for human pose estimation in videos. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 5725–5734.
- Li, Y.; Zhang, S.; Wang, Z.; Yang, S.; Yang, W.; Xia, S.-T.; and Zhou, E. 2021. Tokenpose: Learning keypoint tokens for human pose estimation. In *Proceedings of the IEEE/CVF International conference on computer vision*, 11313–11322.
- Liu, Y.; Qin, G.; Chen, H.; Cheng, Z.; and Yang, X. 2024. Causality-Inspired Invariant Representation Learning for Text-Based Person Retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 14052–14060.
- Liu, Z.; Chen, H.; Feng, R.; Wu, S.; Ji, S.; Yang, B.; and Wang, X. 2021. 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.; Wu, S.; Xu, C.; Wang, X.; Zhu, L.; Wu, S.; and Feng, F. 2022b. Copy motion from one to another: Fake motion video generation. *arXiv preprint arXiv:2205.01373*.
- Long, S.; Zhao, Z.; Pi, J.; Wang, S.; and Wang, J. 2023. Beyond attentive tokens: Incorporating token importance and diversity for efficient vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10334–10343.
- Qiao, S.; Chen, L.-C.; and Yuille, A. 2021. Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10213–10224.
- Rafi, U.; Leibe, B.; and Gall, J. 2020. Self-supervised key-point correspondences for multi-person pose estimation and tracking in videos. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, 36–52. Springer.
- 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.
- Song, J.; Wang, L.; Van Gool, L.; and Hilliges, O. 2017. Thin-slicing network: A deep structured model for pose estimation in videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4220–4229.
- 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 conference on computer vision and pattern recognition*, 5693–5703.
- Tang, F.; Xu, Z.; Huang, Q.; Wang, J.; Hou, X.; Su, J.; and Liu, J. 2023. DuAT: Dual-aggregation transformer network for medical image segmentation. In *PRCV*.
- Tang, F.; Xu, Z.; Qu, Z.; Feng, W.; Jiang, X.; and Ge, Z. 2024. Hunting Attributes: Context Prototype-Aware Learning for Weakly Supervised Semantic Segmentation. In *CVPR*.
- 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.; 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, K.; Engstrom, L.; Ilyas, A.; and Madry, A. 2020. Noise or signal: The role of image backgrounds in object recognition. *arXiv preprint arXiv:2006.09994*.
- Xu, C.; Tan, R. T.; Tan, Y.; Chen, S.; Wang, X.; and Wang, Y. 2023. Auxiliary tasks benefit 3d skeleton-based human motion prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9509–9520.
- 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.
- Xu, Z.; Tang, F.; Chen, Z.; Zhou, Z.; Wu, W.; Yang, Y.; Liang, Y.; Jiang, J.; Cai, X.; and Su, J. 2024. Polyp-Mamba: Polyp Segmentation with Visual Mamba. In *MIC-CAI*. Springer.
- 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*.
- Yang, Y.; Ren, Z.; Li, H.; Zhou, C.; Wang, X.; and Hua, G. 2021. Learning dynamics via graph neural networks for human pose estimation and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8074–8084.
- Yao, J.; Lai, Y.; Kou, H.; Wu, T.; and Liu, R. 2024. QE-BEV: Query evolution for bird’s eye view object detection in varied contexts. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 2927–2935.
- Yao, J.; Li, C.; and Xiao, C. 2024. Swift Sampler: Efficient Learning of Sampler by 10 Parameters. *arXiv preprint arXiv:2410.05578*.
- Zeng, W.; Jin, S.; Liu, W.; Qian, C.; Luo, P.; Ouyang, W.; and Wang, X. 2022. Not all tokens are equal: Human-centric visual analysis via token clustering transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11101–11111.
- Zhang, Z.; Ji, Y.; and Liu, C. 2023. Knowledge-aware causal inference network for visual dialog. In *Proceedings of the 2023 ACM International Conference on Multimedia Retrieval*, 253–261.
- Zhang, Z.; Zhang, W.; Li, Y.; and Bai, T. 2024. Caption-Aware Multimodal Relation Extraction with Mutual Information Maximization. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 1148–1157.