

DeciWatch: A Simple Baseline for 10× Efficient 2D and 3D Pose Estimation

Ailing Zeng¹, Xuan Ju¹, Lei Yang², Ruiyuan Gao¹, Xizhou Zhu², Bo Dai³, and Qiang Xu¹

¹The Chinese University of Hong Kong, ²Sensetime Group Limited,
³Shanghai AI Laboratory
{alzeng, qxu}@cse.cuhk.edu.hk

Abstract. This paper proposes a simple baseline framework for video-based 2D/3D human pose estimation that can achieve 10× efficiency improvement over existing works without any performance degradation, named *DeciWatch*. Unlike current solutions that estimate each frame in a video, *DeciWatch* introduces a simple yet effective *sample-denoise-recover* framework that only watches sparsely sampled frames, taking advantage of the continuity of human motions and the lightweight pose representation. Specifically, *DeciWatch* uniformly samples less than 10% video frames for detailed estimation, denoises the estimated 2D/3D poses with an efficient Transformer architecture, and then accurately recovers the rest of the frames using another Transformer-based network. Comprehensive experimental results on three video-based human pose estimation, body mesh recovery tasks and efficient labeling in videos with four datasets validate the efficiency and effectiveness of *DeciWatch*. Code is available at <https://github.com/cure-lab/DeciWatch>.

Keywords: Human Pose Estimation, Video Analysis, Efficiency

1 Introduction

2D/3D human pose estimation [36,7,63] has numerous applications, such as surveillance, virtual reality, and autonomous driving. Various high-performance image-based pose estimators [40,51,48,25,30,29] are proposed in the literature, but they are associated with substantial computational costs.

There are two main approaches to improving the efficiency of human pose estimators so that they can be deployed on resource-scarce edge devices (e.g., smart cameras). A straightforward way to improve the efficiency is designing more compact models, such as numerous light-weighted image-level pose estimators [3,35,21,61,6,60,42,62,4,18,54] (see Fig. 1(a)(i)) and video-level pose estimators [41,9] (see Fig. 1(a)(ii)) introduced in previous literature. However, when estimating on a video, such approaches inevitably lead to a sub-optimal solution for efficiency improvement due to the frame-by-frame estimation scheme. In contrast, a promising but rarely explored direction is leveraging the semantic redundancy among frames of videos, where we can feed only keyframes to heavy

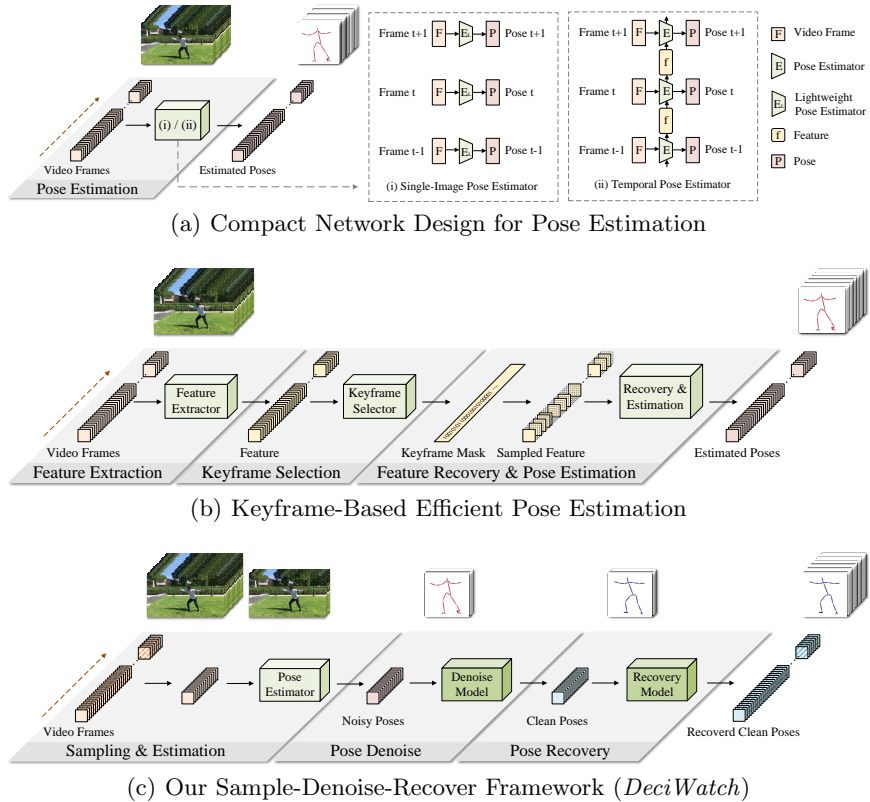


Fig. 1. The workflows of three types of efficient pose estimation frameworks. (a) is compact model designs. The (green) pose estimation module has two design strategies: (i) shows single-frame efficient methods [3,35,21,61,6,60,42,62,4,18] that use lightweight models to reduce the costs of each frame; (ii) presents some temporal efficient strategies [9,41] that utilize feature similarities among consecutive frames via RNNs to decrease feature extraction cost. (b) is the keyframe-based efficient framework[59,10]. They first select about 30%~40% keyframes in a video by watching all frames, then recover the whole sequence based on features of selected keyframes. (c) is the proposed efficient *sample-denoise-recover* framework *DeciWatch* with 5%~10% frames watched.

and high-performance modules and recover or estimate the rest of the frames with light-weighted modules [59,10] (see Fig. 1(b)). While the computational efficiency of these works is improved due to the use of keyframes, they still need to conduct costly feature extraction on each frame for keyframe selection, making it hard to further reduce their computational complexity.

To achieve highly efficient 2D/3D pose estimation *without the need of watching every frame in a video*, we propose a novel framework based on the continuity of human motions, which conducts pose estimation only on sparsely sampled video frames. Since these detected poses ineluctably contain various noises, they

will affect the effectiveness of the recovery. Subsequently, poses of those sampled frames should be denoised before recovered, where we formulate the three-step *sample-denoise-recover* framework. By doing so, the problem in the *recover* stage is similar to the long-standing motion completion task in the computer graphics literature [31,24,46,28,2,8]. However, there are two main differences: (i). our objective is to achieve highly efficient pose estimation, and hence we could only afford lightweight models for pose recovery on frames that are not processed by pose estimators; (ii). most existing motion completion works assume ground-truth poses on the given keyframes. In contrast, the visible frames in our task could have untrustworthy poses with challenging occlusion or rarely seen actions.

This work proposes a simple yet effective baseline framework (see Fig. 1(c)) that watches sparsely sampled frames for highly efficient 2D and 3D video-based human pose estimation. We empirically show that we could maintain and even improve the pose estimation accuracy, with less than 10% frames calculated with the costly pose estimator. We name the proposed framework *DeciWatch*, and the contributions of this work include:

- To the best of our knowledge, this is the first work that considers sparsely sampled frames in video-based pose estimation tasks. *DeciWatch* is compatible with any given single-frame pose estimator, achieving $10\times$ efficiency improvement without any performance degradation. Moreover, the pose sequence obtained by *DeciWatch* is much smoother than existing solutions as it naturally models the continuity of human motions.
- We propose a novel sample-denoise-recover pipeline in *DeciWatch*. Specifically, we uniformly sample less than 10% of video frames for estimation, denoise the estimated 2D/3D poses with an efficient Transformer architecture named *DenoiseNet*, and then accurately recover the poses for the rest of the frames using another Transformer network called *RecoverNet*. Thanks to the lightweight pose representation, the two subnets in our design are much smaller than the costly pose estimator.
- We verify the efficiency and effectiveness of *DeciWatch* on three human pose estimation, body recovery tasks, and efficient labeling in videos with four widely-used datasets and five popular single-frame pose estimators as backbones. We also conduct extensive ablation studies and point out future research directions to further enhance video-based tasks’ efficiency.

2 Related Work

2.1 Efficient Human Pose Estimation

Efficient attempts at human pose estimation can be divided into image-based and video-based. Image-based efficient pose estimators [3,35,21,61,6,60,42,62,4,54] mainly focus on employing well-designed network structures [3,60,42,62,4,54,56], knowledge distillation [35,21,18], or low-resolution features [61,6,33] to reduce model capacity and decrease spatial redundancies, where they may suffer from accuracy reduction, especially in the cases of complex and rare poses. Moreover,

when dealing with videos, these methods reveal their limitations for having to estimate poses frame-by-frame. Their outputs also suffer from unavoidable jitters because they lack the capability of using temporal information.

To cope with video inputs, other attempts exploit temporal co-dependency among consecutive frames to decrease unnecessary calculations. However, only a few video-based efficient estimation methods [9,41,59,10] are proposed in the literature, and they mainly target on 2D pose estimation. In particular, DKD [41] introduces a lightweight distillator to online distill the pose knowledge via leveraging temporal cues from the previous frame. In addition to using local information of adjacent frames, KFP [59] designs a keyframe proposal network that selects informative keyframes after estimating the whole sequence, and then applies a learned dictionary to recover the entire pose sequence. Lastly, MAPN [10] exploits the readily available motion and residual information stored in the compressed streams to dramatically boost the efficiency, and all the residual frames will be calculated by a dynamic gate.

These proposed methods reduce computation costs by employing adaptive operations on different frames, *i.e.*, complex operations on indispensable frames and simple ones on the rest. Despite obtaining efficiency improvement, they still fail to push the efficiency to a higher level since they ignore the fact that it is not necessary to watch each frame. What’s more, relying on image features as intermediate representation is heavy for calculation.

2.2 Motion Completion

Motion completion is widely explored in the area of computer graphics, generally including motion capture data completion [20,44,32,12,50,1,31,24,46] and motion in-filling [11,14,15,14,17,28,2,8,55], which has great significance in the film, animation, and game applications. To be specific, points or sequences missing often occur in motion capture due to technical limitations and occlusions of markers. Accordingly, existing approaches include traditional methods (e.g. linear, Cubic Spline, Lagrange, and Newton’s polynomial interpolation, low-rank matrix completion) [20,44,32,12] and learning-based methods (e.g., Recurrent Neural Networks (RNNs)) [31,24]. Motion in-filling aims to complete the absent poses with specific keyframe constraints. RNNs [11,14,15,14,52] and convolutional models [53,28] are commonly used in motion in-filling. Recently, Generative adversarial learning [17,27] and autoencoder [28,2] are also introduced for realistic and naturalistic output. Some recent works[8,19] also introduce self-attention models to infill the invisible frames.

Although both general motion completion and our target are to recover the full pose sequence, there are two main differences. On the one hand, the objective of motion completion is to generate diverse or realistic motions under certain assumptions, e.g., a recurring or repeated motion like walking. They may fail when motions are aperiodic and complex. In contrast, our goal is to achieve high efficiency in video-based pose estimation, where the benchmarks are usually from real-life videos. On the other hand, motion completion assumes having ground-

truth poses as inputs rather than estimated poses. Current designs may not be able to handle unreliable and noisy poses generated from deep models.

3 Method

3.1 Problem Definition and Overview

Given an input video $\mathcal{I} = \{\mathbf{I}^t\}_{t=1}^T$ of length T , a pose estimation framework computes the corresponding sequence of poses $\hat{\mathcal{P}} = \{\hat{\mathbf{P}}^t\}_{t=1}^T$, aiming to minimize the distance between the estimated poses $\hat{\mathcal{P}}$ and the ground-truth poses \mathcal{P} . $\hat{\mathbf{P}}^t$ could be any human pose representation, including 2D keypoint position, 3D keypoint position, and 6D rotation matrix.

The main target of this work is to set a baseline for efficient video-based pose estimation without compromising accuracy. As shown in Fig. 1(c), we devise a three-step *sample-denoise-recover* flow to process video-based pose estimation efficiently and effectively. As adjacent frames usually contain redundant information and human motion is continuous, *DeciWatch* first samples a small percentage of frames (e.g., 10%) $\mathbf{I}^{sampled}$ and applies existing pose estimators [51,39,25,30,29] thereon to obtain the corresponding poses. However, recovering the full pose sequence from sparsely observed poses is challenging, especially when the poses are estimated by networks and often contain noise. Relying on a few poses to recover the entire sequence, the quality of sampled poses is the key. To tackle the challenge, we introduce two subnets, *DenoiseNet* and *RecoverNet*. Specifically, *DenoiseNet* refines sparse poses from pose estimator. Then *RecoverNet* performs motion recovery based on the refined sparse poses to recover the whole pose sequence, with the intuition that humans can perceive complete motion information through a small number of keyframes. With this new mechanism, the computation cost can be reduced significantly by watching only a small number of frames, which replaces high-cost image feature extraction and pose estimation with a low-cost pose recovery.

3.2 Getting Sampled Poses

Different from the previous keyframe-based efficient frameworks [59,10] using each frame’s feature to select keyframes, we use a uniform sampling that watches one frame in every N frame to select sparse frames $\mathbf{I}^{sampled}$ as a baseline strategy. Due to the redundancy in consecutive frames and continuity of human poses, a uniform sampling strategy under a certain ratio is capable of keeping enough information for recovery. Then we can estimate $\mathbf{I}^{sampled}$ by any existing pose estimators, such as SimplePose [51] for 2D poses, FCN [39] for 3D poses, and PARE [29] for 3D body mesh recovery, to get sparse poses $\hat{\mathbf{P}}_{noisy}^{sampled} \in \mathbb{R}^{\frac{T}{N} \times (K \cdot D)}$. K is the number of keypoints, and D is the dimensions for each keypoint. Notably, we experimentally show that uniform sampling can surpass complex keyframe selection methods from both efficiency and accuracy in Sec. 4.2.

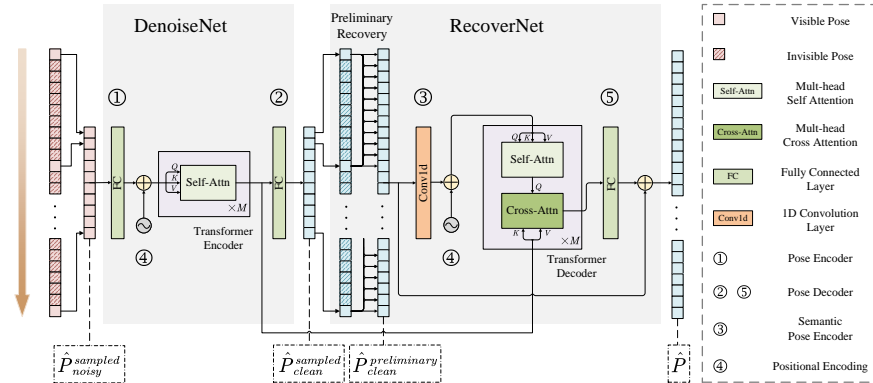


Fig. 2. Illustration of the *denoise and recovery* subnets. First, we denoise the sparsely sampled poses $\hat{\mathbf{P}}_{noisy}^{sampled}$ into a clean poses $\hat{\mathbf{P}}_{clean}^{sampled}$ by a transformer-based *DenoiseNet* to handle the dynamically various noises. Then, after a preliminary pose recovery, we embed the sequence into temporal semantic pose tokens and put them into another transformer-based *RecoverNet* that can leverage spatio-temporal correlations to recover realistic and natural poses.

3.3 Denoising the Sampled Poses

Motion completion often resorts to ground-truth sparse poses for infilling the whole sequence. However, in our scenario, the sampled poses are obtained by single-frame pose estimators, inevitably leading to noisy, sparse poses. Consequently, the quality of sparse poses is crucial for motion recovery. Before recovering the full motion, we develop a denoising network to refine the sampled poses $\hat{\mathbf{P}}_{noisy}^{sampled}$ to clean poses $\hat{\mathbf{P}}_{clean}^{sampled}$. Due to the temporal sparseness and noisy jitters, the key designs of *DenoiseNet* lie in two aspects: (i) A *dynamic* model for handling diverse possible pose noises; (ii) *Global* temporal receptive fields to capture useful Spatio-temporal information while suppressing distracting noises. Based on these two considerations, local operations, like convolutional or recurrent networks, are not well suited. Intuitively, Transformer-based models [49] are capable of capturing the global correlations among discrete tokens, so we use Transformer-based encoder modules to relieve noises from the sparse poses. The denoise process can be formulated as:

$$\hat{\mathbf{F}}_{clean}^{sampled} = \mathbf{TransformerEncoder} \left(\hat{\mathbf{P}}_{noisy}^{sampled} \mathbf{W}_{DE} + \mathbf{E}_{pos} \right) \quad (1)$$

As demonstrated in the left block of Fig. 2, after being encoded through a linear projection matrix $\mathbf{W}_{DE} \in \mathbb{R}^{(K \cdot D) \times C}$, summed with a positional embedding $\mathbf{E}_{pos} \in \mathbb{R}^{\frac{T}{N} \times C}$, and then processed by the **TransformerEncoder** composed of M Multi-head Self-Attention blocks as in [49], input noisy poses $\hat{\mathbf{P}}_{noisy}^{sampled}$ are embedded into a clean feature $\hat{\mathbf{F}}_{clean}^{sampled} \in \mathbb{R}^{\frac{T}{N} \times C}$, where C is the embedding dimension. Dropout, Layer Normalization, and Feedforward layers are the same

as the original Transformer. Lastly, we use another linear projection matrix $\mathbf{W}_{DD} \in \mathbb{R}^{C \times (K \cdot D)}$ to obtain refined sparse poses.

$$\hat{\mathbf{P}}_{clean}^{sampled} = \hat{\mathbf{F}}_{clean}^{sampled} \mathbf{W}_{DD} \quad (2)$$

The learnable parameters in *DenoiseNet* are trained by minimizing $\hat{\mathbf{P}}_{clean}^{sampled}$ with sampled ground-truth poses $\mathbf{P}^{sampled}$.

3.4 Recovering the Sampled Poses

After getting the sparse clean poses $\hat{\mathbf{P}}_{clean}^{sampled} \in \mathbb{R}^{\frac{T}{N} \times (K \cdot D)}$, we use another Spatio-temporal subnet, *RecoverNet*, to recover the absent poses. In order to learn the consistent temporal correlations, a simple temporal upsampling (e.g., a linear transformation $\mathbf{W}_{PR} \in \mathbb{R}^{T \times \frac{T}{N}}$) is applied to perform preliminary sequence recovery to get $\hat{\mathbf{P}}_{clean}^{preliminary} \in \mathbb{R}^{T \times (K \cdot D)}$.

$$\hat{\mathbf{P}}_{clean}^{preliminary} = \mathbf{W}_{PR} \hat{\mathbf{P}}_{clean}^{sampled} \quad (3)$$

To make the recovery more realistic and accurate, we adopt another transformer-based network for detailed poses recovery. Unlike the previous pose transformers [64], we bring temporal semantics into pose encoding to encode the neighboring D frames' poses into pose tokens via a temporal 1D convolutional layer. The main architecture of *RecoverNet* is also the same as Transformer, which employs M multi-head self-attention blocks.

$$\hat{\mathbf{P}} = \mathbf{TransformerDecoder} \left(\mathbf{Conv1d} \left(\hat{\mathbf{P}}_{clean}^{preliminary} \right) + \mathbf{E}_{pos}, \hat{\mathbf{F}}_{clean}^{sampled} \right) \mathbf{W}_{RD}, \quad (4)$$

where the pose decoder is $\mathbf{W}_{RD} \in \mathbb{R}^{C \times (K \cdot D)}$. As illustrated in the second block marked as *RecoverNet* in Fig.2, we draw key information in the Cross-Attention block by leveraging denoised features $\hat{\mathbf{F}}_{clean}^{sampled}$.

Efficiency calculation. The computational costs of *DeciWatch* is from three parts: (i) using existing backbones to estimate the sampled poses $\hat{\mathbf{P}}_{noisy}^{sampled}$, (ii) using *DenoiseNet* to get clean sampled poses $\hat{\mathbf{P}}_{clean}^{sampled}$, and (iii) using *RecoverNet* to recover the clean sampled poses to the complete pose sequence $\hat{\mathbf{P}}^t$. To summarize, FLOPs of *DeciWatch* is:

$$FLOPs = \frac{1}{T} (T/N * f(E) + T * (f(D) + f(R))), \quad (5)$$

where $f(\cdot)$ calculates the model's per frame FLOPs. $f(E)$, $f(D)$ and $f(R)$ represent per frame FLOPs of pose estimators, *DenoiseNet* and *RecoverNet*, respectively. Using poses instead of image features as representation makes two subnets computational efficient. Notably, $(f(D) + f(R)) \ll f(E)$ (more than $10^4 \times$. Details can be find in Table 1 and 2). Since *DeciWatch* samples very few frames in step 1, the mean FLOPs can be reduced to $1/N$ compared with those watch-every-frame methods, resulting in a $10 \times$ speedup overalls.

3.5 Loss Function

We follow recent 3D pose estimation methods [43,57] to apply a simple L1 regression loss to minimize the errors between \mathbf{P}^t and $\hat{\mathbf{P}}^t$ for 2D or 3D pose estimation. Particularly, to learn the noisy patterns from sampled estimated poses, we further add an L1 loss between sparse estimated poses $\hat{\mathbf{P}}_{clean}^{sampled}$ and the corresponding ground-truth poses $\mathbf{P}^{sampled}$. Therefore, the objective function is defined as follows.

$$\mathcal{L} = \lambda \left(\frac{1}{T} \sum_{t=1}^T |\hat{\mathbf{P}}^t - \mathbf{P}^t| \right) + \frac{1}{(T/N)} \sum_{n=1}^{T/N} |\hat{\mathbf{P}}_{clean}^{sampled(n)} - \mathbf{P}^{sampled(n)}|, \quad (6)$$

where λ is a scalar to balance the losses between *RecoverNet* and *DenoiseNet*. We set $\lambda = 5$ by default.

4 Experiments¹

4.1 Experimental Settings

Datasets. We verify our baseline framework on three tasks. For 2D pose estimation, we follow existing video-based efficient methods [59,10] using dataset Sub-JHMDB [23]. For 3D pose estimation, we choose the most commonly used dataset Human3.6M [22]. For 3D body recovery, we evaluate on an in-the-wild dataset 3DPW [38] and a dance dataset AIST++ [34] with fast-moving and diverse actions.

Evaluation metrics. For 2D pose estimation, we follow previous works [41,59,10] adopting the Percentage of the Correct Keypoints (*PCK*), where the matching threshold is set as 20% of the bounding box size under pixel level. For 3D pose estimation and body recovery, following [26,30,29,39,58], we report Mean Per Joint Position Error (*MPJPE*) and the mean Acceleration error (*Accel*) to respectively measure the localization precision and smoothness. Besides, we report efficiency metrics mean FLOPs (G) per frame, the number of parameters and the inference time tested on a single TITAN Xp GPU.

Implementation details. To facilitate the training and testing in steps 2 and 3, we first prepare the detected poses on both training and test sets offline. The uniform sampling ratio is set to 10% by default, which means watching one frame in every $N = 10$ frames in videos. To deal with different input video lengths, we input non-overlapping sliced windows with fixed window sizes. It is important to make sure the first and last frames are visible, so the input and output window sizes are both $(N * Q + 1)$, where Q is the average number of visible frames in a window. We set $Q = 1$ for 2D poses due to the short video length of the 2D dataset and $Q = 10$ for others. We change embedding dimension C and video length T to adapt different datasets and estimators, which influence FLOPs

¹ Due to the limit of pages, we present *data description, comprehensive results of different sampling ratios, the effect of hyper-parameters, generalization ability, qualitative results, and failure cases analyses* in the supplementary material.

slightly. For *DenoiseNet*, we apply $M = 5$ transformer blocks with embedding dimension $C = 128$ by default. For *RecoverNet*, we use the same settings as *DenoiseNet*. The temporal kernel size of the semantic pose encoder is 5. For more details, please refer to the supplementary material. All experiments can be conducted on a single TITAN Xp GPU.

4.2 Comparison with Efficient Video-based Methods

Existing efficient video-based pose estimation methods [41,59,10] only validate on 2D poses. In this section, we compare the accuracy and the efficiency of *DeciWatch* with SOTAs. We follow their experiment settings for fair comparisons and use the same pose estimator SimplePose [51].

Table 1. Comparison on Sub-JHMDB [23] dataset with existing video-based efficient methods [41,59,10] for 2D pose estimation. R stands for ResNet backbone [16]. $Ratio$ represents the sampling ratio. The pose estimator of *DeciWatch* is the single-frame model SimplePose (R50) [51]. Best results are in bold.

Sub-JHMDB dataset - 2D Pose Estimation										
Methods	Head	Sho.	Elb.	Wri.	Hip	Knee	Ank.	Avg. \uparrow	FLOPs(G) \downarrow	$Ratio$
Luo <i>et al.</i> [37]	98.2	96.5	89.6	86.0	98.7	95.6	90.0	93.6	70.98	100%
DKD (R50) [41]	98.3	96.6	90.4	87.1	99.1	96.0	92.9	94.0	8.65	100%
KFP (R50) [59]	95.1	96.4	95.3	91.3	96.3	95.6	92.6	94.7	10.69 [†]	44.5%
KFP (R18) [59]	94.7	96.3	95.2	90.2	96.4	95.5	93.2	94.5	7.19 [†]	40.8%
MAPN (R18) [10]	98.2	97.4	91.7	85.2	99.2	96.7	92.2	94.7	2.70	35.2%
SimplePose [51]	97.5	97.8	91.1	86.0	99.6	96.8	92.6	94.4	11.96	100%
DeciWatch	99.8	99.5	99.7	99.7	98.7	99.4	96.5	98.8	1.196+0.0005 [‡]	10.0%
DeciWatch	99.9	99.8	99.6	99.7	98.6	99.6	96.6	98.9	0.997+0.0005 [‡]	8.3%
DeciWatch	98.4	98.3	98.2	98.7	97.5	98.3	95.2	97.5	0.598+0.0005[‡]	5.0%

[†] The results are recalculated according to $Ratio$ and their tested FLOPs for SimplePose (i.e., 11.96G).

[‡] Tested with ptflops v0.5.2 [47].

As shown in Table 1, our approach shows significantly increased accuracy with the highest efficiency, achieving more than $20\times$ improvement in the computation cost on the Sub-JHMDB dataset. Compared to the SOTA method [10], we surpass them by 4.3% and 4.4% on average PCK (Avg.) with 55.7% and 77.9% reduction in FLOPs. Our improvement mainly comes from elbows (Elb.) (from 91.7% to 99.6%) and ankles (Ank.) (from 92.2% to 96.6%) under a 8.3% ratio. These outer joints usually move faster than inner joints (e.g., Hips), which may cause motion blur and make estimators hard to detect precisely. However, previous efficient video-based pose estimation methods did not consider a denoising or refinement strategy. *DeciWatch* uses *DenoiseNet* to reduce noises. Then, *RecoverNet* interpolates the sparse poses using the assumption of continuity of motion without watching blurry frames, showing the superiority of *DeciWatch*.

To further verify the effectiveness of the denoise scheme in *DeciWatch*, we input the full sequence of outputs from SimplePose, which means the *Ratio* is 100%, and the result of *PCK* is 99.3. The additional improvement in accuracy shows that *DeciWatch* can also be used as an effective denoise/refinement model to further calibrate the output positions. Based on the above observations, using a lightweight *DeciWatch* in a regression manner to further refine heatmap-based 2D pose estimation methods can be a promising refinement strategy. For efficiency, the total number of parameters in *DenoiseNet* and *RecoverNet* is 0.60M and the inference time is about 0.58ms/frame.

Besides, we argue that PCK@0.05 with lower thresholds will be better to reflect the effectiveness of the methods since the commonly used metric PCK@0.2 appears saturated. We report PCK@0.05 and PCK@0.1 of *DeciWatch* in Supp.

4.3 Boosting Single-frame Methods

The used single-frame pose estimators: We compare *DeciWatch* with the following single-frame pose estimators [39,30,25,29] that watch each frame when estimating a video. We first introduce these methods as follows.

- FCN [39] is one of the most important 2D-to-3D methods with multiple fully connected layers along the spatial dimension.
- SPIN [30] is one of the most commonly used methods, which combines SMPL optimization in the training process.
- EFT [25] is trained on augmented data compared with SPIN [30] to get better performance and generalization ability.
- PARE [29] proposes a part-guided attention mechanism to handle partial occlusion scenes, achieving the state-of-the-art on many benchmarks.

The comparisons: We demonstrate the comparison results in Table 2 at sampling ratios of 10% ($N = 10$) and 6.7% ($N = 15$). To be specific, when the sampling ratio is 10%, *DeciWatch* can reduce *MPJPE* by about 2% to 3% for most estimators and reduce *Accel* by about 73% to 92%, indicating *DeciWatch* achieves higher precision and smoothness with about 10% FLOPs. Moreover, with 6.7% watched frames, *DeciWatch* still has the capability to recover the complete pose sequence with competitive results. For the AIST++ dataset, we surprisingly find that training on sparse poses and recovering them can significantly improve output qualities by 33.8% and 23.6% with a sampling ratio of 10% and 6.7% respectively. This indicates that our method is capable of datasets with fast movements and difficult actions, such as Hip-hop or Ballet dances.

In general, we attribute the high efficiency of *DeciWatch* to the use of lightweight and temporal continuous poses representation rather than the heavy features used by previous works [41,59,10]. Meanwhile, the superior effectiveness, especially for motion smoothness, comes from its ability to capture spatio-temporal dynamic relations in the denoising and recovery process and the well-designed sample-denoise-recover steps. Additionally, the inference speeds in step 2 and 3 are about 0.1ms/frame, significantly faster than image feature extraction.

Table 2. Comparing *DeciWatch* with existing single-image 3D pose estimators on Human3.6M [22], 3DPW [38], and AIST++ [34] datasets. Pose estimators used in *DeciWatch* keep the same as the corresponding methods.

Methods	<i>MPJPE</i> ↓	<i>Accel</i> ↓	FLOPs(G)↓	<i>Ratio</i>
Human3.6M [22] - 3D Pose Estimation				
FCN [39]	54.6	19.2	6.21	100.0%
DeciWatch	52.8 ↓1.8(3.3%)	1.5 ↓17.7(92.2%)	0.621+0.0007	10.0%
DeciWatch	53.5 ↓1.1(2.0%)	1.4 ↓17.8(92.7%)	0.414+0.0007	6.7%
3DPW [38] - 3D Body Recovery				
SPIN [30]	96.9	34.7	4.13	100.0%
DeciWatch	93.3 ↓3.6(3.7%)	7.1 ↓27.6(79.5%)	0.413+0.0004	10.0%
DeciWatch	96.7 ↓0.2(0.2%)	6.9 ↓27.8(80.1%)	0.275+0.0004	6.7%
EFT [25]	90.3	32.8	4.13	100.0%
DeciWatch	89.0 ↓1.3(1.4%)	6.8 ↓26.0(79.3%)	0.413+0.0004	10.0%
DeciWatch	92.3 ↑2.0(2.2%)	6.6 ↓26.2(79.9%)	0.275+0.0004	6.7%
PARE [29]	78.9	25.7	15.51	100.0%
DeciWatch	77.2 ↓1.7(2.2%)	6.9 ↓18.8(73.2%)	1.551+0.0004	10.0%
DeciWatch	80.7 ↑1.8(2.2%)	6.7 ↓18.6(73.9%)	1.034+0.0004	6.7%
AIST++ [34] - 3D Body Recovery				
SPIN [30]	107.7	33.8	4.13	100.0%
DeciWatch	71.3 ↓36.4(33.8%)	5.7 ↓28.1(83.1%)	0.413+0.0007	10.0%
DeciWatch	82.3 ↓25.4(23.6%)	5.5 ↓28.3(83.7%)	0.275+0.0007	6.7%

All estimation results are re-implemented or tested by us for fair comparisons.

4.4 Comparison with Motion Completion Techniques

The third step of *DeciWatch* is similar to motion completion/interpolation as introduced in Sec. 2.2. To assess existing interpolation methods quantitatively, we compare our model with four traditional methods and one of the latest learning-based interpolation methods [2] based on Conditional Variational Auto-Encoder(CVAE). The original experiments in the CAVE-based model are based on the ground-truth of the Human3.6M dataset [22] (marked as CVAE [2]-R.★). We compare two additional settings on the same dataset: (i). CVAE [2]-R. inputs estimated 3D poses rather than ground-truth 3D poses and uses Random sampling; (ii). CVAE [2]-U. inputs estimated 3D poses and use Uniform sampling, which is the same setting as *DeciWatch*. For a fair comparison, we adjusted the sampling ratio of training and testing to be consistent as 20%, 10%, and 5%.

In Table 3, *DeciWatch* outperforms all methods. Specifically, we find that the results of the CVAE-based model are even twice as bad as the traditional methods at all ratios, especially with estimated poses inputs and uniform sampling. This is because CVAE-based methods try to encode a long sequence of motion into an embedding and then recover them, which is practically difficult to embed well and recover precisely for a specific video. Instead, our method and

Table 3. Comparison of *MPJPE* with existing motion completion methods on Human3.6M dataset [22] for 3D pose estimation. Noted that [2] is originally trained and tested on ground-truth 3D poses (noted by \star) with random sampling (CVAE [2]-R.), we retrain their model with detected 3D poses to keep the same uniform sampling as us (CVAE [2]-U.). We use FCN [39] as the single-frame estimator to generate the sparse detected results, and its *MPJPE* is 54.6mm.

<i>Ratio</i>	Nearest	Linear	Quadratic	Cubic-Spline	CVAE-R. \star	CVAE-R.	CVAE-U.	<i>DeciWatch</i>
20%	54.4	54.4	<u>54.3</u>	54.5	87.4	114.1	119.4	52.8 \downarrow _{1.8(3.2%)}
10%	54.7	<u>54.3</u>	55.2	54.4	99.1	119.2	121.5	52.8 \downarrow _{1.8(3.2%)}
5%	57.6	57.5	<u>57.2</u>	57.3	134.9	140.5	123.1	54.4 \downarrow _{0.2(0.3%)}

All estimation results are re-implemented or tested by us for fair comparisons.

the traditional interpolation strategies directly utilize the continuity of human poses as a priori, making the interpolation process easier.

Owing to the relatively low *MPJPE* and the slow motions in the Human3.6M dataset, *DeciWatch* only have limited improvement over traditional methods. We evaluate on a more challenging dance dataset AIST++ [34] in Table 4. A tremendous lift is revealed to over 30% under a 10% ratio. The improvement of *DeciWatch* is from: *DeciWatch* can learn to minimize errors with data-driven training, especially poses with high errors, while traditional methods have no such prior knowledge to decrease the errors from both visible and invisible poses.

Table 4. Comparison of *MPJPE* with traditional interpolation methods on AIST++ dataset [34]. We use 3D pose estimator SPIN [30] as the single-frame estimator, and its *MPJPE* is 107.7mm.

<i>Ratio</i>	Nearest	Linear	Quadratic	Cubic-Spline	<i>DeciWatch</i>
20%	106.7	<u>104.6</u>	105.8	106.8	67.6 \downarrow _{39.7(37.0%)}
10%	108.3	<u>106.3</u>	108.2	108.9	71.3 \downarrow _{36.0(33.6%)}
5%	123.2	120.7	<u>119.9</u>	121.2	90.8 \downarrow _{16.5(15.4%)}

4.5 Ablation Study

As a baseline framework, we do not emphasize the novelty of network design but provide some possible designs in each step for further research. We have explored the impact of different pose estimators in Table 1 and 2 in previous sections. This section will explore how designs in steps 2 and 3 influence the results. All experiments use the same input window length at 101 and a 10% sampling ratio by default. We keep the same setting in both training and testing. **Impact of sampling ratio and input window size.** Due to space limitation, we discuss this part *in Supp.*. We summarize the key observations as follows :(i). With the increase in sampling ratio, the *MPJPEs* first drop before rising, and they are at the lowest when the sampling ratio is about 20%. *Accels* will decrease

constantly. These observations give us a new perspective that in pose estimation, *it is not essential to watch all frames for achieving a better and smoother performance.* (ii). Besides, the *MPJPE* of *DeciWatch* surpasses the original baseline even when the sampling ratio is about 8%. (iii). Lastly, *DeciWatch* is robust to different window sizes from 11 to 201 frames.

Impact of sampling strategies.

Although we use uniform sampling one frame for every N frame, there are other sampling strategies that can be adopted without watching each frame, such as (i). uniform sampling 2 or 3 frames for every N frames, which contains velocity and acceleration information (named as *U.-2* and *U.-3*); (ii). random sampling (*R.*); (iii). combining uniform sampling with random sampling (*U.-R.*). From Table 5, *U.-2* and *U.-3* get the worse results compared to *U.(Ours)* because intervals between visible frames become longer, and the information in two or three

adjacent frames is too similar to be helpful for the recovery. Moreover, random sampling shows is capable of recovery since a long invisible period may appear, which is hard for model learning. Combining uniform sampling (half of the frames) to avoid long invisible periods can slightly decrease the error in random sampling. In summary, uniform sampling one frame for every N frame (*U.(Ours)*) surpasses all other sampling strategies under the same model.

Impact of denoise and recovery subnets. In Table 6, we comprehensively verify the effectiveness of *DenoiseNet* and *RecoverNet* at 10% sampling ratio on three datasets. When we remove any part of the two subnets, the results deteriorate to various degrees. Removing *RecoverNet* means we only use a preliminary recovery via a temporal linear layer, which leads to unsatisfying results as discussed in Sec. 4.4. In fact, *RecoverNet* is very important for the whole framework since it is supervised by the entire sequence’s ground-truth, especially for the fast-moving dance dataset AIST++ [34]. *DenoiseNet* can remove noises in advance while giving a better initial pose sequence to *RecoverNet*, which can reduce the burden in the recovery stage. In summary, the two subnets are both essential and effectively improve the final performance.

4.6 An application: Efficient Pose Labeling in Videos

A large amount of labeled data leads to the success of deep models. However, labeling each frame in videos is labor-intensive and high cost. It is also hard to guarantee continuity among adjacent frames, especially for 3D annotations. Due to the efficiency and smoothness of the pose sequences recovered by *DeciWatch*, reducing the need for dense labeling could be a potential application. We verify

Table 5. Comparison of *MPJPE* with different sampling strategies on 3DPW dataset with EFT [25] pose estimator (*MPJPE* is 90.3mm). *U.-2* and *U.-3* are uniform sampling 2 or 3 frames for every N frames. *U.-R.* conducts both uniform sampling and random sampling. *R.* is random sampling.

<i>Ratio</i>	<i>U.(Ours)</i>	<i>U.-2</i>	<i>U.-3</i>	<i>U.-R.</i>	<i>R.</i>
20%	87.2	89.3	91.5	94.2	97.1
10%	89.0	96.3	103.4	101.3	104.4

Table 6. Exploring impacts of the two *DenoiseNet* and *RecoverNet* subnets with 10% sampling ratio on three dataset and the corresponding backbones. *Ori.* means the original estimator (watching all frames) with 100% sampling ratio. No *RecoverNet* is preliminarily recovered via a temporal linear layer.

Dataset w/Backbone	Ori. (100%)	No <i>DenoiseNet</i>	No <i>RecoverNet</i>	<i>DeciWatch</i>
Human3.6M w/FCN [39]	54.6	54.5	54.7	52.8
3DPW w/PARE [29]	78.9	79.8	81.0	77.2
AIST++ w/SPIN [25]	107.7	91.5	95.3	71.3

the effectiveness of this application on the Human3.6M and AIST++ dataset by directly inputting the sparse ground-truth 3D positions into the *RecoverNet* of *DeciWatch*. Due to limited pages, please see Supp. for details.

5 Conclusion and Future Work

This work proposes a sample-denoise-recover flow as a simple baseline framework for highly efficient video-based 2D/3D pose estimation. Thanks to the lightweight representation and continuity characteristics of human poses, this method can watch one frame in every 10 frame and achieve nearly 10× improvement in efficiency while maintaining competitive performance, as validated in the comprehensive experiments across various video-based human pose estimation and body mesh recovery tasks. There are many opportunities to further improve the proposed baseline solution:

Adaptive sampling and dynamic recovery. In *DeciWatch*, we use a simple uniform sampling strategy for all the joints. In practice, the movements of different joints under different actions vary greatly. Consequently, an adaptive sampling strategy has the potential to further boost the efficiency of video-based pose estimators. For instance, combining multi-modality information (e.g., WIFI, sensors) to relieve visual computation can be interesting. Correspondingly, how to design a dynamic recovery network that can handle non-uniformly sampled poses is an interesting yet challenging problem to explore.

High-performance pose estimator design. While this work emphasizes the efficiency of pose estimators, our results show that watching fewer frames with our framework could achieve better per-frame precision compared with watching each frame. This is in line with the recent findings on multi-view pose estimation methods [5,13,45], showing better results without calculating every possible view simultaneously. We attribute the above phenomena to the same intrinsic principle that it is likely to achieve better results by discarding some untrustworthy estimation results. Therefore, designing such a strategy to achieve the best pose estimation performance is an interesting problem to explore.

Acknowledgement. This work is supported in part by Shenzhen-Hong Kong-Macau Science and Technology Program (Category C) of Shenzhen Science Technology and Innovation Commission under Grant No. SGDX2020110309500101, and Shanghai AI Laboratory.

References

1. Burke, M., Lasenby, J.: Estimating missing marker positions using low dimensional kalman smoothing. *Journal of biomechanics* **49**(9), 1854–1858 (2016)
2. Cai, Y., Wang, Y., Zhu, Y., Cham, T.J., Cai, J., Yuan, J., Liu, J., Zheng, C., Yan, S., Ding, H., et al.: A unified 3d human motion synthesis model via conditional variational auto-encoder. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 11645–11655 (2021)
3. Cao, Z., Hidalgo, G., Simon, T., Wei, S.E., Sheikh, Y.: Openpose: realtime multi-person 2d pose estimation using part affinity fields. *IEEE transactions on pattern analysis and machine intelligence* **43**(1), 172–186 (2019)
4. Choi, S., Choi, S., Kim, C.: Mobilehumanpose: Toward real-time 3d human pose estimation in mobile devices. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2328–2338 (2021)
5. Chu, H., Lee, J.H., Lee, Y.C., Hsu, C.H., Li, J.D., Chen, C.S.: Part-aware measurement for robust multi-view multi-human 3d pose estimation and tracking. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. pp. 1472–1481 (June 2021)
6. Dai, H., Shi, H., Liu, W., Wang, L., Liu, Y., Mei, T.: Fasterpose: A faster simple baseline for human pose estimation. *arXiv preprint arXiv:2107.03215* (2021)
7. Desmarais, Y., Mottet, D., Slangen, P., Montesinos, P.: A review of 3d human pose estimation algorithms for markerless motion capture. *Computer Vision and Image Understanding* **212**, 103275 (2021)
8. Duan, Y., Shi, T., Zou, Z., Lin, Y., Qian, Z., Zhang, B., Yuan, Y.: Single-shot motion completion with transformer. *arXiv preprint arXiv:2103.00776* (2021)
9. Fan, Z., Liu, J., Wang, Y.: Adaptive computationally efficient network for monocular 3d hand pose estimation. In: *European Conference on Computer Vision*. pp. 127–144. Springer (2020)
10. Fan, Z., Liu, J., Wang, Y.: Motion adaptive pose estimation from compressed videos. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 11719–11728 (2021)
11. Fragkiadaki, K., Levine, S., Felsen, P., Malik, J.: Recurrent network models for human dynamics. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 4346–4354 (2015)
12. Gløersen, Ø., Federolf, P.: Predicting missing marker trajectories in human motion data using marker intercorrelations. *PloS one* **11**(3), e0152616 (2016)
13. Gundavarapu, N.B., Srivastava, D., Mitra, R., Sharma, A., Jain, A.: Structured aleatoric uncertainty in human pose estimation. In: *CVPR Workshops*. vol. 2, p. 2 (2019)
14. Harvey, F.G., Pal, C.: Recurrent transition networks for character locomotion. In: *SIGGRAPH Asia 2018 Technical Briefs*. pp. 1–4 (2018)
15. Harvey, F.G., Yurick, M., Nowrouzezahrai, D., Pal, C.: Robust motion in-betweening. *ACM Transactions on Graphics (TOG)* **39**(4), 60–1 (2020)
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
17. Hernandez, A., Gall, J., Moreno-Noguer, F.: Human motion prediction via spatio-temporal inpainting. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 7134–7143 (2019)

18. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
19. Ho, H.I., Chen, X., Song, J., Hilliges, O.: Render in-between: Motion guided video synthesis for action interpolation. arXiv preprint arXiv:2111.01029 (2021)
20. Howarth, S.J., Callaghan, J.P.: Quantitative assessment of the accuracy for three interpolation techniques in kinematic analysis of human movement. *Computer methods in biomechanics and biomedical engineering* **13**(6), 847–855 (2010)
21. Hwang, D.H., Kim, S., Monet, N., Koike, H., Bae, S.: Lightweight 3d human pose estimation network training using teacher-student learning. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 479–488 (2020)
22. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence* **36**(7), 1325–1339 (2013)
23. Jhuang, H., Gall, J., Zuffi, S., Schmid, C., Black, M.J.: Towards understanding action recognition. In: Proceedings of the IEEE international conference on computer vision. pp. 3192–3199 (2013)
24. Ji, L., Liu, R., Zhou, D., Zhang, Q., Wei, X.: Missing data recovery for human mocap data based on a-lstm and ls constraint. In: 2020 IEEE 5th International Conference on Signal and Image Processing (ICSIP). pp. 729–734. IEEE (2020)
25. Joo, H., Neverova, N., Vedaldi, A.: Exemplar fine-tuning for 3d human model fitting towards in-the-wild 3d human pose estimation. In: 2021 International Conference on 3D Vision (3DV). pp. 42–52. IEEE (2021)
26. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7122–7131 (2018)
27. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019)
28. Kaufmann, M., Aksan, E., Song, J., Pece, F., Ziegler, R., Hilliges, O.: Convolutional autoencoders for human motion infilling. In: 2020 International Conference on 3D Vision (3DV). pp. 918–927. IEEE (2020)
29. Kocabas, M., Huang, C.H.P., Hilliges, O., Black, M.J.: Pare: Part attention regressor for 3d human body estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11127–11137 (2021)
30. Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2252–2261 (2019)
31. Kucherenko, T., Beskow, J., Kjellström, H.: A neural network approach to missing marker reconstruction in human motion capture. arXiv preprint arXiv:1803.02665 (2018)
32. Lai, R.Y., Yuen, P.C., Lee, K.K.: Motion capture data completion and denoising by singular value thresholding. In: Eurographics (Short Papers). pp. 45–48 (2011)
33. Li, J., Bian, S., Zeng, A., Wang, C., Pang, B., Liu, W., Lu, C.: Human pose regression with residual log-likelihood estimation. In: ICCV (2021)
34. Li, R., Yang, S., Ross, D.A., Kanazawa, A.: Ai choreographer: Music conditioned 3d dance generation with aist++. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 13401–13412 (October 2021)

35. Li, Z., Ye, J., Song, M., Huang, Y., Pan, Z.: Online knowledge distillation for efficient pose estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11740–11750 (2021)
36. Liu, W., Bao, Q., Sun, Y., Mei, T.: Recent advances in monocular 2d and 3d human pose estimation: A deep learning perspective. arXiv preprint arXiv:2104.11536 (2021)
37. Luo, Y., Ren, J., Wang, Z., Sun, W., Pan, J., Liu, J., Pang, J., Lin, L.: Lstm pose machines. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5207–5215 (2018)
38. von Marcard, T., Henschel, R., Black, M.J., Rosenhahn, B., Pons-Moll, G.: Recovering accurate 3d human pose in the wild using imus and a moving camera. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 601–617 (2018)
39. Martinez, J., Hossain, R., Romero, J., Little, J.J.: A simple yet effective baseline for 3d human pose estimation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2640–2649 (2017)
40. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: European conference on computer vision. pp. 483–499. Springer (2016)
41. Nie, X., Li, Y., Luo, L., Zhang, N., Feng, J.: Dynamic kernel distillation for efficient pose estimation in videos. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6942–6950 (2019)
42. Osokin, D.: Real-time 2d multi-person pose estimation on cpu: Lightweight open-pose. arXiv preprint arXiv:1811.12004 (2018)
43. Pavllo, D., Feichtenhofer, C., Grangier, D., Auli, M.: 3d human pose estimation in video with temporal convolutions and semi-supervised training. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7753–7762 (2019)
44. Reda, H.E.A., Benaoumeur, I., Kamel, B., Zoubir, A.F.: Mocap systems and hand movement reconstruction using cubic spline. In: 2018 5th International Conference on Control, Decision and Information Technologies (CoDIT). pp. 1–5. IEEE (2018)
45. Shuai, H., Wu, L., Liu, Q.: Adaptively multi-view and temporal fusing transformer for 3d human pose estimation. arXiv preprint arXiv:2110.05092 (2021)
46. Skurowski, P., Pawlyta, M.: Gap reconstruction in optical motion capture sequences using neural networks. *Sensors* **21**(18), 6115 (2021)
47. Sovrasov, V.: Flops counter for convolutional networks in pytorch framework (Mar 2022), <https://github.com/sovrasov/flops-counter.pytorch>, original-date: 2018-08-17T09:54:59Z
48. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5693–5703 (2019)
49. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
50. Wu, Q., Boulanger, P.: Real-time estimation of missing markers for reconstruction of human motion. In: 2011 XIII Symposium on Virtual Reality. pp. 161–168. IEEE (2011)
51. Xiao, B., Wu, H., Wei, Y.: Simple baselines for human pose estimation and tracking. In: Proceedings of the European conference on computer vision (ECCV). pp. 466–481 (2018)

52. Xu, J., Wang, M., Gong, J., Liu, W., Qian, C., Xie, Y., Ma, L.: Exploring versatile prior for human motion via motion frequency guidance. In: 2021 International Conference on 3D Vision (3DV). pp. 606–616. IEEE (2021)
53. Yan, S., Li, Z., Xiong, Y., Yan, H., Lin, D.: Convolutional sequence generation for skeleton-based action synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4394–4402 (2019)
54. Yu, C., Xiao, B., Gao, C., Yuan, L., Zhang, L., Sang, N., Wang, J.: Lite-hrnet: A lightweight high-resolution network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10440–10450 (2021)
55. Yuan, Y., Iqbal, U., Molchanov, P., Kitani, K., Kautz, J.: Glamr: Global occlusion-aware human mesh recovery with dynamic cameras. arXiv preprint arXiv:2112.01524 (2021)
56. Zeng, A., Sun, X., Huang, F., Liu, M., Xu, Q., Lin, S.C.F.: Srnet: Improving generalization in 3d human pose estimation with a split-and-recombine approach. In: ECCV (2020)
57. Zeng, A., Sun, X., Yang, L., Zhao, N., Liu, M., Xu, Q.: Learning skeletal graph neural networks for hard 3d pose estimation. In: Proceedings of the IEEE International Conference on Computer Vision (2021)
58. Zeng, A., Yang, L., Ju, X., Li, J., Wang, J., Xu, Q.: Smoothnet: A plug-and-play network for refining human poses in videos. arXiv preprint arXiv:2112.13715 (2021)
59. Zhang, Y., Wang, Y., Camps, O., Sznaiier, M.: Key frame proposal network for efficient pose estimation in videos. In: European Conference on Computer Vision. pp. 609–625. Springer (2020)
60. Zhang, Z., Tang, J., Wu, G.: Simple and lightweight human pose estimation. arXiv preprint arXiv:1911.10346 (2019)
61. Zhao, L., Wang, N., Gong, C., Yang, J., Gao, X.: Estimating human pose efficiently by parallel pyramid networks. *IEEE Transactions on Image Processing* **30**, 6785–6800 (2021)
62. Zheng, C., Mendieta, M., Wang, P., Lu, A., Chen, C.: A lightweight graph transformer network for human mesh reconstruction from 2d human pose. arXiv preprint arXiv:2111.12696 (2021)
63. Zheng, C., Wu, W., Yang, T., Zhu, S., Chen, C., Liu, R., Shen, J., Kehtarnavaz, N., Shah, M.: Deep learning-based human pose estimation: A survey. arXiv preprint arXiv:2012.13392 (2020)
64. Zheng, C., Zhu, S., Mendieta, M., Yang, T., Chen, C., Ding, Z.: 3d human pose estimation with spatial and temporal transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11656–11665 (2021)