

UPose3D: Uncertainty-Aware 3D Human Pose Estimation with Cross-View and Temporal Cues

Vandad Davoodnia^{1,2}, Saeed Ghorbani², Marc-André Carbonneau²,
Alexandre Messier², and Ali Etemad¹

¹ Queen’s University, Canada

² Ubisoft LaForge, Canada

{vandad.davoodnia, ali.etemad}@queensu.ca

{saeed.ghorbani, marc-andre.carbonneau2, alexandre.messier}@ubisoft.com

Abstract. We introduce UPose3D, a novel approach for multi-view 3D human pose estimation, addressing challenges in accuracy and scalability. Our method advances existing pose estimation frameworks by improving robustness and flexibility without requiring direct 3D annotations. At the core of our method, a pose compiler module refines predictions from a 2D keypoints estimator that operates on a single image by leveraging temporal and cross-view information. Our novel cross-view fusion strategy is scalable to any number of cameras, while our synthetic data generation strategy ensures generalization across diverse actors, scenes, and viewpoints. Finally, UPose3D leverages the prediction uncertainty of both the 2D keypoint estimator and the pose compiler module. This provides robustness to outliers and noisy data, resulting in state-of-the-art performance in out-of-distribution settings. In addition, for in-distribution settings, UPose3D yields performance rivalling methods that rely on 3D annotated data while being the state-of-the-art among methods relying only on 2D supervision.

Keywords: Markerless Motion Capture · Multi-view Human Pose Estimation · Uncertainty Modeling · Temporal Learning

1 Introduction

Multi-view 3D human pose estimation is a challenging task in computer vision that involves determining the 3D position of human body landmarks given videos or images from multiple synchronized cameras [14, 19, 60]. Compared to monocular setups, multi-view pose estimation leverages information from different viewpoints, alleviating the single-camera ambiguity and improving accuracy in challenging situations. This robustness is crucial in precision-demanding applications like markerless motion capture, essential to industries such as video gaming and film-making, where sub-centimeter accuracy is often required.

The conventional 3D keypoint estimation process involves two stages. Firstly, 2D landmarks are extracted from each camera view. This is followed by triangulation using the known camera parameters to infer 3D points [19, 66]. However, the

accuracy of such methods heavily relies on the precision of independent 2D predictions across views, which is problematic in scenarios with complex body-part interactions or severe occlusions. Outlier mitigation techniques such as RANdom-SAmple Consensus (RANSAC) [43] and refinement algorithms [60] offer some robustness but cannot fully address these inherent limitations. Recent advances in deep learning models that use cross-view fusion strategies [14, 47, 55] have yielded promising 3D pose estimation results. For instance, Epipolar Transformers [14] shows the benefits of leveraging cross-view information using epipolar geometry. However, the scalability of such approaches is often hindered with additional cameras, requiring advanced techniques to maximize the agreement between several model outputs [47]. Another research direction aims to leverage rich temporal information [55, 58] to enhance pose estimation accuracy. For instance, numerous works show the impact of using large temporal context for monocular 3D pose estimation [48, 70]. Similarly, methods like MFT-Transformers [55], focusing on multi-view fusion and temporal modelings, can yield improvements over single-frame methods. However, such approaches require access to large annotated 3D datasets (multi-view video streams and corresponding 3D pose coordinates) during training, which is especially scarce in outdoor and in-the-wild settings. Furthermore, these models are often trained on limited pose and camera variations, hindering their generalization to novel views.

To address the challenges of viewpoint scalability and reliance on 3D annotated training data, we introduce UPose3D, a new method for 3D human pose estimation. Our method leverages 2D keypoints and their uncertainties from two sources to improve robustness to outliers and noisy data. These sources include: *a*) direct 2D pose estimation from RGB images, and *b*) a pose compiler module that utilizes consistency across views and over time. Additionally, we introduce a new cross-view fusion strategy to ensure scalability to different numbers of cameras before our pose compiler. Specifically, we project the keypoints from all available views onto a reference view to obtain a 2D point cloud for each joint. These are then fed to a point cloud transformer module to learn cross-view representations. These features are then passed to a spatiotemporal encoder to efficiently process temporal and skeleton information from temporal windows. To train our pose compiler without relying on 3D annotated datasets, we generate synthetic data simulating realistic multi-view human pose recordings from a large-scale motion capture dataset. This approach promotes generalization across diverse camera configurations and postures, overcoming the limitations of real-world, multi-view 3D annotated datasets.

To evaluate the performance of our proposed model, we use four widely used public datasets, Human3.6m [18], RICH [16], HUMBI [64, 65], and CMU Panoptic [22], across two separate experiment setups. More specifically, we assess the performance of our method in both in-distribution (InD) and out-of-distribution (OoD) settings to better evaluate generalizability to new environments and multi-view camera setups. These experiments demonstrate that our approach outperforms prior state-of-the-art solutions in OoD while achieving competitive performance in InD settings. Next, we provide detailed ablation experiments

demonstrating the impact of uncertainty modeling and our pose compiler in improving multi-view pose estimation robustness to outliers. Finally, we perform several experiments to showcase the impact of the number of camera views and larger time windows. In summary, our method achieves a high level of view-point flexibility and robustness without requiring direct 3D annotations. Our contributions can be summarized as follows:

- We present UPose3D, a 3D human pose estimation pipeline for multi-view setups that achieves state-of-the-art results in OoD settings and performs competitively in InD evaluations.
- Our method uses a novel uncertainty-aware 3D pose estimation algorithm that uses normalizing flows to leverage 2D pose distribution modeling. Experiments demonstrate that this approach is more effective than prior works using heatmaps in terms of accuracy and computational costs. This also allows our pipeline to scale to different camera setups with *constant* runtime.
- We propose a new training strategy that relies only on synthetic multi-view motion sequences generated online from motion capture data.

2 Related Works

3D Pose Estimation. Traditionally, triangulation techniques like RANSAC have been used for 3D human pose estimation [19, 66]. However, these methods are generally not directly differentiable and their integration into deep learning pipelines is hindered. Therefore, recent research has explored more flexible, soft-predictive models, such as volumetric 3D keypoint representations [19] and cross-view feature fusion strategies [47]. Another approach [2] introduces a stochastic framework for human pose triangulation that relies on 3D pose hypothesis generation, scoring, and selection from 2D detection of several camera views. However, the accuracy of 2D pose detectors limits most such approaches. As a result, some methods incorporate epipolar geometry for pose consistency via self-supervised [27] and semi-supervised [63] learning. More advanced techniques have explored feature fusion strategies using epipolar lines across camera pairs [14, 66], showing significant improvements. However, these methods require complex re-projection in view pairs, limiting their scalability to a large number of cameras or with different placements due to lack of variety in the training set.

To avoid relying on large parameters, a recent work [55] proposed a multi-view temporal transformer network for end-to-end 3D pose estimation by leveraging cross-view feature fusion. Other works have conducted experiments on using human pose priors using generative models such as GANs [25] and Diffusion probabilistic models [15] to perform 3D pose estimations from a variety of inputs. However, most such approaches rely on 3D pose annotations from in-studio collected datasets to train their model. Another research track explores the benefits of uncertainty modeling for 3D human pose estimation for enhancing the performance in occlusion-heavy scenarios [3, 29]. For example, Residual Log-likelihood Estimation (RLE) [31], has been proposed to model the underlying




Fig. 1: We illustrate the key stages of UPose3D. It begins with extracting 2D keypoints and uncertainties from the multi-view videos. The keypoints are then projected onto each reference view using epipolar geometry. Our pose compiler is then used to refine the predictions by leveraging cross-view and spatiotemporal information. Finally, the 3D pose is obtained using the keypoint and uncertainty predictions of each stage.

distribution of 2D keypoints in human pose estimation via regression. This approach leverages a re-parameterization technique and normalizing flows to learn keypoint uncertainty and achieves similar performance to heatmap-based [4, 57] and SimCC-based [33] techniques. We exploit the ability of RLE to estimate a differentiable likelihood of keypoints, which other methods lack.

Transformers in Pose Estimation. The Transformer architecture and its self-attention mechanism have significantly advanced Natural Language Processing and Computer Vision. Self-attention’s ability to capture long-range dependencies makes it invaluable for 3D pose estimation, which requires careful consideration of spatial, temporal, and multi-view information. Recent works effectively leverage Transformers for 3D pose estimation in both single-camera setups [35, 71] and for handling spatiotemporal information within single images [32, 38, 68]. Additionally, transformers show promise in aggregating multi-view clues via epipolar geometry [14, 41]. Despite the impressive performance of transformers in a variety of tasks, their memory requirements cause an obstacle in processing all spatial, temporal, and multi-view information together. As a result, a group of researchers has adopted a criss-cross attention mechanism to limit each attention layer’s receptive field without sacrificing the network’s overall receptive field [17]. Recent work has also adopted the criss-cross attentions [58] to process temporal and joint information for human 3D pose estimation, showing superior performance compared to parallel or concurrent models.

3 Method

Upose3D determines the 3D coordinates of body joints from one or more multi-view frames. Figure 1 depicts an overview of our proposed method. At first, an RLE-based 2D pose estimator (Sec. 3.1) extracts 2D keypoints and their corresponding uncertainties from the input images. The keypoints are then projected

onto other views using epipolar geometry (Sec. 3.2). From there, the pose compiler (Sec. 3.3) refines the keypoint locations and their uncertainties by leveraging spatiotemporal and cross-view information. Lastly, using outputs from the 2D pose estimator and the pose compiler, an iterative algorithm obtains the final 3D poses (Sec. 3.4). This section finished with a description of our multi-view training data synthesis (Sec. 3.5).

3.1 2D Pose Estimation

The first step of our method is estimating 2D joint position in each image $\mathcal{I}_{i,t}$ from all camera views $i \in \mathcal{V} = \{1, 2, \dots, V\}$ and frames $t \in \mathcal{T} = \{1, 2, \dots, T\}$. Similar to [8, 30], our method implements a single layer Residual Log-likelihood Estimation (RLE) head [31] on top of an off-the-shelf backbone (*e.g.*, CPN [5]). Aside from being computationally effective and robust to occlusion, RLE provides uncertainty $\hat{\sigma}$ for each joint position prediction $\hat{\mu}$. Specifically, the RLE predicts a distribution $P_{\Theta}(x|\mathcal{I})$ that models the probability of the ground truth appearing in position x using a normalizing flow model [52] with learnable parameters Θ . The $\hat{\mu}$ produced by this module is refined in the next stages by leveraging cross-view and temporal information. The estimated $\hat{\mu}$, $\hat{\sigma}$, along with the normalizing flow parameters Θ , are used for 3D keypoint estimation in Sec. 3.4.

3.2 Cross-view Projection

Next, we leverage the information from multiple camera views by projecting the 2D keypoints from one view to another using epipolar geometry. We derive a fundamental matrix $F_{ij} \in \mathbb{R}^{3 \times 3}$ relating two camera views $i, j \in \mathcal{V}$ from known intrinsic and extrinsic camera parameters. For each predicted keypoint $\hat{\mu}_j$ in view j , we obtain the epipolar line in the reference view i by $I_{ij} = F_{ij}^T \hat{\mu}_j$ [12]. Next, we find the closest point on I_{ij} to the keypoint $\hat{\mu}_i$ in view i . This point represents the projection of the keypoint from view j onto the reference view i . By repeating this cross-view projection for all available views i , joints $k \in \mathcal{J} = \{1, 2, \dots, J\}$, and time frames t , we create 2D point clouds $\mathcal{C}_{i,k,t}$ containing the projected keypoint. Finally, for consistency between various scales and views, we normalize the elements of \mathcal{C} using the subject’s bounding box within their reference view.

3.3 Pose Compiler

The pose compiler module aggregates multi-view information embedded in the 2D point clouds and leverages temporal information (*i.e.*, joint positions across time). As shown in Fig. 1, the pose compiler consists of 2 main parts: a point cloud encoder and a spatiotemporal encoder. The point cloud encoder first extracts a feature vector $f_{i,k,t}$ that describes a point cloud $\mathcal{C}_{i,k,t}$. Our implementation is inspired by the naïve Point Cloud Transformers [11, 51], but we modify its architecture to preserve coordinate information. Specifically, we use multi-head attention and a residual connection from the features of the reference view to




Fig. 2: Architecture of the proposed pose compiler module consisting of a point cloud encoder and a spatiotemporal encoder with criss-cross attention. Tensor sizes depend on the batch size B , temporal window length T , number of joints J , camera views V and the point cloud feature dimensionality H .

the output of the last max-pooling layer. For the detailed architecture of our point cloud encoder, please refer to Fig. 2.

Next, we concatenate the feature vectors from every frame and joint in each reference view $\{\mathbf{f}_i\}$ with temporal and spatial position embeddings, *i.e.*, frame and skeleton joint ID. The results are then passed to our spatiotemporal encoder that implements an RLE head to provide position $\hat{\mu}'_{i,k,t}$ and uncertainty $\hat{\sigma}'_{i,k,t}$ estimations. Our spatiotemporal encoder is a transformer [61], with specific attention block modifications that accommodate the large dimensionality of the tensors containing all temporal and skeleton information. As detailed in Fig. 2, we replace the standard attention module with criss-cross attentions [17, 58], which approximates full temporal and spatial dependencies while being more memory efficient than full attention. We train the RLE head of the spatiotemporal encoder using the same strategy as [31]. Finally, similarly to Sec. 3.1, the RLE head simultaneously maximizes and learns a distribution $P_{\Theta'}(x|\mathcal{C})$ that represents the probability of the appearance of a keypoint in position x using a normalizing flow model with learnable parameters Θ' . In the next section, we use the estimated $\hat{\mu}'$, $\hat{\sigma}'$, and Θ' for 3D pose estimation.

3.4 3D Pose Estimation

As the final step, our method obtains the 3D keypoints using position estimates from 2D frames $(\hat{\mu}, \hat{\sigma})$, as well as refined versions from the pose compiler $(\hat{\mu}', \hat{\sigma}')$. Specifically, we use the estimated keypoints as labels when sampling from the associated keypoint density functions ($P_{\Theta}(x|\mathcal{I})$ and $P_{\Theta'}(x|\mathcal{C})$) during Maximum Likelihood Estimation (MLE). Therefore, the total loss function of MLE for each joint and each time frame is defined as:

$$\mathcal{L}_{mle} = -\log \prod_{i \in \mathcal{V}} P_{\Theta}(u_i|\mathcal{I}) \Big|_{u_i = \hat{\mu}_i} - \log \prod_{i \in \mathcal{V}} P_{\Theta'}(u_i|\mathcal{C}) \Big|_{u_i = \hat{\mu}'_i}, \quad (1)$$

where u_i is the projection of variable 3D point U onto each camera view i . Minimizing this loss function increases the likelihood of U appearing close to the ground-truth 3D keypoints U_g without 3D supervision. To solve this non-convex optimization problem, we initialize U with a Direct Linear Transformation (DLT) algorithm [12] and iteratively refine it using an optimizer, *i.e.*, L-BFGS [37].

3.5 Training and Multi-view Data Synthesis

Our proposed pose compiler does not require 3D annotated data, given that it operates on point clouds. Therefore, we can use any animation data to create a synthetic training set. The synthesis of multi-view training data starts by randomly positioning several cameras in 3D space to look roughly at the human body’s center. During this process, we limit the camera height, distance, tilt, pitch, and yaw within a representative range of a normal multi-view recording setup. Next, we extract the 3D landmarks and project them onto each camera to obtain ground-truth 2D keypoints μ_g . We then corrupt these keypoints through extensive augmentations. Refer to our Supplementary Materials for more details on our data synthesis. Finally, we obtain the point clouds training data from the corrupted 2D keypoint using the cross-view projection from Sec. 3.2.

4 Experiments

4.1 Datasets

We compare our proposed approach with prior works on the Human3.6m [18] dataset and measure how our method generalizes to different skeleton configurations and unseen outdoor environments on the RICH [16] dataset. We use the COCO WholeBody [21, 36] and MPII [1] for training our 2D pose estimator in different scenarios (refer to the Supplementary Materials for more information [69]). The pose compiler is trained using motion capture data from the AMASS dataset [42] to simulate multi-view training data. Finally, we analyze the viewpoint scalability of our approach on the CMU Panoptic [22] and HUMBI [64, 65] datasets. The details of these datasets are as follows.

Human3.6m. The Human3.6m [18] dataset is a standard benchmark for evaluating the performance of 3D human pose estimation solutions in multi-view [55] and single-view [70] settings. We report our InD performance in Protocol-I using subjects 1, 5, 6, 7, and 8 for training and 9 and 11 for testing.

RICH. Real scenes, Interaction, Contact, and Humans (RICH) dataset [16] is a dataset of multi-view videos with accurate bodies and 3D-scanned scenes obtained using markerless motion capture technology. We extract 3D ground-truth keypoints from the provided SMPL-x [46] bodies for evaluation. We report OoD performance on the test set, which contains 7 subjects in 52 scenarios and 3 environments, including a construction site, a gym, and a lecture hall.

CMU Panoptic. The CMU Panoptic [22] dataset is a large collection of human body poses and interactions recorded from multiple views. The dataset provides

videos from 480 VGA and 30 HD cameras and ground-truth 3D poses obtained from markerless 2D pose estimation and triangulation. We use a small subset of the validation set in line with [19] to analyze UPose3D’s camera scalability.

HUMBI. The HUman Multiview Behavioral Imaging (HUMBI) [64,65] dataset is a large-scale multi-view dataset designed for modeling and reverse-rendering of the human body in different expressions. It contains 107 HD recordings of 772 subjects with natural clothing. We use this dataset only for OoD evaluations during our scalability analysis. Specifically, we evaluate on the first 80 subjects by extracting ground-truth 3D keypoints from the provided SMPL [39] parameters.

AMASS. The Archive of Motion Capture as Surface Shapes (AMASS) [42] is a large collection of 3D human motion capture datasets. It contains over 40 hours of recording from 300 subjects spanning 11,000 actions. We only use the training set of this dataset following prior works [7,46,50]. Since the extraction of landmarks requires a specialized joint regressor for each skeleton configuration, in our InD experiments, we use the SMPL+H [39] body model and a 17-joint regressor provided by [44]. For our OoD experiments, we use the SMPL-x [46] body model and a whole-body 133-joint regressor from [7].

4.2 Evaluation Metrics

We adopt the standard evaluation metrics from prior works for our InD experiments on the Human3.6m [18] dataset. Mean Per Joint Position Error (MPJPE) represents the distance between the ground-truth and predicted pose in millimeters, while its Procrustes Aligned version (PA-MPJPE) and the Normalized variant (N-MPJPE) show how well the predicted pose fits the ground-truth keypoints after a similarity transformation. Next, we report the Translation Aligned error (TA-MPJPE) and PA-MPJPE for our OoD experiments on the RICH [16] dataset following prior works [26,28,34,35,54,59].

4.3 Implementation Details

We implement our pipeline using PyTorch [45] and MMPose [6]. Here, we provide the details of our neural networks and the computation costs of our pose compiler. We encourage readers to refer to our Supplementary Materials for information on our multi-view data synthesis and augmentation strategies.

To compare our method with prior works, we choose a CPN [5] backbone based on ResNet152 [13] with an input size 384×384 and an RLE head [31] for 2D pose estimation. We fine-tune our 2D pose estimator jointly on Human3.6m [18] and MPII [1] datasets starting from the weights of the checkpoints provided by [19]. For our OoD experiments, we use an HRNet-W48 [57] with 384×288 input size and train it from scratch. Our point cloud encoder uses 4 multi-head self-attention blocks, each with 4 attention heads. Additionally, we reduce its hidden dimension to 64 to accommodate the relatively smaller size of our point clouds. Following the point cloud encoder, the features are concatenated with two positional embeddings of size 64 to represent the time frames and different skeleton joints. Our spatiotemporal encoder employs four criss-cross transformer

blocks [17, 58] with 4 attention heads each. We set the hidden layer dimension of our transformers to 64 and the hidden dimension of prediction heads to 1024. We then train our pose compiler module using AdamW optimizer [40] with a batch size of 64 on the AMASS [42] dataset for both OoD and InD evaluations. We use a learning rate of $4e-5$ with a warm-up factor of $1e-4$ for the first 500 iterations and a cosine annealing scheduler over the next 20,000 iterations. The network training takes about 6 hours on an NVIDIA 2080 RTX GPU.

Unlike prior works that rely on rendering techniques for synthetic data generation [9, 66], we generate samples online with varying numbers of camera views (up to 8) during training without significant computation overhead. Furthermore, we increase the diversity of the motion capture data by applying several augmentations to the motion capture data before 3D keypoint extraction, such as random 180-degree rotations and mirroring around the mid-hip point.

4.4 Baselines

On Human3.6m [18], we compare our UPose3D with methods that infer 3D human keypoints using only 2D supervised models [5, 14, 23, 27, 47, 60]. Additionally, we provide a summary of methods that rely on direct 3D annotations [10, 19, 20, 49, 55, 66], except for a few [10, 20, 48] that use uncalibrated cameras. Moreover, we compare our work with methods that only rely on multi-view 2D keypoints along with the camera parameters [14, 41, 47, 67]. Please refer to our Supplementary Materials for a more in-depth comparison with weakly-supervised and semi-supervised approaches [9, 27, 53, 56, 62], multi-view methods without that do not rely on camera parameters [10, 20, 48], and monocular 3D pose estimation methods [26, 28, 34, 35, 54, 59]. We also report and compare the performance of our method against triangulation approaches, such as DLT, RANSAC, and RPSM [47]. On the RICH dataset [16], we compare our model with triangulation approaches and replicate the performance of AdaFuse [66], which is one of the top-performing methods on the Human3.6m dataset [18].

5 Results

5.1 In-Distribution Performance

In Tab. 1, we present a comparison of our InD results with other state-of-the-art methods on the Human3.6m dataset, divided into three categories: *a*) methods that, unlike ours, are trained directly with 3D annotated data; *b*) methods that leverage only 2D supervision; and *c*) methods that rely on 2D supervision along with the camera parameters. We observe that UPose3D outperforms all other methods that rely on 2D supervision. In addition, it yields results that are competitive to state-of-the-art methods relying on 3D supervision. For example, our method achieves similar results to MTF-Transformers [55], which, like our pose compiler, takes in multi-view and temporal 2D keypoints as input but estimates the 3D poses directly using a deep network with $10.1M$ parameters. Furthermore,

Table 1: The comparison of our method in InD settings against prior multi-view works on the full test set of the Human3.6m dataset. The errors are reported in *mm*.

Method	Backbone	Frames	MPJPE↓	PA-MPJPE↓	N-MPJPE↓
<i>3D Supervision</i>					
Learnable Triangulation [19]	ResNet152	1	20.7	17.0	-
Canonical Fusion [49]	ResNet152	1	21.0	-	-
AdaFuse [66]	ResNet152	1	19.5	-	-
TesseTrack [48]	HRNet [57]	5	18.7	-	-
MTF-Transformer+ [55]	CPN	1	26.5	-	-
MTF-Transformer+ [55]	CPN	27	25.8	-	-
Flex [10]	ResNet152	All	30.2	-	-
Jiang <i>et al.</i> [20]	ResNet152	1	27.8	-	-
<i>2D Supervision</i>					
EpipolarPose [27]	ResNet50	1	55.08	47.91	54.90
AniPose [23, 60]	GT	1	-	75.0	103.0
MetaPose [60]	PoseNet [24]	1	-	32.0	49.0
DLT [47]	ResNet152	1	36.3	-	-
DLT [5]	CPN	1	30.5	27.6	29.8
UPose3D (Ours)	ResNet152	1	31.0	29.0	31.2
UPose3D (Ours)	ResNet152	27	29.9	27.2	29.8
UPose3D (Ours)	CPN	1	26.9	24.1	26.2
UPose3D (Ours)	CPN	27	26.4	23.4	25.6
<i>2D Supervision + Camera Parameters</i>					
TransFusion [41]	ResNet50	1	25.8	-	-
DLT* [13]	ResNet50	1	71.2	62.8	70.9
UPose3D (Ours)*	ResNet50	27	57.3	51.7	57.2
Cross-view Fusion [47]	ResNet152	1	26.2	-	-
Epipolar Transformers [14]	ResNet152	1	19.0	-	-
TR Loss† [67]	HRNet [57]	1	25.8	-	-
UPose3D (Ours)	CPN	27	26.2	23.1	25.4

- denotes the error was not reported in the original work.

* denotes training from scratch with no additional data.

† denotes using 1% of unlabeled testing data during training.

among methods that use the camera parameters of the Human3.6m dataset, our method performs similarly to TR Loss [67], which uses 1% of the unlabeled *testing* set during training. Although Epipolar Transformers [14] perform better than ours due to their effective feature fusion strategy, it relies on precise camera parameters and has constraints on the viewing angle of neighbouring cameras. The consistent performance of our approach in both groups indicates the effectiveness of our proposed pose compiler and training strategies in obtaining precise keypoints regardless of the target dataset’s camera setup.

We also show that the choice of the 2D keypoint estimator backbone is an important factor for accuracy and that CPN is preferred over ResNet152. Interestingly, a vanilla triangulation algorithm, DLT, can outperform prior works given accurate 2D predictions from strong pose estimators. Combining strong pose estimators with our method yields even better performance. For instance, UPose3D improves the MPJPE by 3.6 *mm* compared to vanilla triangulation using a CPN backbone (DLT+CPN). Additionally, we experiment with a ResNet-50 model trained from scratch in the third group (2D supervised with camera parameters) and observe a 19.5% improvement when UPose3D is used. Finally, we demonstrate the positive impact of the temporal frame window, which reduces the MPJPE by 0.5 *mm* when using 27 frames instead of a single frame.

Table 2: Comparison of our method in OoD setting on the RICH [16] dataset against prior works. The training sources are denoted as (2D pose estimator, Pose compiler).

Method	Training sources	MPJPE↓	PA-MPJPE↓
AdaFuse* [66]	(Human3.6m+MPII, Human3.6m)	524.0	85.8
Ours ($T = 27$)	(Human3.6m+MPII, Human3.6m)	51.8	43.6
Ours ($T = 1$)	(Human3.6m+MPII, AMASS)	49.9	42.3
HRNet-W48+DLT*	(COCO, N/A)	66.0	55.1
HRNet-W48+Grid Search*	(COCO, N/A)	64.4	54.9
Ours ($T = 1$)	(COCO, AMASS)	36.2	33.4
Ours ($T = 27$)	(COCO, AMASS)	34.7	32.0

* denotes our implementation of prior works.

5.2 Out-of-Distribution Generalization

To evaluate the performance of UPose3D in OoD settings, we compare it to the best-performing baseline on the Human3.6m [18] dataset, namely AdaFuse [66]. Additionally, we report the results for triangulation techniques applied to 2D pose estimators in Tab. 2. In this experiment setup, neither the models nor their components are trained on the RICH [16] dataset. First, we observe that our method, despite the training source dataset, obtains the best results by achieving half of the error of the next best method. Additionally, we observe that AdaFuse performs poorly on this dataset despite being considered state-of-the-art on the Human3.6m [18] dataset. Our analysis of AdaFuse performance on different clips of RICH [16] indicates that this method performs well only if all the viewpoints’ predictions are within a reasonable range, whereas occasional noisy predictions cause large triangulation errors. In contrast, our solution deduces the correct pose by effectively discarding views with low confidence. In conclusion, when training on in-studio datasets, methods such as AdaFuse fail to adapt to in-the-wild environments. At the same time, our approach uses uncertainty modeling to better generalize to unseen camera configurations. UPose3D is not limited to the source dataset skeleton configuration because synthetic data are used for training, and it performs consistently between InD and OoD evaluations.

5.3 Ablation Study

To investigate the impact of each component, we systematically remove them during *inference* and report the results with a CPN backbone on the test set of Human3.6m [18]. We perform all experiments using a model with a temporal window of 27 frames. Please refer to our Supplementary Materials for additional analysis on temporal length, spatiotemporal encoder’s architecture, different point cloud formulations, and initialization strategies for 3D estimation. **Pose Compiler.** We employ the pose compiler module to improve the keypoint and uncertainty predictions using cross-view and spatiotemporal information. By ablating this component, we use the original keypoints and uncertainties predicted by our 2D pose estimator to perform MLE. As shown in Tab. 3, this experiment results in an additional 9.28 *mm* error resulting in a higher error

Table 3: Ablation experiments on the Human3.6m dataset with $T = 27$.

Method	MPJPE↓	PA-MPJPE↓
UPose3D	26.42	23.42
w/o compiler	37.14	33.90
w/o image branch	69.90	50.97
w/o compiler uncertainty	26.42	23.58
w/o image uncertainty	27.61	24.88
w/o uncertainty	48.11	41.20
w/o image branch	77.25	54.02
w/o compiler	30.48	27.63

compared to the DLT algorithm (see Tab. 1). This shows the effectiveness of our pose compiler for 3D pose estimation using normalizing flows.

Image Branch. Removing the image branch from the pipeline causes a significant rise in the error, seen in Tab. 1, indicating the significance of keeping the original predictions for the final estimation.


Uncertainty Modeling. We employ the normalizing flows of our RLE heads during the MLE loss minimization stage to incorporate the uncertainties of our predictions toward improving the model’s robustness. We employ the RLE head’s normalizing flows module to model the uncertainties, improving robustness. To study the effect of uncertainty modeling, we remove the RLE heads from the 2D pose estimator and pose compiler branches. Accordingly, we first remove the uncertainties from our pose compiler and replace our likelihood loss function (Eq. (1)) with a reprojection distance loss. Table 1 shows that the PA-MPJPE error slightly rises, but the MPJPE remains unchanged. However, it should be noted that removing the pose compiler uncertainty during *training* causes more performance degradation by resulting in 26.8 *mm*, which shows the significance of uncertainty modeling. Next, we remove the image pose estimator uncertainties, which results in a 1.2 *mm* rise in the error. Finally, we completely remove the uncertainties from both branches, effectively reducing our problem to a classic triangulation problem without confidence that can be solved via the DLT algorithm. This variation shows over 20 *mm* higher error than UPose3D. In our final two experiments, we analyze the image and pose compiler predictions without uncertainty modeling. The first experiment shows the highest error, while the second experiment results in the same performance as a simple DLT algorithm without cross-view fusion.

5.4 Computation Costs

In Tab. 4, we compare the total computation costs of our method in comparison to state-of-the-art 3D approaches in single-frame scenarios. Compared to others, UPose3D has fewer parameters and significantly less computational cost, especially with more cameras. The computational cost of our 10-frame model can be broken down to 508.5*G* FLOPs for CPN, 0.385*G* for pose compiler, and 8.8 ± 1.25 *G* for the optimization process. Due to the computational cost fluctuations caused by the L-BFGS optimizer, we measure it by averaging 20 runs using randomly selected cameras from the HUMBI dataset [64, 65].

Table 4: Comparison of computation costs.

Method	Params(M) \downarrow	FLOPs for 4 Cams(G) \downarrow	FLOPs for 10 Cams(G) \downarrow
Learnable Triangulation [19]	80.6	716.1	1326.9
Epipolar Transformers [14]	68.1	406.5	1016.2
MTF-Transformers [55]	78.6	407.0	1017.8
AdaFuse [66]	69.7	595.0	1487.6
Ours	65.4	208.7	517.7

**Fig. 3:** We demonstrate the scalability of UPose3D to the number of cameras.

5.5 Viewpoint Scalability

In the context of multi-view 3D human pose estimation, scalability may be defined as the capability of a method to handle data from an increasing number of camera views effectively. Specifically, such a solution should have an almost constant runtime with constantly decreasing error when more cameras are used. Accordingly, to test the scalability of our method, we evaluate it with different numbers of cameras on the CMU Panoptic [22] and the HUMBI [64, 65] datasets and present the results in Fig. 3. For this experiment, we use the HRNet-W48 [57] backbone, similar to our experiments on the RICH dataset, to analyze scalability in OoD scenarios. Additionally, we report and compare our error and runtime with two triangulation techniques, RANSAC and DLT. We tune the RANSAC implementation of [19] and run it for 20 iterations. The error analysis shows that as the number of views increases, UPose3D’s error continuously improves while DLT and RANSAC reach a plateau. Next, we discard the common factor of the 2D pose estimators for runtime scalability analysis. The inference time of our pipeline with a single batch size remains constant despite the number of cameras. However, increasing the batch size causes runtime improvements at the cost of memory and runtime variability.

5.6 Qualitative Analysis

Fig. 4 shows the outputs of our method on challenging examples, where we sample the keypoint distributions in every pixel to demonstrate the keypoint likelihoods from the 2D pose estimator and our pose compiler as heatmaps. Despite choosing challenging scenes, we observe that our predictions are still close to the ground-truth keypoints, indicating that our method does not produce any significant outliers within its output even on the unseen samples of RICH [16].




Fig. 4: Illustration of UPose3D on Human3.6m [18] (right) and RICH [16] (left) datasets, showing the accurate 3D pose estimated by our UPose3D (top) compared to ground-truth (bottom).

6 Conclusion

This paper presents UPose3D, a multi-view 3D human pose estimation method designed to address the challenges in generalization, scalability, and over-reliance on real-world 3D annotated data. It includes a novel cross-view fusion strategy that scales well with varying camera numbers and volume sizes. Additionally, UPose3D integrates a pose compiler that learns to predict keypoint positions and uncertainties given the cross-view and temporal information. Finally, the uncertainties and predicted key points from two sources of the image branch and pose compiler are used for uncertainty-aware 3D pose estimation. UPose3D outperforms state-of-the-art approaches in OoD setup and achieves competitive InD results without any 3D pose annotations. Therefore, UPose3D may positively impact applications where performance and robustness are crucial.

Future Work. Our primary focus in this paper was the fidelity of estimated output 3D poses from multi-view inputs without requiring 3D supervision. We acknowledge, however, that our method could be optimized for *real-time* applications. A promising avenue to achieve faster inference involves the exploration of specialized second-order optimizers that use a deep-learning neural network to estimate the Hessian matrix during likelihood maximization. Furthermore, depth and trajectory modalities can be explored for additional supervision in the likelihood function to reduce noisy predictions.

Acknowledgements

This work was performed during Vandad Davoodnia’s internship at Ubisoft Laforge, partially funded by Mitacs through the Accelerate program.

References

1. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2D human pose estimation: New benchmark and state of the art analysis. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3686–3693 (2014). <https://doi.org/10.1109/CVPR.2014.471>
2. Bartol, K., Bojanić, D., Petković, T., Pribanić, T.: Generalizable human pose triangulation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11028–11037 (2022). <https://doi.org/10.1109/CVPR52688.2022.01075>
3. Bramlage, L., Karg, M., Curio, C.: Plausible uncertainties for human pose regression. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 15133–15142 (2023). <https://doi.org/10.1109/ICCV51070.2023.01389>
4. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2D pose estimation using part affinity fields. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7291–7299 (2017). <https://doi.org/10.1109/CVPR.2017.143>
5. Chen, Y., Wang, Z., Peng, Y., Zhang, Z., Yu, G., Sun, J.: Cascaded pyramid network for multi-person pose estimation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7103–7112 (2018). <https://doi.org/10.1109/CVPR.2018.00742>
6. Contributors, M.: Openmmlab pose estimation toolbox and benchmark. <https://github.com/open-mmlab/mmpose> (2020)
7. Davoodnia, V., Ghorbani, S., Messier, A., Etemad, A.: Skelformer: Markerless 3D pose and shape estimation using skeletal transformers. arXiv preprint arXiv:2404.12625 (2024). <https://doi.org/10.48550/arXiv.2404.12625>
8. Dwivedi, S.K., Schmid, C., Yi, H., Black, M.J., Tzionas, D.: Poco: 3D pose and shape estimation with confidence. In: IEEE International Conference on 3D Vision (3DV). pp. 85–95 (2024). <https://doi.org/10.1109/3DV62453.2024.00115>
9. Gong, X., Song, L., Zheng, M., Planche, B., Chen, T., Yuan, J., Doermann, D., Wu, Z.: Progressive multi-view human mesh recovery with self-supervision. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 676–684 (2023). <https://doi.org/10.1609/aaai.v37i1.25144>
10. Gordon, B., Raab, S., Azov, G., Giryas, R., Cohen-Or, D.: Flex: Extrinsic parameters-free multi-view 3D human motion reconstruction. In: European Conference on Computer Vision (ECCV). pp. 176–196. Springer (2022). https://doi.org/10.1007/978-3-031-19827-4_11
11. Guo, M.H., Cai, J.X., Liu, Z.N., Mu, T.J., Martin, R.R., Hu, S.M.: Pct: Point cloud transformer. *Computational Visual Media* **7**, 187–199 (2021). <https://doi.org/10.1007/s41095-021-0229-5>
12. Hartley, R., Zisserman, A.: Multiple view geometry in computer vision. Cambridge University Press, 2 edn. (2004). https://doi.org/10.1108/k.2001.30.9_10.1333.2

13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016). <https://doi.org/10.1109/CVPR.2016.90>
14. He, Y., Yan, R., Fragkiadaki, K., Yu, S.I.: Epipolar transformers. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7779–7788 (2020). <https://doi.org/10.1109/CVPR42600.2020.00780>
15. Holmquist, K., Wandt, B.: Diffpose: Multi-hypothesis human pose estimation using diffusion models. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 15977–15987 (2023). <https://doi.org/10.1109/ICCV51070.2023.01464>
16. Huang, C.H.P., Yi, H., Höschle, M., Safroshkin, M., Alexiadis, T., Polikovskiy, S., Scharstein, D., Black, M.J.: Capturing and inferring dense full-body human-scene contact. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13274–13285 (2022). <https://doi.org/10.1109/CVPR52688.2022.01292>
17. Huang, Z., Wang, X., Huang, L., Huang, C., Wei, Y., Liu, W.: Ccnet: Criss-cross attention for semantic segmentation. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 603–612 (2019). <https://doi.org/10.1109/ICCV.2019.00069>
18. 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). <https://doi.org/10.1109/TPAMI.2013.248>
19. Isakov, K., Burkov, E., Lempitsky, V., Malkov, Y.: Learnable triangulation of human pose. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 7718–7727 (2019). <https://doi.org/10.1109/ICCV.2019.00781>
20. Jiang, B., Hu, L., Xia, S.: Probabilistic triangulation for uncalibrated multi-view 3D human pose estimation. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 14850–14860 (2023). <https://doi.org/10.1109/ICCV51070.2023.01364>
21. Jin, S., Xu, L., Xu, J., Wang, C., Liu, W., Qian, C., Ouyang, W., Luo, P.: Whole-body human pose estimation in the wild. In: European Conference on Computer Vision (ECCV). pp. 196–214. Springer (2020). https://doi.org/10.1007/978-3-030-58545-7_12
22. Joo, H., Liu, H., Tan, L., Gui, L., Nabbe, B., Matthews, I., Kanade, T., Nobuhara, S., Sheikh, Y.: Panoptic studio: A massively multiview system for social motion capture. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3334–3342 (2015). <https://doi.org/10.1109/ICCV.2015.381>
23. Karashchuk, P., Rupp, K.L., Dickinson, E.S., Walling-Bell, S., Sanders, E., Azim, E., Brunton, B.W., Tuthill, J.C.: Anipose: a toolkit for robust markerless 3D pose estimation. *Cell Reports* **36**(13) (2021). <https://doi.org/10.1016/j.celrep.2021.109730>
24. Kendall, A., Grimes, M., Cipolla, R.: PoseNet: A convolutional network for real-time 6-dof camera relocalization. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2938–2946 (2015). <https://doi.org/10.1109/ICCV.2015.336>
25. Kocabas, M., Athanasiou, N., Black, M.J.: Vibe: Video inference for human body pose and shape estimation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5253–5263 (2020). <https://doi.org/10.1109/CVPR42600.2020.00530>

26. Kocabas, M., Huang, C.H.P., Hilliges, O., Black, M.J.: Pare: Part attention regressor for 3D human body estimation. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 11127–11137 (2021). <https://doi.org/10.1109/ICCV48922.2021.01094>
27. Kocabas, M., Karagoz, S., Akbas, E.: Self-supervised learning of 3D human pose using multi-view geometry. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1077–1086 (2019). <https://doi.org/10.1109/CVPR.2019.00117>
28. Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2252–2261 (2019). <https://doi.org/10.1109/ICCV.2019.00234>
29. Kundu, J.N., Seth, S., YM, P., Jampani, V., Chakraborty, A., Babu, R.V.: Uncertainty-aware adaptation for self-supervised 3D human pose estimation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20448–20459 (2022). <https://doi.org/10.1109/CVPR52688.2022.01980>
30. Li, J., Bian, S., Xu, C., Chen, Z., Yang, L., Lu, C.: Hybrik-x: Hybrid analytical-neural inverse kinematics for whole-body mesh recovery. arXiv preprint arXiv:2304.05690 (2023). <https://doi.org/10.48550/arXiv.2304.05690>
31. Li, J., Bian, S., Zeng, A., Wang, C., Pang, B., Liu, W., Lu, C.: Human pose regression with residual log-likelihood estimation. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 11025–11034 (2021). <https://doi.org/10.1109/ICCV48922.2021.01084>
32. Li, W., Liu, H., Tang, H., Wang, P., Van Gool, L.: Mhformer: Multi-hypothesis transformer for 3D human pose estimation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13147–13156 (2022). <https://doi.org/10.1109/CVPR52688.2022.01280>
33. Li, Y., Yang, S., Liu, P., Zhang, S., Wang, Y., Wang, Z., Yang, W., Xia, S.T.: Simcc: A simple coordinate classification perspective for human pose estimation. In: European Conference on Computer Vision (ECCV). pp. 89–106. Springer (2022). https://doi.org/10.1007/978-3-031-20068-7_6
34. Li, Z., Liu, J., Zhang, Z., Xu, S., Yan, Y.: Cliff: Carrying location information in full frames into human pose and shape estimation. In: European Conference on Computer Vision (ECCV). pp. 590–606. Springer (2022). https://doi.org/10.1007/978-3-031-20065-6_34
35. Lin, K., Wang, L., Liu, Z.: End-to-end human pose and mesh reconstruction with transformers. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1954–1963 (2021). <https://doi.org/10.1109/CVPR46437.2021.00199>
36. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European Conference on Computer Vision (ECCV). pp. 740–755. Springer (2014). https://doi.org/10.1007/978-3-319-10602-1_48
37. Liu, D.C., Nocedal, J.: On the limited memory bfgs method for large scale optimization. *Mathematical Programming* **45**(1-3), 503–528 (1989). <https://doi.org/10.1007/BF01589116>
38. Llopart, A.: Liftformer: 3D human pose estimation using attention models. arXiv preprint arXiv:2009.00348 (2020). <https://doi.org/10.48550/arXiv.2009.00348>
39. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. *ACM Transactions on Graphics* **34**(6), 1–16 (2015). <https://doi.org/10.1145/2816795.2818013>

40. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (ICLR) (2019), <https://openreview.net/forum?id=Bkg6RiCqY7>
41. Ma, H., Chen, L., Kong, D., Wang, Z., Liu, X., Tang, H., Yan, X., Xie, Y., Lin, S.Y., Xie, X.: Transfusion: Cross-view fusion with transformer for 3D human pose estimation. In: British Machine Vision Conference (BMVC). British Machine Vision Association (2021). <https://doi.org/10.48550/arXiv.2110.09554>
42. Mahmood, N., Ghorbani, N., Troje, N.F., Pons-Moll, G., Black, M.J.: Amass: Archive of motion capture as surface shapes. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 5442–5451 (2019). <https://doi.org/10.1109/ICCV.2019.00554>
43. Martínez-Otzeta, J.M., Rodríguez-Moreno, I., Mendiádua, I., Sierra, B.: Ransac for robotic applications: A survey. *Sensors* **23**(1), 327 (2022). <https://doi.org/10.3390/s23010327>
44. Moon, G., Choi, H., Lee, K.M.: Neuralannot: Neural annotator for 3D human mesh training sets. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2299–2307 (2022). <https://doi.org/10.1109/CVPRW56347.2022.00256>
45. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems (NeurIPS)* **32** (2019), <https://pytorch.org/>
46. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3D hands, face, and body from a single image. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10975–10985 (2019). <https://doi.org/10.1109/CVPR.2019.01123>
47. Qiu, H., Wang, C., Wang, J., Wang, N., Zeng, W.: Cross view fusion for 3D human pose estimation. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 4342–4351 (2019). <https://doi.org/10.1109/ICCV.2019.00444>
48. Reddy, N.D., Guigues, L., Pishchulin, L., Eledath, J., Narasimhan, S.G.: Tes-track: End-to-end learnable multi-person articulated 3D pose tracking. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 15190–15200 (2021). <https://doi.org/10.1109/CVPR46437.2021.01494>
49. Remelli, E., Han, S., Honari, S., Fua, P., Wang, R.: Lightweight multi-view 3D pose estimation through camera-disentangled representation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6040–6049 (2020). <https://doi.org/10.1109/CVPR42600.2020.00608>
50. Rempe, D., Birdal, T., Hertzmann, A., Yang, J., Sridhar, S., Guibas, L.J.: Humor: 3D human motion model for robust pose estimation. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 11488–11499 (2021). <https://doi.org/10.1109/ICCV48922.2021.01129>
51. Ren, J., Pan, L., Liu, Z.: Benchmarking and analyzing point cloud classification under corruptions. In: International Conference on Machine Learning (ICML). pp. 18559–18575. PMLR (2022), <https://proceedings.mlr.press/v162/ren22c.html>
52. Rezende, D., Mohamed, S.: Variational inference with normalizing flows. In: International Conference on Machine Learning (ICML). pp. 1530–1538. PMLR (2015), <https://proceedings.mlr.press/v37/rezende15.html>

53. Rhodin, H., Spörri, J., Katircioglu, I., Constantin, V., Meyer, F., Müller, E., Salzmann, M., Fua, P.: Learning monocular 3D human pose estimation from multi-view images. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8437–8446 (2018). <https://doi.org/10.1109/CVPR.2018.00880>
54. Shen, Z., Cen, Z., Peng, S., Shuai, Q., Bao, H., Zhou, X.: Learning human mesh recovery in 3D scenes. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 17038–17047 (2023). <https://doi.org/10.1109/CVPR52729.2023.01634>
55. Shuai, H., Wu, L., Liu, Q.: Adaptive multi-view and temporal fusing transformer for 3D human pose estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(4), 4122–4135 (2022). <https://doi.org/10.1109/TPAMI.2022.3188716>
56. Sun, J.J., Karashchuk, L., Dravid, A., Ryou, S., Fereidooni, S., Tuthill, J., Katsaggelos, A., Brunton, B.W., Gkioxari, G., Kennedy, A., Yue, Y., Perona, P.: BKinD-3D: self-supervised 3D keypoint discovery from multi-view videos. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9001–9010 (2023). <https://doi.org/10.1109/CVPR52729.2023.00869>
57. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5693–5703 (2019). <https://doi.org/10.1109/CVPR.2019.00584>
58. Tang, Z., Qiu, Z., Hao, Y., Hong, R., Yao, T.: 3D human pose estimation with spatio-temporal criss-cross attention. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4790–4799 (2023). <https://doi.org/10.1109/CVPR52729.2023.00464>
59. Tripathi, S., Müller, L., Huang, C.H.P., Taheri, O., Black, M.J., Tzionas, D.: 3D human pose estimation via intuitive physics. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4713–4725 (2023). <https://doi.org/10.1109/CVPR52729.2023.00457>
60. Usman, B., Tagliasacchi, A., Saenko, K., Sud, A.: Metapose: Fast 3D pose from multiple views without 3D supervision. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6759–6770 (2022). <https://doi.org/10.1109/CVPR52688.2022.00664>
61. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)* **30** (2017). <https://doi.org/10.48550/arXiv.1706.03762>
62. Wandt, B., Rudolph, M., Zell, P., Rhodin, H., Rosenhahn, B.: Canonpose: Self-supervised monocular 3D human pose estimation in the wild. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13294–13304 (2021). <https://doi.org/10.1109/CVPR46437.2021.01309>
63. Yao, Y., Jafarian, Y., Park, H.S.: Monet: Multiview semi-supervised keypoint detection via epipolar divergence. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 753–762 (2019). <https://doi.org/10.1109/ICCV.2019.00084>
64. Yoon, J.S., Yu, Z., Park, J., Park, H.S.: Humbi: A large multiview dataset of human body expressions and benchmark challenge. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(1), 623–640 (2021). <https://doi.org/10.1109/TPAMI.2021.3138762>

65. Yu, Z., Yoon, J.S., Lee, I.K., Venkatesh, P., Park, J., Yu, J., Park, H.S.: Humbi: A large multiview dataset of human body expressions. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2990–3000 (2020). <https://doi.org/10.1109/CVPR42600.2020.00306>
66. Zhang, Z., Wang, C., Qiu, W., Qin, W., Zeng, W.: Adafuse: Adaptive multiview fusion for accurate human pose estimation in the wild. *International Journal of Computer Vision* **129**, 703–718 (2021). <https://doi.org/10.1007/s11263-020-01398-9>
67. Zhao, J., Yu, T., An, L., Huang, Y., Deng, F., Dai, Q.: Triangulation residual loss for data-efficient 3D pose estimation. *Advances in Neural Information Processing Systems (NeurIPS)* **36** (2024), <https://openreview.net/forum?id=gLwjBDsE3G>
68. Zheng, C., Zhu, S., Mendieta, M., Yang, T., Chen, C., Ding, Z.: 3D human pose estimation with spatial and temporal transformers. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 11656–11665 (2021). <https://doi.org/10.1109/ICCV48922.2021.01145>
69. Zhou, X., Huang, Q., Sun, X., Xue, X., Wei, Y.: Towards 3D human pose estimation in the wild: a weakly-supervised approach. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 398–407 (2017). <https://doi.org/10.1109/ICCV.2017.51>
70. Zhu, W., Ma, X., Liu, Z., Liu, L., Wu, W., Wang, Y.: Motionbert: A unified perspective on learning human motion representations. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 15085–15099. IEEE (2023). <https://doi.org/10.1109/ICCV51070.2023.01385>
71. Zhu, Y., Xu, X., Shen, F., Ji, Y., Gao, L., Shen, H.T.: Posegtac: Graph transformer encoder-decoder with atrous convolution for 3D human pose estimation. In: International Joint Conference on Artificial Intelligence (IJCAI). pp. 1359–1365 (2021). <https://doi.org/10.24963/ijcai.2021/188>