

---

# Direct Multi-view Multi-person 3D Pose Estimation

---

Tao Wang<sup>1,2\*</sup>, Jianfeng Zhang<sup>2\*</sup>, Yujun Cai<sup>1</sup>, Shuicheng Yan<sup>1</sup>, Jiashi Feng<sup>1</sup>,  
<sup>1</sup>Sea AI Lab <sup>2</sup>National University of Singapore,  
twangnh@gmail.com,  
zhangjianfeng@u.nus.edu,  
{caiyj,yansc,fengjs}@sea.com

## Abstract

We present **Multi-view Pose transformer (MvP)** for estimating multi-person 3D poses from multi-view images. Instead of estimating 3D joint locations from costly volumetric representation or reconstructing the per-person 3D pose from multiple detected 2D poses as in previous methods, MvP directly regresses the multi-person 3D poses in a clean and efficient way, without relying on intermediate tasks. Specifically, MvP represents skeleton joints as learnable query embeddings and let them progressively attend to and reason over the multi-view information from the input images to directly regress the actual 3D joint locations. To improve the accuracy of such a simple pipeline, MvP presents a hierarchical scheme to concisely represent query embeddings of multi-person skeleton joints and introduces an input-dependent query adaptation approach. Further, MvP designs a novel geometrically guided attention mechanism, called *projective attention*, to more precisely fuse the cross-view information for each joint. MvP also introduces a RayConv operation to integrate the view-dependent camera geometry into the feature representations for augmenting the projective attention. We show experimentally that our MvP model outperforms the state-of-the-art methods on several benchmarks while being much more efficient. Notably, it achieves 92.3% AP<sub>25</sub> on the challenging Panoptic dataset, improving upon the previous best approach [40] by 9.8%. MvP is general and also extendable to recovering human mesh represented by the SMPL model, thus useful for modeling multi-person body shapes. Code and models are available at <https://github.com/sail-sg/mvp>.

## 1 Introduction

Multi-view multi-person 3D pose estimation aims to localize 3D skeleton joints for each person instance in a scene from multi-view camera inputs. It is a fundamental task that benefits many real-world applications (such as surveillance, sportscast, gaming and mixed reality) and is mainly tackled by reconstruction-based [6, 14, 4] and volumetric [40] approaches in previous literature, as shown in Fig. 1 (a) and (b). The former first estimates 2D poses in each view independently and then aggregates them and reconstructs their 3D counterparts via triangulation or a 3D pictorial structure model. The volumetric approach [40] builds a 3D feature volume through heatmap estimation and 2D-to-3D un-projection at first, based on which instance localization and 3D pose estimation are performed for each person instance individually. Though with notable accuracy, the above paradigms are inefficient due to highly relying on those intermediate tasks. Moreover, they estimate 3D pose for each person separately, making the computation cost grow linearly with the number of persons.

Targeted at a more simplified and efficient pipeline, we were wondering if it is possible to *directly* regress 3D poses from multi-view images without relying on any intermediate task? Though conceptually attractive, adopting such a direct mapping paradigm is highly non-trivial as it remains

---

\*Equal Contribution.

unclear how to perform skeleton joints detection and association for multiple persons within a single stage. In this work, we address these challenges by developing a novel **Multi-view Pose transformer (MvP)** model which significantly simplifies the multi-person 3D pose estimation. Specifically, MvP represents each skeleton joint as a learnable positional embedding, named *joint query*, which is fed into the model and mapped into final 3D pose estimation directly (Fig. 1 (c)), via a specifically designed attention mechanism to fuse multi-view information and globally reason over the joint predictions to assign them to the corresponding person instances. We develop a novel hierarchical query embedding scheme to represent the multi-person joint queries. It shares joint embedding across different persons and introduces person-level query embedding to help the model in learning both person-level and joint-level priors. Benefiting from exploiting the person-joint relation, the model can more accurately localize the 3D joints. Further, we propose to update the joint queries with input-dependent scene-level information (*i.e.*, globally pooled image features from multi-view inputs) such that the learnt joint queries can adapt to the target scene with better generalization performance.

To effectively fuse the multi-view information, we propose a geometrically-guided projective attention mechanism. Instead of applying full attention to densely aggregate features across spaces and views, it projects the estimated 3D joint into 2D anchor points for different views, and then selectively fuses the multi-view local features near to these anchors to precisely refine the 3D joint location. we propose to encode the camera rays into the multi-view feature representations via a novel RayConv operation to integrate multi-view positional information into the projective attention. In this way, the strong multi-view geometrical priors can be exploited by projective attention to obtain more accurate 3D pose estimation.

Comprehensive experiments on 3D pose benchmarks Panoptic [19], as well as Shelf and Campus [1] demonstrate our MvP works very well. Notably, it obtains 92.3% AP<sub>25</sub> on the challenging Panoptic dataset, improving upon the previous best approach VoxelPose [40] by 9.8%, while achieving nearly 2× speed up. Moreover, the design ethos of our MvP can be easily extended to more complex tasks—we show that a simple body mesh branch with SMPL representation [28] trained on top of a pre-trained MvP can achieve competitively qualitative results.

Our contributions are summarized as follows: 1) We strive for simplicity in addressing the challenging multi-view multi-person 3D pose estimation problem by casting it as a **direct regression problem** and accordingly develop a novel Multi-view Pose transformer (MvP) model, which achieves state-of-the-art results on the challenging Panoptic benchmark. 2) Different from query embedding designs in most transformer models, we propose a more tailored and concise hierarchical joint query embedding scheme to enable the model to effectively encode person-joint relation. Additionally, we mitigate the commonly faced generalization issue by a simple query adaptation strategy. 3) We propose a novel projective attention module along with a RayConv operation for fusing multi-view information effectively, which we believe are also inspiring for model designs in other multi-view 3D tasks.

## 2 Related Works

**3D Human Pose Estimation** 3D pose estimation from monocular inputs [29, 30, 49, 35, 38, 31, 46, 10, 47] is an ill-posed problem as multiple 3D predictions may result in the same 2D projection. To alleviate such projective ambiguities, multi-view methods have been explored. Research works on single-person scenes use either multi-view geometry [11] for feature fusion [36, 13] and triangulation [16, 37], or pictorial structure models for fast and robust 3D pose reconstruction [34, 36], achieving promising results. However, it is more challenging as we progress towards multi-person scenes. Current approaches mainly exploit a multi-stage pipeline for multi-person tasks, including reconstruction-based [6, 4, 14, 21, 26] and volumetric [40] paradigms. Despite their notable accuracy, these methods suffer expensive computation cost from the intermediate tasks, such as cross-view matching and heatmap back-projection. Moreover, the total computation cost grows linearly with the number of persons in the scene, making them hardly scalable for larger scenes. Different from all previous approaches that rely on a multi-stage pipeline with computation redundancy, our method views multi-person 3D pose estimation as a **direct regression problem** based on a novel Multi-view Pose transformer model, enables an intermediate task-free single stage solution.

**Attention and Transformers** Driven by the recent success in natural language fields, there have been growing interests in exploring the Transformers for computer vision tasks, such as image recognition [8] and generation [18], as well as more complicated object detection [3, 51] and video

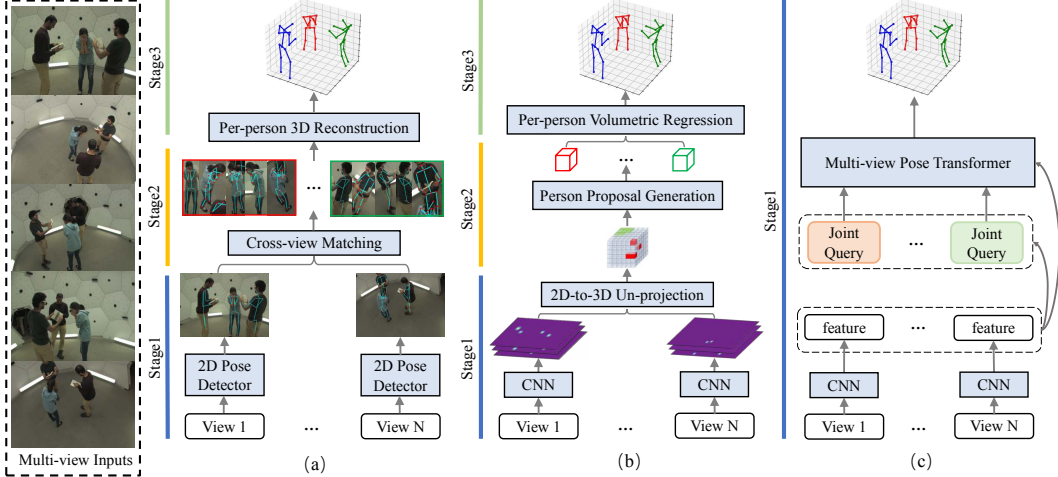


Figure 1: Difference between our method and others for multi-view multi-person 3D pose estimation. Existing methods adopt complex multi-stage pipelines that are either (a) reconstruction-based or (b) volumetric representation based, which incur heavy computation burden. (c) Our method solves this task as a **direct regression problem** without relying on any intermediate task by a novel Multi-view Pose Transformer, and largely simplifies the pipeline and boosts the efficiency.

instance segmentation [42]. However, multi-person 3D pose estimation has not been explored along this direction. In this study, we propose a novel Multi-view Pose Transformer architecture with a joint query embedding scheme and a projective attention module to regress 3D skeleton joints from multi-view images directly, delivering a simplified and effective pipeline.

### 3 Multi-view Pose Transformer (MvP)

To build a direct multi-person 3D pose estimation framework from multi-view images, we introduce a novel Multi-view Pose transformer (MvP). MvP takes in the multi-view feature representations, and transforms them into groups of 3D joint locations directly (Fig. 2 (a)), delivering multi-person 3D pose results, with the following carefully designed query embedding and attention schemes for detecting and grouping the skeleton joints.

#### 3.1 Joint Query Embedding Scheme

Inspired by transformers [41], MvP represents each skeleton joint as a learnable positional embedding, which is fed into the transformer decoder and mapped into final 3D joint location by jointly attending to other joints and the multi-view information (Fig. 2 (a)). The learnt embeddings encode *a priori* knowledge about the skeleton joints and we name them as *joint queries*. MvP develops the following concise query embedding scheme.

**Hierarchical Query Embeddings** The most straightforward way for designing joint query embeddings is to maintain a learnable query vector for each joint per person. However, we empirically find this scheme does not work well, likely because such a naive strategy cannot share the joint-level knowledge between different persons.

To tackle this problem, we develop a hierarchical query embedding scheme to explicitly encode the person-joint relation for better generalization to different scenes. The hierarchical embedding offers joint-level information sharing across different persons and reduces the learnable parameters, helping the model to learn useful knowledge from the training data, and thus generalize better. Concretely, instead of using the set of independent joint queries  $\{\mathbf{q}_m\}_{m=1}^M \subset \mathbb{R}^C$ , we employ a set of person level queries  $\{\mathbf{h}_n\}_{n=1}^N \subset \mathbb{R}^C$ , and a set of joint level queries  $\{\mathbf{l}_j\}_{j=1}^J \subset \mathbb{R}^C$  to represent different persons and different skeleton joints, where  $C$  denotes the feature dimension,  $N$  is the number of persons,  $J$  is the number of joints per person, and  $M = NJ$ . Then the query of joint  $j$  of person  $n$

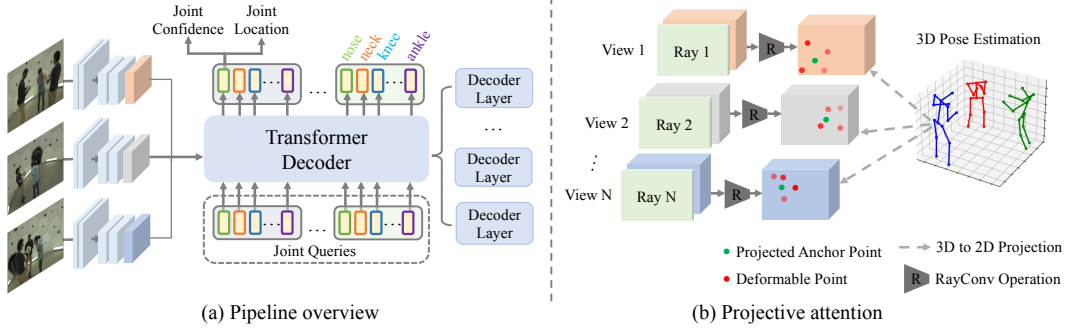


Figure 2: (a) Overview of the proposed MvP model. Upon the multi-view image features from several convolution layers, it deploys a transformer decoder with a stack of decoder layers to map the input joint queries and the multi-view features to 3D poses directly. (b) The projective attention of MvP projects 3D skeleton joints to anchor points (the green dots) on different views and samples deformable points (the red dots) surrounding these anchors to aggregate local contextual features via learned weights (the brighter color density means larger weights).

can be hierarchically formulated as

$$\mathbf{q}_n^j = \mathbf{h}_n + \mathbf{l}_j. \quad (1)$$

With such a hierarchical embedding scheme, the number of learnable query embedding parameters is reduced from  $NJC$  to  $(N + J)C$ .

**Input-dependent Query Adaptation** In the above, the learned joint query embeddings are shared for all the input images, independent of their contents, and thus may not generalize well on the novel target data. To address this limitation, we propose to augment the joint queries with input-dependent scene-level information in both model training and deployment, such that the learnt joint queries can be adaptive to the target data and generalize better. Concretely, we augment the above joint queries with a globally pooled feature vector  $\mathbf{g} \in \mathbb{R}^C$  from the multi-view image feature representations:

$$\mathbf{q}_n^j = \mathbf{g} + \mathbf{h}_n + \mathbf{l}_j. \quad (2)$$

Here  $\mathbf{g} = \text{Concat}(\text{Pool}(\mathbf{Z}_1), \dots, \text{Pool}(\mathbf{Z}_V))\mathbf{W}^g$ , where  $\mathbf{Z}_v$  denotes image feature from  $v$ -th view and  $V$  is the total number of camera views; Concat and Pool denote concatenation and pooling operations, and  $\mathbf{W}^g$  is a learnable linear weight.

### 3.2 Projective Attention for Multi-view Feature Fusion

It is crucial to aggregate complementary multi-view information to transform the joint embeddings into accurate 3D joint locations. We consider the dot product attention mechanism of transformers [41] to fuse the multi-view image features. However, naively applying such dot product attention densely over all spatial locations and camera views will incur enormous computation cost. Moreover, such dense attention is difficult to optimize and delivers poor performance empirically since it does not exploit any 3D geometric knowledge.

Therefore, we propose a geometrically-guided multi-view projective attention scheme, named projective attention. The core idea is to take the 2D projection of the estimated 3D joint location as the anchor point in each view, and only fuse the local features near those projected 2D locations from different views. Motivated by the deformable convolution [5, 50], we adopt an adaptive deformable sampling strategy to gather the localized context information in each camera view, as shown in Fig. 2 (b). Other local attention operations [48, 44, 43] can also be adopted as an alternative. Formally, given joint query feature  $\mathbf{q}$  and 3D joint position  $\mathbf{y}$ , the projective attention is defined as

$$\begin{aligned} \text{PAttention}(\mathbf{q}, \mathbf{y}, \{\mathbf{Z}_v\}_{v=1}^V) &= \text{Concat}(\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_V)\mathbf{W}^P, \\ \text{where } \mathbf{f}_v &= \sum_{k=1}^K \mathbf{a}(k) \cdot \mathbf{Z}_v(\Pi(\mathbf{y}, \mathbf{C}_v) + \Delta\mathbf{p}(k))\mathbf{W}^f. \end{aligned} \quad (3)$$

Here the view-specific feature  $\mathbf{f}_v$  is obtained by aggregating features from  $K$  discrete offsetted sampling points from an anchor point  $\mathbf{p} = \Pi(\mathbf{y}, \mathbf{C}_v)$ , located by projecting the current 3D joint

location  $\mathbf{y}$  to 2D, where  $\Pi : \mathbb{R}^3 \rightarrow \mathbb{R}^2$  denotes perspective projection [11] and  $\mathbf{C}_v$  the corresponding camera parameters.  $\mathbf{W}^P$  and  $\mathbf{W}^f$  are learnable linear weights. The attention weight  $\mathbf{a}$  and the offset to the projected anchor point  $\Delta\mathbf{p}$  are estimated from the fusion of query feature  $\mathbf{q}$  and the view-dependent feature at the projected anchor point  $\mathbf{Z}_v(\mathbf{p})$ , *i.e.*,  $\mathbf{a} = \text{Softmax}((\mathbf{q} + \mathbf{Z}_v(\mathbf{p}))\mathbf{W}^a)$  and  $\Delta\mathbf{p} = (\mathbf{q} + \mathbf{Z}_v(\mathbf{p}))\mathbf{W}^p$ , where  $\mathbf{W}^a$  and  $\mathbf{W}^p$  are learnable linear weights. If the projected location and the offset are fractional, we use bilinear interpolation to obtain the corresponding feature  $\mathbf{Z}_v(\mathbf{p})$  or  $\mathbf{Z}_v(\mathbf{p} + \Delta\mathbf{p}(t))$ .

The projective attention incorporates two geometrical cues, *i.e.*, the corresponding 2D spatial locations across views from the 3D to 2D projection and the deformed neighborhood of the anchors from the learned offsets to gather view-adaptive contextual information. Unlike naive attention where the query feature densely interacts with the multi-view key features across all the spatial locations, the projective attention is more selective for the interaction between the query and each view—only the features from locations near to the projected anchors are aggregated, and thus is much more efficient.

**Encoding Multi-view Positional Information with RayConv** The positional encoding [41] is an important component of the transformer, which provides positional information of the input sequence. However, a simple per-view 2D positional encoding scheme cannot encode the multi-view geometrical information. To tackle this limitation, we propose to encode the camera ray directions that represent positional information in 3D space into the multi-view feature representations. Concretely, the camera ray direction  $\mathbf{R}_v$ , generated with the view-specific camera parameters, is concatenated channel-wisely to the corresponding image feature representation  $\mathbf{Z}_v$ . Then a standard convolution is applied to obtain the updated feature representation  $\hat{\mathbf{Z}}_v$ , with the view-dependent geometric information:

$$\hat{\mathbf{Z}}_v = \text{Conv}(\text{Concat}(\mathbf{Z}_v, \mathbf{R}_v)). \quad (4)$$

We name the operation as *RayConv*. With it, the obtained feature representation  $\hat{\mathbf{Z}}_v$  is used for the projective attention by replacing  $\mathbf{Z}_v$  in Eqn. (3).

Such drop-in replacement introduces negligible computation, while injecting strong multi-view geometrical prior to augment the projective attention scheme, thus helping more precisely predict the refined 3D joint position.

### 3.3 Architecture

Our overall architecture (Fig. 2 (a)) is pleasantly simple. It adopts a convolution neural network, designed for 2D pose estimation [45], to obtain high-resolution image features  $\{\mathbf{Z}_v\}_{v=1}^V$  from multi-view inputs  $\{\mathbf{I}_v\}_{v=1}^V$ . The features are then fed into the transformer decoder consisting of multiple decoder layers to predict the 3D joint locations. Each layer conducts a self-attention to perform pair-wise interaction between all the joints from all the persons in the scene; a projective attention to selectively gather the complementary multi-view information; and a feed-forward regression to predict the 3D joint positions and their confidence scores. Specifically, the transformer decoder applies a *multi-layer progressive regression scheme*, *i.e.*, each decoder layer outputs 3D joint offsets to refine the input 3D joint positions from previous layer.

**Extending to Body Mesh Recovery** MvP learns skeleton joints feature representations and is extendable to recovering human mesh with a parametric body mesh model [28]. Specifically, after average pooling on the joint features into per-person feature, a feed-forward network is used to predict the corresponding body mesh represented by the parametric SMPL model [28]. Similar to the joint location prediction, the SMPL parameters follow multi-layer progressive regression scheme.

### 3.4 Training

MvP infers a fixed set of  $M$  joint locations for  $N$  different persons, where  $M = NJ$ . The main training challenge is how to associate the skeleton joints correctly for different person instances. Unlike the post-hoc grouping of detected skeleton joints as in bottom-up pose estimation methods [32, 24], MvP learns to directly predict the multi-joint 3D human pose in a group-wise fashion as shown in Fig. 2 (a). This is achieved by a grouped matching strategy during model training.

**Grouped Matching** Given the predicted joint positions  $\{\mathbf{y}_m\}_{m=1}^M \subset \mathbb{R}^3$  and associated confidence scores  $\{s_m\}_{m=1}^M$ , we group every consecutive  $J$ -joint predictions into per-person pose estimation

$\{\mathbf{Y}_n\}_{n=1}^N \subset \mathbb{R}^{J \times 3}$ , and average their corresponding confidence scores to obtain the per-person confidence scores  $\{p_n\}_{n=1}^N$ . The same grouping strategy is used during inference.

The ground truth set  $\mathbf{Y}^*$  of 3D poses of different person instances is smaller than the prediction set of size  $N$ , which is padded to size  $N$  with empty element  $\emptyset$ . Then we find a bipartite matching between the prediction set and the ground truth set by searching for a permutation of  $\hat{\sigma} \in \aleph_N$  that achieves the lowest matching cost:

$$\hat{\sigma} = \arg \min_{\sigma \in \aleph_N} \sum_{n=1}^N \mathcal{L}_{\text{match}}(\mathbf{Y}_n^*, \mathbf{Y}_{\sigma(n)}). \quad (5)$$

We consider both the regressed 3D joint position and confidence score for the matching cost:

$$\mathcal{L}_{\text{match}}(\mathbf{Y}_n^*, \mathbf{Y}_{\sigma(n)}) = -p_i + \mathcal{L}_1(\mathbf{Y}_n^*, \mathbf{Y}_{\sigma(n)}) \quad (6)$$

where  $\mathbf{Y}_n^* \neq \emptyset$ , and  $\mathcal{L}_1$  computes the  $L_1$  loss error. Following [3, 39], we employ the Hungarian algorithm [25] to compute the optimal assignment  $\hat{\sigma}$  with the above matching cost.

**Objective Function** We compute the *Hungarian loss* with the obtained optimal assignment  $\hat{\sigma}$ :

$$\mathcal{L}_{\text{Hungarian}}(\mathbf{Y}^*, \mathbf{Y}) = \sum_{n=1}^N [\mathcal{L}_{\text{conf}}(\mathbf{Y}_n^*, p_{\hat{\sigma}(n)}) + \mathbb{1}_{\{\mathbf{Y}_n^* \neq \emptyset\}} \lambda \mathcal{L}_{\text{pose}}(\mathbf{Y}_n^*, \mathbf{Y}_{\hat{\sigma}(n)})]. \quad (7)$$

Here  $\mathcal{L}_{\text{conf}}$  and  $\mathcal{L}_{\text{pose}}$  are losses for confidence score and pose regression, respectively.  $\lambda$  balances the two loss terms. We use focal loss [27] for confidence prediction which adaptively balances the positive and negative samples. For pose regression, we compute  $L_1$  loss for 3D joints and their projected 2D joints in different views.

To learn multi-layer progressive regression, the above matching and loss are applied for each decoder layer. The total loss is thus  $\mathcal{L}_{\text{total}} = \sum_{l=1}^L \mathcal{L}_{\text{Hungarian}}^l$ , where  $\mathcal{L}_{\text{Hungarian}}^l$  denotes loss of the  $l$ -th decoder layer and  $L$  is the number of decoder layers. When extending MvP to body mesh recovery, we apply  $L_1$  loss for 3D joints from the SMPL model and their 2D projections, as well as an adversarial loss following HMR [22, 17, 47] due to lack of GT SMPL parameters.

## 4 Experiments

In this section, we aim to answer following questions. 1) Can MvP provide both efficient and accurate multi-person 3D pose estimation? 2) How does the proposed attention mechanism help multi-view multi-person skeleton joints information fusing? 3) How does each individual design choice affect model performance? To this end, we conduct extensive experiments on several benchmark datasets.

**Datasets** *Panoptic* [20] is a large-scale benchmark with 3D skeleton joint annotations. It captures daily social activities in an indoor environment. We conduct extensive experiments on Panoptic to evaluate and analyze our approach. Following VoxelPose [40], we use the same data sequences except ‘160906\_band3’ in the training set due to broken images. Unless otherwise stated, we use five HD cameras (3, 6, 12, 13, 23) in our experiments. All results reported in the experiments follow the same data setup. We use Average Precision (AP) and Recall [40], as well as Mean Per Joint Position Error (MPJPE) as evaluation metrics. *Shelf* and *Campus* [1] are two multi-person datasets capturing indoor and outdoor environments, respectively. We split them into training and testing sets following [1, 6, 40]. We report Percentage of Correct Parts (PCP) for these two datasets.

**Implementation Details** Following VoxelPose [40], we adopt a pose estimation model [45] build upon ResNet-50 [12] for multi-view image features extraction. Unless otherwise stated, we use a

Table 1: Result on the Panoptic dataset. MvP is more accurate and faster than VoxelPose.

Methods	AP <sub>25</sub>	AP <sub>50</sub>	AP <sub>100</sub>	AP <sub>150</sub>	Recall <sub>@500</sub>	MPJPE[mm]	Time[ms]
VoxelPose [40]	84.0	96.4	<b>97.5</b>	<b>97.8</b>	98.1	17.8	320
MvP (Ours)	<b>92.3</b>	<b>96.6</b>	<b>97.5</b>	97.7	<b>98.2</b>	<b>15.8</b>	<b>170</b>



stack of six transformer decoder layers. The model is trained for 40 epochs, with the Adam optimizer of learning rate  $10^{-4}$ . During inference, a confidence threshold of 0.1 is used to filter out redundant predictions. Please refer to supplementary for more implementation details.

#### 4.1 Main Results

**Panoptic** We first evaluate our MvP model on the challenging Panoptic dataset and compare it with the state-of-the-art VoxelPose model [40]. As shown in Table 1, Our MvP achieves 92.3 AP<sub>25</sub>, improving upon VoxelPose by 9.8%, and achieves much lower MPJPE (15.8 v.s 17.8). Moreover, MvP only requires 170ms to process a multi-view input, about 2× faster than VoxelPose<sup>2</sup>. These results demonstrate both accuracy and efficiency advantages of MvP from estimating 3D poses of multiple persons in a direct regression paradigm. To further demonstrate efficiency of MvP, we compare its inference time with VoxelPose’s when processing different numbers of person instances. As shown in Fig. 3, the inference time of VoxelPose grows linearly with the number of persons in the scene due to the per-person regression paradigm. In contrast, MvP keeps constant inference time no matter how many instances in the scene. Notably, it takes only 185ms for MvP to process scenes even with 100 person instances (the blue line), demonstrating its great potential to handle crowded scenarios.

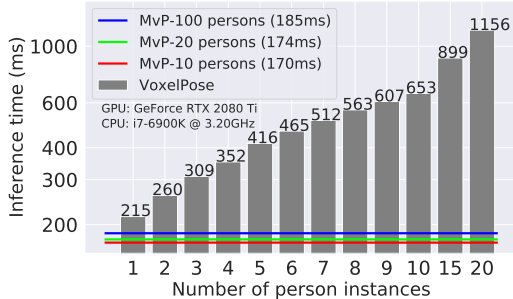


Figure 3: Inference time versus the number of person instances. Benefiting from its direct inference framework, MvP maintains almost constant inference time regardless of the number of persons.

**Shelf and Campus** We further compare our MvP with state-of-the-art approaches on the Shelf and Campus datasets. The reconstruction-based methods [2, 9, 6] use 3D pictorial model [2, 6] or conditional random field [9] within a multi-stage paradigm; and the volumetric approach VoxelPose [40] highly relies on computationally intensive intermediate tasks. As shown in Table 2, our MvP achieves the best performance in all the actors on the Shelf dataset. Moreover, it obtains a comparable result on the Campus dataset as VoxelPose [40] without relying on any intermediate task. These results further confirm the effectiveness of MvP for estimating 3D poses of multiple persons directly.

Table 2: Results (in PCP) on Shelf and Campus datasets.

Methods	Shelf				Campus			
	Actor 1	Actor 2	Actor 3	Average	Actor 1	Actor 2	Actor 3	Average
Belagiannis <i>et al.</i> [2]	75.3	69.7	87.6	77.5	93.5	75.7	84.4	84.5
Ershadi <i>et al.</i> [9]	93.3	75.9	94.8	88.0	94.2	92.9	84.6	90.6
Dong <i>et al.</i> [6]	98.8	94.1	<b>97.8</b>	96.9	97.6	93.3	98.0	96.3
VoxelPose [40]	<b>99.3</b>	94.1	97.6	97.0	97.6	93.8	<b>98.8</b>	<b>96.7</b>
MvP (Ours)	<b>99.3</b>	<b>95.1</b>	<b>97.8</b>	<b>97.4</b>	<b>98.2</b>	<b>94.1</b>	97.4	96.6

#### 4.2 Visualization

**3D Pose and Body Mesh Estimation** We visualize some 3D pose estimations of MvP on the challenging Panoptic dataset in Fig. 4. It can be observed that MvP is robust to large pose deformation (the 1st example) and severe occlusion (the 2nd example), and can achieve geometrically plausible results w.r.t. different viewpoints (the rightmost column). Moreover, MvP is extendable to body mesh recovery and can achieve fairly good reconstruction results (the 2nd and 4th rows). All these results verify both effectiveness and extendability of MvP. Please see supplementary for more examples.

<sup>2</sup>We count averaged per-sample inference time in millisecond on Panoptic test set. For all methods, the time is counted on GPU GeForce RTX 2080 Ti and CPU Intel i7-6900K @ 3.20GHz.



Figure 4: Example 3D pose estimations from Panoptic dataset. The left four columns show the multi-view inputs and the corresponding body mesh estimations. The rightmost column shows the estimated 3D poses from two different viewpoints. Best viewed in color.

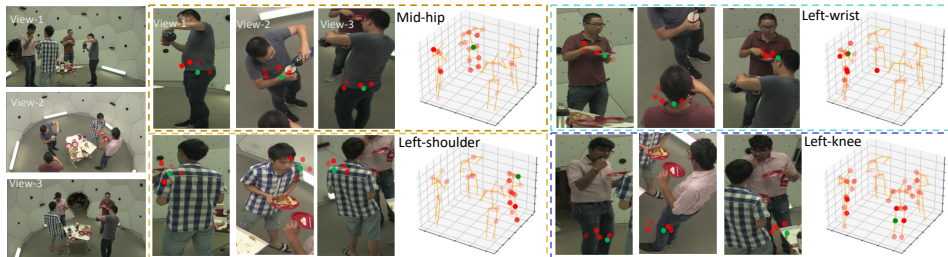


Figure 5: Visualization of projective attention and self-attention on example skeleton joints. The attention weights are obtained with the 4-th decoder layer of a trained model. *Projective attention* (in the cropped image triplets): the green points denote the projected anchor points in each camera view, and the red points denote the offsetted spatial locations, with brighter color for stronger attention. *Self-attention* (in the 3D skeleton plots): example skeleton joint (green) to all the other skeleton joints (red) in the scene. The color density indicates attention weight. Best viewed in color and  $2\times$  zoom.

**Attention Mechanism** We visualize the projective attention and the self-attention in Fig. 5. Benefiting from the 3D-to-2D projection, the projective attention can accurately locate the skeleton joint in each camera view (the green point) based on the current estimated 3D joint location. We observe it learns to gather adaptive local context information (the red points) with the deformable sampling operation. For instance, when regressing the 3D position of mid-hip (the 1st example), the projective attention selectively attends to informative joints such as the left and right hips as well as thorax, which offers sufficient contextual information for accurate estimation. We also visualize the self-attention, which learns pair-wise interaction between all the skeleton joints in the scene. From the 3D plot in Fig. 5, we can observe a certain skeleton joint mainly attends to other joints of the same person instance (more opaque). It also attends to joints from other person instances, but with less attention (more transparent). This phenomenon is reasonable as the skeleton joints of a human body are strongly correlated to each other, *e.g.*, with certain pose priors and bone length.

### 4.3 Ablation

**Importance of RayConv** MvP introduces RayConv to encode multi-view geometric information, *i.e.*, camera ray directions into image feature representations. As shown in Table 3a, if removing RayConv, the performance drops significantly—4.8 decrease in  $AP_{25}$  and 1.6 increase in MPJPE. This indicates the multi-view geometrical information is important for the model to more precisely



Table 3: Ablations on Panoptic. In (b), *Hier.* denotes the hierarchical query embedding scheme, *Hier.+ad.* means further adding the adaptation strategy. Please see supplement for more ablations.

<i>RConv</i>	AP <sub>25</sub>	AP <sub>100</sub>	MPJPE
w/	92.3	97.5	15.8
w/o	87.5	96.2	17.4

(a) The effect of *RayConv.* w/o means removing RayConv.

<i>Query</i>	AP <sub>25</sub>	AP <sub>100</sub>	MPJPE
Per-joint	67.4	84.7	41.2
Hier.	82.5	93.2	19.5
Hier.+ad.	92.3	97.5	15.8

(b) Different joint query embedding schemes.

<i>Thr.</i>	AP <sub>25</sub>	AP <sub>100</sub>	MPJPE
0.0	93.1	98.5	16.3
0.1	92.3	97.5	15.8
0.2	91.1	96.2	15.5
0.4	89.2	93.7	15.0

(c) Different confidence threshold during evaluation.

<i>Dec.</i>	AP <sub>25</sub>	AP <sub>100</sub>	MPJPE
2	6.3	92.5	49.6
3	63.4	95.6	22.8
4	86.8	96.8	17.5
5	91.8	97.6	16.2
6	92.3	97.5	15.8
7	92.0	97.5	15.9

(d) Number of decoder layers.

<i>Cam.</i>	AP <sub>25</sub>	AP <sub>100</sub>	MPJPE
1	4.7	61.0	93.8
2	37.7	93.0	34.8
3	71.8	95.1	21.1
4	84.1	96.7	19.3
5	92.3	97.5	15.8

(e) Number of camera views.

<i>K</i>	AP <sub>25</sub>	AP <sub>100</sub>	MPJPE
1	88.6	96.3	18.2
2	89.3	97.5	17.4
4	92.3	97.7	15.8
8	84.4	91.1	20.3

(f) Number of deformable points *K*.

localize the skeleton joints in 3D space. Without RayConv, the transformer decoder cannot accurately capture positional information in 3D space, resulting in performance drop.

**Importance of Hierarchical Query Embedding** As shown in Table 3b, compared with the straightforward and unstructured per-joint query embedding scheme, the proposed hierarchical query embedding boosts the performance sharply—14.1 increase in AP<sub>25</sub> and 23.4 decrease in MPJPE. Its advantageous performance clearly verifies introducing the person-level queries to collaborate with the joint-level queries can better exploit human body structural information and improve model to better localize the joints. Upon the hierarchical query embedding scheme, adding the query adaptation strategy further improves the performance significantly, reaching AP<sub>25</sub> of 92.3 and MPJPE of 15.8. This shows the proposed approach effectively adapts the query embeddings to the target scene and such adaptation is indeed beneficial for the generalization of MvP to novel scenes.

**Different Model Designs** We also examine effects of varying the following designs of the MvP model to gain better understanding on them.

**Confidence Threshold** During inference, a confidence threshold is used to filter out the low-confidence and erroneous pose predictions, and obtain the final result. Adopting a higher confidence will select the predictions in a more restrictive way. As shown in Table 3c, a higher confidence threshold brings lower MPJPE as it selects more accurate predictions; but it also filters out some true positive predictions and thus reduces the average precision.

**Number of Decoder Layers** Decoder layers are used for refining the pose estimation. Stacking more decoder layers thus gives better performance (Table 3d). For instance, the MPJPE is as high as 49.6 when using only two decoder layers, but it is significantly reduced to 22.8 when using three decoder layers. This clearly justifies the progressive refinement strategy of our MvP model is effective. However the benefit of using more decoder layers diminishes when the number of layers is large enough, implying the model has reached the ceiling of its model capacity.

**Number of Camera Views** Multi-view inputs provide complementary information to each other which is extremely useful when handling some challenging environment factors in 3D pose estimation like occlusions. We vary the number of camera views to examine whether MvP can effectively fuse and leverage multi-view information to continuously improve the pose estimation quality (Table 3e). As expected, with more camera views, the 3D pose estimation accuracy monotonically increases, demonstrating the capacity of MvP in fusing multi-view information.

**Number of Deformable Sampling Points** Table 3f shows the effect of the number of deformable sampling points *K* used in the projective attention. With only one deformable point, MvP already achieves a respectable result, *i.e.*, 88.6 in AP<sub>25</sub> and 17.4 in MPJPE. Using more sampling points

further improves the performance, demonstrating the projective attention is effective at aggregating information from the useful locations. When  $K = 4$ , the model gives the best result. Further increasing  $K$  to 8, the performance starts to drop. It is likely because using too many deformable points introduces redundant information and thus makes the model more difficult to optimize.

## 5 Conclusion

We introduced a direct and efficient model, named Multi-view Pose transformer (MvP), to address the challenging multi-view multi-person 3D human pose estimation problem. Different from existing methods relying on tedious intermediate tasks, MvP substantially simplifies the pipeline into a direct regression one by carefully designing the transformer-alike model architecture with a novel hierarchical joint query embedding scheme and projective attention mechanism. We conducted extensive experiments to verify its superior performance and speed over the well-established baselines.

We empirically found MvP needs sufficient data for model training since it learns the 3D geometry implicitly. In the future, we will study how to enhance the data-efficiency of MvP by leveraging the strategy like self-supervised pre-training or exploring more advanced approaches. Similar to prior works, we also found MvP suffers from performance drop for cross-camera generalization, that is, generalizing on novel camera views. We will explore approaches like disentangling camera parameters and multi-view feature learning to improve this aspect. Besides, we will explore the large-scale applications of MvP and further extend it to other relevant tasks. Thanks to its efficiency, MvP would be scalable to handle very crowded scenes with many persons. Moreover, the framework of MvP is general and thus extensible to other 3D modeling tasks like dense mesh recovery of common objects.

## References

- [1] Vasileios Belagiannis, Sikandar Amin, Mykhaylo Andriluka, Bernt Schiele, Nassir Navab, and Slobodan Ilic. 3d pictorial structures for multiple human pose estimation. In *CVPR*, 2014.
- [2] Vasileios Belagiannis, Sikandar Amin, Mykhaylo Andriluka, Bernt Schiele, Nassir Navab, and Slobodan Ilic. 3d pictorial structures revisited: Multiple human pose estimation. *IEEE transactions on pattern analysis and machine intelligence*, 38(10):1929–1942, 2015.
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020.
- [4] He Chen, Pengfei Guo, Pengfei Li, Gim Hee Lee, and Gregory Chirikjian. Multi-person 3d pose estimation in crowded scenes based on multi-view geometry. In *ECCV*, 2020.
- [5] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *ICCV*, 2017.
- [6] Junting Dong, Wen Jiang, Qixing Huang, Hujun Bao, and Xiaowei Zhou. Fast and robust multi-person 3d pose estimation from multiple views. In *CVPR*, 2019.
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv*, 2020.
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv*, 2020.
- [9] Sara Ershadi-Nasab, Erfan Noury, Shohreh Kasaei, and Esmaeil Sanaei. Multiple human 3d pose estimation from multiview images. *Multimedia Tools and Applications*, 77(12):15573–15601, 2018.
- [10] Kehong Gong, Jianfeng Zhang, and Jiashi Feng. Poseaug: A differentiable pose augmentation framework for 3d human pose estimation. In *CVPR*, 2021.
- [11] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, New York, NY, USA, 2 edition, 2003.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [13] Yihui He, Rui Yan, Katerina Fragkiadaki, and Shoubo Yu. Epipolar transformers. In *CVPR*, 2020.
- [14] Congzhen Huang, Shuai Jiang, Yang Li, Ziyue Zhang, Jason Traish, Chen Deng, Sam Ferguson, and Richard Yi Da Xu. End-to-end dynamic matching network for multi-view multi-person 3d pose estimation. In *ECCV*, 2020.

- [15] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014.
- [16] Karim Isakov, Egor Burkov, Victor Lempitsky, and Yury Malkov. Learnable triangulation of human pose. In *ICCV*, 2019.
- [17] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *CVPR*, 2020.
- [18] Yifan Jiang, Shiyu Chang, and Zhangyang Wang. Transgan: Two transformers can make one strong gan. *arXiv*, 2021.
- [19] Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *ICCV*, 2015.
- [20] Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, Sean Banerjee, Timothy Godisart, Bart Nabbe, Iain Matthews, et al. Panoptic studio: A massively multiview system for social interaction capture. *IEEE transactions on pattern analysis and machine intelligence*, 41(1):190–204, 2017.
- [21] Abdolrahim Kadkhodamohammadi and Nicolas Padoy. A generalizable approach for multi-view 3d human pose regression. *Machine Vision and Applications*, 32(1):1–14, 2021.
- [22] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018.
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICCV*, 2015.
- [24] Sven Kreiss, Lorenzo Bertoni, and Alexandre Alahi. Pifpaf: Composite fields for human pose estimation. In *CVPR*, 2019.
- [25] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [26] Jiahao Lin and Gim Hee Lee. Multi-view multi-person 3d pose estimation with plane sweep stereo. In *CVPR*, 2021.
- [27] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017.
- [28] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015.
- [29] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *ICCV*, 2017.
- [30] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Trans. on Graphics*, 36(4):44, 2017.
- [31] Xuecheng Nie, Jianfeng Zhang, Shuicheng Yan, and Jiashi Feng. Single-stage multi-person pose machines. In *ICCV*, 2019.
- [32] George Papandreou, Tyler Zhu, Liang-Chieh Chen, Spyros Gidaris, Jonathan Tompson, and Kevin Murphy. Personlab: Person pose estimation and instance segmentation with a bottom-up, part-based, geometric embedding model. In *ECCV*, 2018.
- [33] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NeurIPS*, 2017.
- [34] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Harvesting multiple views for marker-less 3d human pose annotations. In *CVPR*, 2017.
- [35] Alin-Ionut Popa, Mihai Zanfir, and Cristian Sminchisescu. Deep multitask architecture for integrated 2d and 3d human sensing. In *CVPR*, 2017.
- [36] Haibo Qiu, Chunyu Wang, Jingdong Wang, Naiyan Wang, and Wenjun Zeng. Cross view fusion for 3d human pose estimation. In *ICCV*, 2019.
- [37] Edoardo Remelli, Shangchen Han, Sina Honari, Pascal Fua, and Robert Wang. Lightweight multi-view 3d pose estimation through camera-disentangled representation. In *CVPR*, 2020.
- [38] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *ECCV*, 2018.
- [39] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. *arXiv*, 2014.

- [40] Hanyue Tu, Chunyu Wang, and Wenjun Zeng. Voxelpose: Towards multi-camera 3d human pose estimation in wild environment. In *ECCV*, 2020.
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv*, 2017.
- [42] Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. End-to-end video instance segmentation with transformers. *arXiv*, 2020.
- [43] Felix Wu, Angela Fan, Alexei Baevski, Yann N Dauphin, and Michael Auli. Pay less attention with lightweight and dynamic convolutions. *arXiv*, 2019.
- [44] Zhanghao Wu, Zhijian Liu, Ji Lin, Yujun Lin, and Song Han. Lite transformer with long-short range attention. *arXiv*, 2020.
- [45] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *ECCV*, 2018.
- [46] Jianfeng Zhang, Xuecheng Nie, and Jiashi Feng. Inference stage optimization for cross-scenario 3d human pose estimation. In *NeurIPS*, 2020.
- [47] Jianfeng Zhang, Dongdong Yu, Jun Hao Liew, Xuecheng Nie, and Jiashi Feng. Body meshes as points. In *CVPR*, 2021.
- [48] Hengshuang Zhao, Jiaya Jia, and Vladlen Koltun. Exploring self-attention for image recognition. In *CVPR*, 2020.
- [49] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *ICCV*, 2017.
- [50] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *CVPR*, 2019.
- [51] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaoang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICCV*, 2020.