

Lifting by Image - Leveraging Image Cues for Accurate 3D Human Pose Estimation

Feng Zhou, Jianqin Yin*, Peiyang Li

School of Artificial Intelligence, Beijing University of Posts and Telecommunications, China
{zhoufeng, jqyin, lipeiyang}@bupt.edu.cn

Abstract

The “lifting from 2D pose” method has been the dominant approach to 3D Human Pose Estimation (3DHPE) due to the powerful visual analysis ability of 2D pose estimators. Widely known, there exists a depth ambiguity problem when estimating solely from 2D pose, where one 2D pose can be mapped to multiple 3D poses. Intuitively, the rich semantic and texture information in images can contribute to a more accurate “lifting” procedure. Yet, existing research encounters two primary challenges. Firstly, the distribution of image data in 3D motion capture datasets is too narrow because of the laboratorial environment, which leads to poor generalization ability of methods trained with image information. Secondly, effective strategies for leveraging image information are lacking. In this paper, we give new insight into the cause of poor generalization problems and the effectiveness of image features. Based on that, we propose an advanced framework. Specifically, the framework consists of two stages. First, we enable the keypoints to query and select the beneficial features from all image patches. To reduce the keypoints attention to inconsequential background features, we design a novel Pose-guided Transformer Layer, which adaptively limits the updates to unimportant image patches. Then, through a designed Adaptive Feature Selection Module, we prune less significant image patches from the feature map. In the second stage, we allow the keypoints to further emphasize the retained critical image features. This progressive learning approach prevents further training on insignificant image features. Experimental results show that our model achieves state-of-the-art performance on both the Human3.6M dataset and the MPI-INF-3DHP dataset.

Introduction

Monocular 3D Human Pose Estimation (3DHPE) aims to estimate the relative 3D coordinates of human joints from an image. It is a fundamental computer vision task related to a wide range of applications, including human motion forecasting (Ding and Yin 2022; Liu et al. 2020), human action recognition (Dang, Yang, and Yin 2020), human-centric generation (Cao et al. 2023b, 2022, 2023a) and so on.

In recent years, 3D human pose estimation has been dominated by the “lifting” technique (Martinez et al. 2017). This

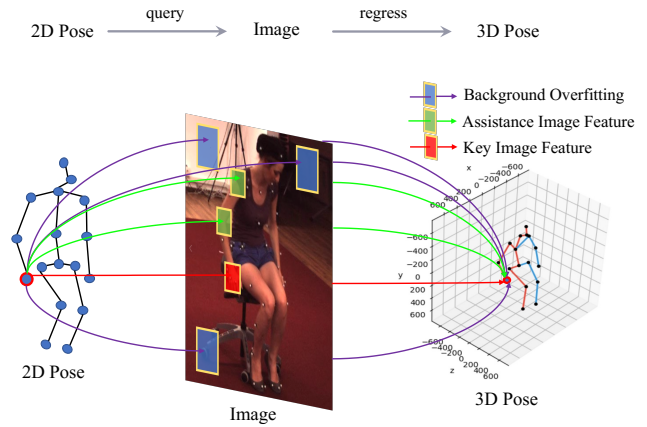


Figure 1: The main idea of this paper is to design a framework that enables 2D poses to regress 3D poses by querying information from the image. The framework is specifically designed based on two key insights: First, excessive attention to dataset-biased background information leads to poor generalization ability. Second, it is not only the image features corresponding to the keypoints that are helpful for the task, but also the associated body structural positions of the keypoints that can provide valuable assistance.

approach consists of two stages. First, utilize off-the-shelf 2D pose estimators (Sun et al. 2019; Newell, Yang, and Deng 2016; Dang et al. 2022) to estimate the 2D pose from the image and then regress the 3D pose from the obtained 2D human pose. Compared to direct estimation, this cascaded approach has the following advantages: 2D estimator is trained on more diverse and extensive 2D human pose datasets, which enables stronger visual perception and generalization ability (Martinez et al. 2017). Besides, the “lifting” can be trained with infinite 2D-3D pairs by setting different camera views (Xu et al. 2021). Nevertheless, estimating 3D pose from 2D pose introduces the depth ambiguity problem, one 2D pose can be mapped to multiple 3D poses.

Intuitively, rich texture and semantic information in images can assist in regressing a more accurate 3D pose from 2D pose. There has been some exploration in this direction. For example, Nie, Wei, and Zhu; Xu et al. segment im-

*Corresponding author




Figure 2: Cross-dataset evaluation between straightforward image-based model and our approach. With continuous training on Human3.6m, the former’s accuracy on 3DHP decreased, highlighting poor generalization ability.

age patches around keypoint locations to aid in generating the 3D pose. Likewise, Zhao et al.; Liu et al. introduced a method of superimposing extracted image features around keypoints’ position onto 2D keypoints to offer complementary information to the network. Yet despite the considerable progress, there exist some issues that need to be addressed. Firstly, because 3D human motion datasets were primarily captured in constrained laboratory environments, the distribution of image data is limited. Consequently, methods that are trained with image information tend to suffer from poor generalization ability, as shown in Fig 2. Additionally, effective strategies for leveraging image information are lacking.

This paper gives a novel insight into the cause of weak generalization ability and the specific effectiveness of image information in predicting the 3D pose. Based on that, we propose a novel framework that estimates 3D human pose from 2D pose by leveraging effective image cues, as shown in Fig 1.

To begin, we utilize the attention mechanism (Vaswani et al. 2017) to study the response of human keypoints to the image features. By analyzing the attention maps, we derived two noteworthy insights: 1. In general, for all keypoints, the attention maps exhibit a high-proportion, wide-range, and indiscriminate emphasis on irrelevant background information outside the human body. This may shed light on the weak generalization abilities of image-based methods, as they overly focus on dataset-biased information. 2. For a specific keypoint, its required image features are not confined solely to its own location in the image. Instead, the required positions also encompass body structure positions that provide depth information for that keypoint. For instance, the features of *elbow* can be instrumental in estimating the depth of *wrist* keypoint. This underscores the constraints of previous methods that exclusively concatenate localized image patches or features around keypoints (Nie, Wei, and Zhu 2017; Zhao et al. 2019; Liu et al. 2019).

Based on these understandings, we propose a novel 3D pose estimation framework. The key concept is to allow the keypoints to adaptively focus on critical image features. To

give an overview, the progressive learning framework consists of two stages. We enable the keypoints to query and select the beneficial features from all image patches in the first stage, called “Broad Query”. Then we prune the irrelevant image features (mostly background features). At last, we allow the keypoints to further explore information from these critical image features to obtain an accurate 3D pose, called “Focused Exploration”.

Specifically, in Stage 1, we introduce a Pose-guided Transformer Layer, which effectively reduces the keypoints’ attention to the background. It leverages the pose-to-image attention matrix to allow the image features to reversely query and aggregate the keypoints features. Through our design, the more crucial image features can extract more relevant information from the keypoint, while less important image features, like background patches, receive comparatively less information. Then we proposed an Adaptive Feature Selection Module, which aims to rank and prune the less important image features by the attention mechanism. In Stage 2, the keypoints are allowed to refocus on critical human image features by several Transformer Layers. Through this cascaded approach, the keypoints are empowered to dynamically explore critical features broadly and simultaneously prevent over-training to the background features.

We demonstrate quantitative results by conducting our method on standard 3D human pose benchmarks. Experimental results show that our method outperforms state-of-the-art performance on Human3.6M (Ionescu et al. 2013) and MPI-INF-3DHP (Mehta et al. 2017). Mention that our method not only significantly improves the accuracy of single-frame 3D pose estimation but also outperforms even the accuracy of 3D pose estimation networks based on temporal information. Our contribution can be summarized as follows:

- We propose two novel insights about 3DHPE methods involving image information. For one thing, overly focusing on dataset-biased background leads to poor generalization ability. For another, valuable image patches for estimating specific keypoint’s 3D coordinates are not confined to its exact image location; they extend to areas with structurally related positions.
- We propose a 3DHPE framework leveraging effective image features in two stages: broad query followed by focused exploration. It not only enables keypoints to determine all the necessary image features but also prevents excessive training on the background, thus improving generalization.
- We propose a novel Pose-guided Transformer Layer that effectively improves the keypoints’ ability to significant features. Besides, we propose an Adaptive Feature Selection Module, which adaptively stops irrelevant image features from further training.

Related Work

In the past few years, there has been extensive research on deep-learning-based algorithms for monocular 3D human pose estimation. Methods that directly regress 3D pose from the image are popular in the early stages (Li and Chan




Figure 3: Examples of attention map visualization of all keypoints on image features in two datasets.

Traindata	Testdata	MPJPE↓	Background Attention
H3.6M	H3.6M	30.4	73%
H3.6M	3DHP	74.2	75%

Table 1: Comparison of different test datasets and observed issues of excessive attention to the background and poor generalization.

2015). However, these approaches suffered from limited performance due to their reliance on training and testing within the constraints of 3D Motion Capture data (Xu et al. 2021). To address this limitation, the “lifting” method emerged as the dominant approach, offering better solutions to the problem.

“Lifting” Based 3D Human Pose Estimation

“Lifting” based approaches leverage off-the-shelf 2D human pose estimators trained on large and more diverse 2D datasets. By adopting this, the process of 3D human pose estimation is simplified to lifting the 2D pose to 3D pose without image participation. Martinez et al. first proposed a fully connected residual network in this approach. To handle the issue of depth ambiguity in the lifting process, some methods have leveraged temporal information (Pavlo et al. 2019; Chen et al. 2021) or proposed models with multiple hypotheses (Li and Lee 2019; Li et al. 2022).

Fusion Approach

Apart from the two mainstream approaches, there exist some methods that combine 2D pose with image information. Despite the remarkable attempts made by these methods, they still exhibit certain limitations. For example, some methods did not leverage off-the-shelf 2D pose estimators to generate 2D pose (Zhao et al. 2019; Liu et al. 2019). These approaches not only add a burden to the network but also fail to leverage the benefits of 2D estimators mentioned before. Besides, some methods employ rudimentary approaches to integrate image information. For instance, Nie, Wei, and Zhu; Xu et al. segment image patches around keypoint positions to assist in generating the 3D pose. Similarly, Zhao et al.; Liu et al. overlay image features extracted from keypoints position onto 2D keypoints. Nevertheless, the insight proposed in the next section proves this local concatenation approach might be ineffective. Moreover, Zhou et al.; Gong et al. utilize 2D keypoints heatmap on image to provide extra infor-




Figure 4: Visualization examples of heatmaps depicting the attention of specific keypoints.

mation. Indeed, heatmap only contains limited information, which may not be sufficient to accurately regress 3D poses.

Insight of Image Effect to 3DHPE

In this section, we study the roles and limitations of image features in estimating 3D pose using attention mechanisms. The keypoint-to-image attention map represents which image patches offer beneficial information for estimating the 3D coordinates of that keypoint.

Background Overfitting

Given the task of estimating the relative coordinates of human keypoints, the presence of non-contact backgrounds in 3D capture-environment datasets can be considered a form of dataset-biased noise. When we visualized the average attention maps of keypoints on the images, we observed a wide-ranging and indiscriminate focus on background features, as shown in Fig 3. This reveals the model’s overfitting to background information, which could be a potential cause of the poor generalization of image-based models. We further quantified the proportion of attention on background features, and found very high proportions on both datasets (73%, 75%), as shown in Table 1.

Structural Assistance

Logically speaking, for a specific human body keypoint in the image, the keypoint’s own image features can only provide its 2D coordinates in the image. However, the relative depth coordinate with respect to the pelvis point requires prior knowledge derived from combining other human structure features. We present attention of specific keypoints on image features, examples shown in Fig 4. Not surprisingly, our findings lead to the conclusion: not only the image features of keypoint’s locations are required. It will extend to body structure positions that provide depth information for that keypoint. For instance, the *knee* keypoint gives attention to ankles, the *wrist* keypoint gives attention to elbows and shoulders, and the *elbow* keypoint gives attention to shoulders. Hence, previous methods have been mistaken in their

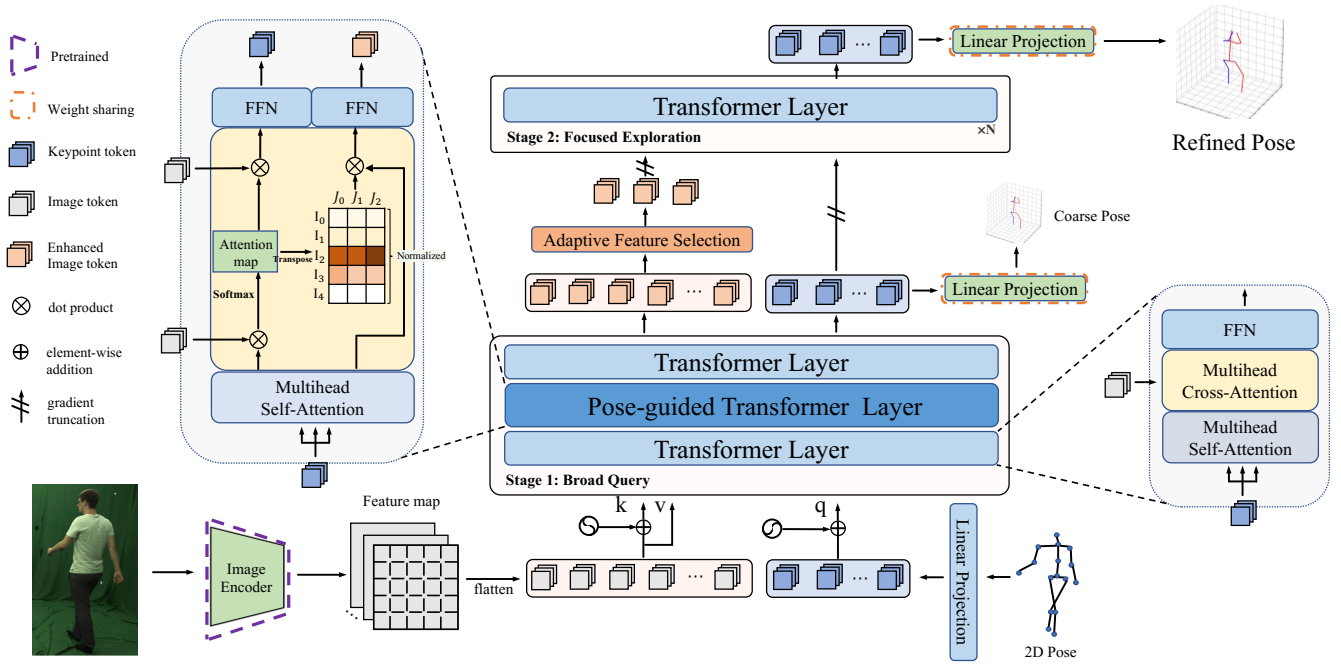


Figure 5: The overview of the proposed network.

assumption that only concatenating image patches or features around keypoints is sufficient.

Method

In this section, we provide a detailed description of the proposed framework, as illustrated in Fig 5. Given a 2D pose $J_{2d} \in \mathbb{R}^{N \times 2}$, our method aims to reconstruct the 3D pose $J_{3d} \in \mathbb{R}^{N \times 3}$ by effectively leveraging the information from a cropped image $I \in \mathbb{R}^{h \times w \times 3}$, where N is the number of keypoints, h, w is the input image size. To accomplish this, we proposed a Progressive Training framework. It consists of two stages. In Stage 1, we allow the keypoints to query beneficial information from all image features under coarse pose supervision until convergence. Subsequently, to counteract the detriment of background features' excessive training on generalization, we employ an Adaptive Feature Selection Module to prune less crucial image features. Then, in Stage 2, keypoints exclusively query the preserved image features to generate a refined pose.

Stage 1: Broad Query

The image I is fed into a 2d-pose-estimation-pretrained image encoder, resulting in features $F_I \in \mathbb{R}^{H \times W \times d}$, which are flattened into tokens $T_I \in \mathbb{R}^{HW \times d}$ with sequence length HW and dimension d . Similarly, the 2D pose J_{2d} is transformed to pose tokens $T_P \in \mathbb{R}^{N \times d}$ by linear projection.

Subsequently, the image tokens and keypoint tokens are fed into three consecutive Transformer Layers. Situated in the middle is the specially crafted Pose-guided Transformer Layer, intended to selectively enhance image tokens while diminishing the keypoints' focus on irrelevant image tokens.

The resulting keypoint tokens are then projected linearly to generate the coarse 3D pose denoted as $J_{3d_1} \in \mathbb{R}^{J \times 3}$.

Transformer Layer consists of three consecutive modules, including Multi-head Self-Attention (MSA), Multi-head Cross-Attention (MCA), and Feed Forward Network (FFN). Multi-head Attention can be formulated as:

$$A(Q, K, V) = \text{softmax}\left(\frac{Q \cdot K^\top}{\sqrt{d}}\right) \cdot V \quad (1)$$

In MSA, keypoints tokens are linearly mapped to Queries $Q \in \mathbb{R}^{N \times d}$, Keys $K \in \mathbb{R}^{N \times d}$, and Values $V \in \mathbb{R}^{N \times d}$. Similarly, in MCA, keypoints tokens are linearly mapped to Queries $Q \in \mathbb{R}^{N \times d}$, Image tokens are linearly mapped into Keys $K \in \mathbb{R}^{HW \times d}$, and Values $V \in \mathbb{R}^{HW \times d}$.

Pose-guided Transformer Layer. We designed a Pose-guided dual attention structure that effectively reduces the keypoints' attention to the background. It leverages the pose-to-image attention matrix to allow the image features to reversely query and aggregate the keypoints features. By our design, the more crucial image patches can obtain more information from the keypoint features.

Specifically, the novel attention mechanism produces two outputs: keypoint tokens \hat{T}_J and enhanced image tokens \hat{T}_I . The update of image tokens will be influenced by the update of keypoint tokens through Attention Map (\mathbb{A}). The formulations are as follows:

$$\begin{aligned} \mathbb{A} &= \text{softmax}\left(\frac{Q \cdot K^\top}{\sqrt{d}}\right) \\ \hat{T}_J &= \mathbb{A} \cdot V_I + T_J \\ \hat{T}_I &= \mathbb{A}^\top \cdot V_J + T_I \end{aligned} \quad (2)$$




Figure 6: Details of Adaptive Feature Selection Module.

Similarly, keypoints tokens are linearly mapped to Queries $Q \in \mathbb{R}^{N \times d}$, Image tokens are linearly mapped into Keys $K \in \mathbb{R}^{HW \times d}$. $V_I \in \mathbb{R}^{HW \times d}$ and $V_J \in \mathbb{R}^{N \times d}$ are Values linearly mapped from images tokens and keypoints tokens, respectively. The attention map $A \in \mathbb{R}^{N \times HW}$ represents the weighting that keypoint tokens assign to image tokens, and it is normalized using the *softmax* function. Similarly, the transposed attention map $A^T \in \mathbb{R}^{HW \times N}$ represents the weights image tokens assign to keypoint tokens. Because of the normalization before transposition, intuitively, the image tokens deemed more important receive a greater overall weight from the keypoint tokens. This signifies that image tokens with higher significance can gather more information from keypoint tokens, while on the contrary, less significant image tokens (typically background tokens) collect limited information.

We choose to replace only the second Transformer Layer with the proposed Pose-guided Transformer Layer for the following reasons: before guiding the update of image features, the keypoint tokens require an initial perception of the image features (by the first transformer layer) to evaluate their significance. The final layer cannot be replaced, as the update of image tokens lacks direct supervision in Stage 1.

Adaptive Feature Selection Module

To prevent less significant image tokens from further training and aggregating critical image tokens. We propose an Adaptive Feature Selection Module, details shown in Fig. 6. We leverage the attention map from the last transformer layer to rank the importance of image tokens. For simplicity, we aggregate the attention weights of all keypoints on image features and set a retention rate, denoted as r ($0 < r < 1$). The top $r \times HW$ image tokens with the highest weights are retained.

Stage 2: Focused Exploration

In this stage, we allow the keypoint tokens to further dig information from selected critical image tokens and generate a refined 3D pose $J_{3d_2} \in \mathbb{R}^{J \times 3}$. Specifically, we freeze the weights trained in the former stage and feed the keypoints tokens and selected image tokens into a new Transformer Block consisting of several Transformer Layers. Then, the output keypoint tokens will be projected to a refined 3D pose J_{3d_2} by Linear Projection.

Loss Function

Our model is trained with Mean Squared Error (MSE) loss.

$$L = \sum_{i=1}^J \|Y_i - \hat{Y}_i\|_2 \quad (3)$$

where Y_i and \hat{Y}_i represent the predicted and ground-truth 3D pose of joint i , respectively.

Experiments

Datasets and Evaluation Metrics

We evaluate our method on two widely-used datasets for 3DHPE: Human3.6M (Ionescu et al. 2013) and MPI-INF-3DHP (Mehta et al. 2017).

Human3.6M (H3.6M) is the largest and most representative benchmark for 3DHPE. Following Martinez et al., we use subject S1, S5, S6, S7 and S8 for training, and S9, S11 for testing. We down-sampled the original frame rate from 50 fps to 5 fps for faster training. The Mean Per Joint Position Error (MPJPE) is computed under two protocols: Protocol 1 computes MPJPE between ground truth and the estimated 3D poses after aligning their root (pelvis) keypoints; Protocol 2 is the MPJPE after aligning the estimated 3D pose with the ground truth using translation, rotation, and scale (P-MPJPE).

MPI-INF-3DHP (3DHP) provides monocular videos of six subjects acting in three different scenes, including indoors and outdoors. This dataset is often used to evaluate the generalization performance of different models. Following the convention, we directly apply our model trained on H3.6M dataset to this dataset without fine-tuning. We report results using three metrics, Mean Per Joint Position Error (MPJPE), Percentage of Correctly estimated Keypoints (PCK) with a threshold of 150 mm, and Area Under the Curve (AUC) a range of PCK thresholds.

Implementation Details

We take HRNet-w32 as our backbone with input size 256×192 , which is pretrained on MS COCO 2017 dataset (Lin et al. 2014), provided by Sun et al.. The retention r is set to 0.3. The number of Transformer layers in Stage 2 is set to 3. For a fair comparison, following previous work (Pavlo et al. 2019; Martinez et al. 2017), we obtain 2D pose detections cascaded pyramid network (CPN) (Chen et al. 2018) and stacked hourglass network (SH) (Newell, Yang, and Deng 2016). We take the ground-truth bounding boxes. Our model is implemented in Pytorch and optimized via Adam. All experiments are conducted on two NVIDIA RTX 3090 GPUs. The initial learning rate is set to 0.001 with a shrink factor of 0.9 per 4 epochs with 128 batch size. We first train the initial interaction stage and image encoder for 20 epochs and freeze them to train the remaining modules.

Comparison with State-of-the-art Methods

Results on Human3.6M. The proposed method is compared with the state-of-the-art methods on Human3.6M. The result and comparison of our model with SH detected,

Method	Dir.	Disc	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg.
Learning (Fang et al. 2018)	50.1	54.3	57.0	57.1	66.6	73.3	53.4	55.7	72.8	88.6	60.3	57.7	62.7	47.5	50.6	60.4
SemGCN (Zhao et al. 2019)*	47.3	60.7	51.4	60.5	61.1	49.9	47.3	68.1	86.2	55.0	67.8	61.0	42.1	60.6	45.3	57.6
Monocular (Xu et al. 2021)*	47.1	<u>52.8</u>	54.2	<u>54.9</u>	63.8	72.5	51.7	<u>54.3</u>	70.9	85.0	58.7	<u>54.9</u>	59.7	43.8	47.1	58.1
Graformer (Zhao, Wang, and Tian 2022)	49.3	53.9	54.1	55.0	63.0	69.8	51.1	53.3	69.4	90.0	<u>58.0</u>	55.2	60.3	47.4	50.6	58.7
Ours *	48.3	51.5	46.1	48.5	53.7	42.8	47.3	59.9	71.0	51.6	52.7	46.1	39.8	39.8	43.9	51.0
VideoPose (Pavlo et al. 2019) †	45.1	47.4	42.0	46.0	49.1	56.7	44.5	44.4	57.2	66.1	47.5	44.8	49.2	32.6	34.0	47.1
GraphSH (Xu and Takano 2021)	45.2	49.9	47.5	50.9	54.9	66.1	48.5	46.3	59.7	71.5	51.4	48.6	53.9	39.9	44.1	51.9
MGCN (Zou and Tang 2021)	45.4	49.2	45.7	49.4	50.4	58.2	47.9	46.0	57.5	63.0	49.7	46.6	52.2	38.9	40.8	49.4
MHFormer (Li et al. 2022) † f=243	39.2	43.1	40.1	40.9	44.9	<u>51.2</u>	40.6	41.3	53.5	<u>60.3</u>	43.7	41.1	43.8	29.8	30.6	43.0
Pose-Oriented (Li et al. 2023)	47.9	50.0	47.1	51.3	51.2	59.5	48.7	46.9	<u>56.0</u>	<u>61.9</u>	51.1	48.9	54.3	40.0	42.9	50.5
diffPose (Gong et al. 2023)	42.8	49.1	45.2	48.7	52.1	63.5	46.3	45.2	58.6	66.3	50.4	47.6	52.0	37.6	40.2	49.7
Ours*	44.9	<u>46.4</u>	42.4	<u>44.9</u>	<u>48.7</u>	40.1	<u>44.3</u>	55.0	58.9	47.1	48.2	<u>42.6</u>	36.9	48.8	40.1	<u>46.4</u>
SemGCN (Zhao et al. 2019)*	37.8	49.4	37.6	40.9	45.1	41.4	40.1	48.3	50.1	42.2	53.5	44.3	40.5	47.3	39.0	43.8
VideoPose (Pavlo et al. 2019) †	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	37.2
GraphSH (Xu and Takano 2021)	35.8	38.1	31.0	35.3	35.8	43.2	37.3	31.7	38.4	45.5	35.4	36.7	36.8	27.9	30.7	35.8
Graformer (Zhao, Wang, and Tian 2022)	32.0	38.0	30.4	34.4	34.7	43.3	35.2	31.4	38.0	46.2	34.2	35.7	36.1	27.4	30.6	35.2
MHFormer (Li et al. 2022) † f=243	27.7	<u>32.1</u>	29.1	28.9	<u>30.0</u>	<u>33.9</u>	33.0	31.2	37.0	<u>39.3</u>	30.0	<u>31.0</u>	29.4	22.2	23.0	<u>30.5</u>
Pose-Oriented (Li et al. 2023)	32.9	38.3	28.3	33.8	34.9	38.7	37.2	30.7	34.5	39.7	33.9	34.7	34.3	26.1	28.9	33.8
diffPose (Gong et al. 2023)	28.8	32.7	27.8	30.9	32.8	38.9	<u>32.2</u>	28.3	<u>33.3</u>	41.0	<u>31.0</u>	32.1	31.5	25.9	27.5	31.6
Ours*	29.5	30.1	25.0	<u>29.0</u>	28.5	28.6	26.9	<u>30.5</u>	31.1	27.7	32.4	27.7	24.8	<u>30.0</u>	<u>25.9</u>	28.6

Table 2: Quantitative comparison with the state-of-the-art methods on Human3.6M under Protocol 1, using SH (Newell, Yang, and Deng 2016) detected 2D poses (top), using CPN (Chen et al. 2018) detected 2D poses (middle), ground truth 2D poses (bottom) as inputs. (†) - uses temporal information. (*) - uses image information. Bold: best; Underlined: second best.

Method	PCK ↑	AUC ↑	MPJPE ↓
Simple (Martinez et al. 2017)	82.6	50.2	88.6
Cascaded (Li et al. 2020)	81.2	46.1	99.7
MGCN (Zou and Tang 2021)	86.1	53.7	-
Pose-Oriented (Li et al. 2023)	84.1	53.7	-
Ours	88.2	59.3	68.9

Table 3: Quantitative comparison with the state-of-the-art methods on MPI-INF-3DHP. Best in Bold.

CPN detected and ground-truth 2D pose are reported in Table 2. Our method outperforms all previous state-of-the-art methods by a large margin under protocol 1 for all single-frame methods (Methods w/o †). Besides, our method even achieves comparable results with the video-based method (Methods w/ †).

Result on MPI-INF-3DHP. To assess the generalization ability, we directly apply our method to 3DHP dataset trained on H3.6M. Table 3 shows the result of different methods with 2D ground-truth input. Our approach achieves the best performance on all metrics (PCK, AUC, and MPJPE). It emphasizes the strong generalization ability.

Qualitative Result. Figure 7 demonstrates some qualitative results on H3.6M dataset compared with Graformer (Zhao, Wang, and Tian 2022) with 2d ground-truth input. Our results achieve almost the same as the ground truth when faced with easy cases. For hard cases, compared with the baseline, our result can effectively reduce the gap between the estimated result and ground-truth.

Ablation Study and Discussion

Effect on Progressive Learning Strategy. We first compared the Coarse Pose and Refined Pose results on both

Strategy	Human3.6m	3DHP
end-to-end w/ fine supervision	32.6	76.4
end-to-end w/ both supervision	30.2	72.4
Coarse Pose	29.2	70.9
Refined Pose	28.6	68.9

Table 4: Ablation study on progressive learning.

	dataset	MPJPE	Background Attention
w/o PGTL	H3.6M	30.4	73%
	3DHP	74.2	75%
w/ PGTL	H3.6M	29.2	60%
	3DHP	70.9	63%

Table 5: Effect on the Pose-guided Transformer Layer (PGTL). Represented in Coarse Pose in MPJPE.

H3.6M and 3DHP on Table 4. It can be seen that Stage 2 brings 0.6mm performance on Human3.6m and 2.0mm performance on 3DHP. To demonstrate that performance improvement is not brought by the increased model parameters, we retrained our method from end to end without coarse pose supervision or with both supervisions. The performance drops 4.0mm or 1.6mm in Human3.6m and drops 7.5mm or 3.5mm in 3DHP.

To demonstrate the ability of our method to improve generalization, we compare the test result curves of end-to-end training and progressive training on the 3DHP dataset as shown in Fig. 8. It can be seen that in the case of training only on key features (Stage 2), the test accuracy of the 3DHP dataset continues to increase (lower in MPJPE). This demonstrates that preventing excessive training on background features can enhance generalization.

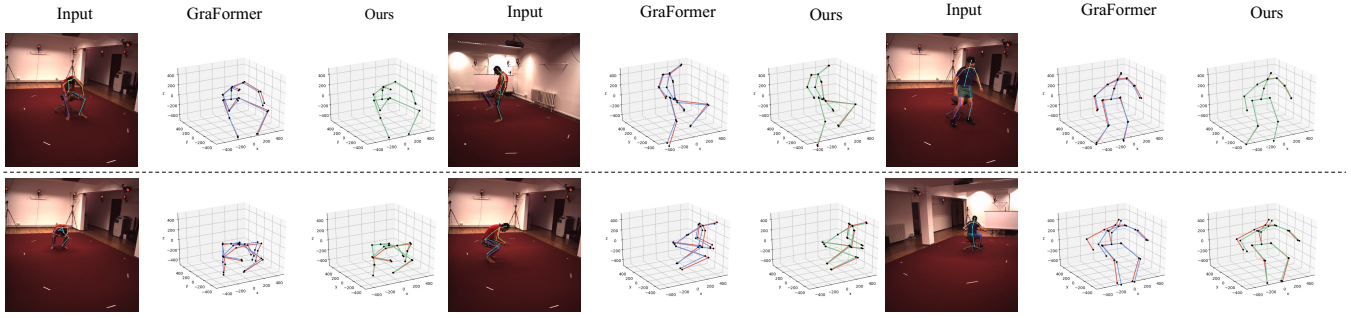


Figure 7: Qualitative results on Human3.6M. Green lines represent our results, results from baseline (Graformer) are represented by blue lines, and Ground-truth is represented in red. Easy case examples are shown at the top, and hard case examples are shown at the bottom.




Figure 8: Effect of progressive learning strategy on test curves on 3DHP dataset.

Transformer Layer	Pose-guided Layer	MPJPE
0	2	30.4
1	1	29.2
2	0	31.3

Table 6: Effect on number of Transformer Layers in Stage 1.

Ablation Study on Pose-guided Transformer Layer. We then diagnose the effect of the designed Pose-guide Transformer Layer. To test the ability to reduce the focus of keypoints on background information, we conducted a quantitative analysis. We took a 30-pixel-radius circle for the positions of each keypoints on the attention map, and defined these areas outside these circles as “background”. We then calculated the average attention of the keypoints to the background and performed statistical analysis on the test set. The result is shown in Table 5. We achieved 1.2mm on H3.6m improvement and 3.3mm improvement on 3DHP by replacing Transformer Layer with Pose-guide Transformer Layer. More importantly, the attention rate to background in the first stage dropped from 73% to 60% on H3.6m dataset, and from 75% to 63% on 3DHP dataset, which shows our approach is effective in enhancing the ability of keypoints to perceive critical features. In addition, we performed an ablation study on the number of layers in the Pose-guided Transformer Layer, shown in Table 6.

Effect on Adaptive Feature Selection Module. We then conduct an ablation study on Adaptive Feature Selection




Figure 9: Visualization of the retained feature (red pixels).

r	0.01	0.3	1	Coarse Pose
MPJPE	28.9	28.6	29.0	29.2

Table 7: Ablation study on the retention rate r .

Module. We first tested the effect of different retention rates (r) on the results, shown in Table 7.

It can be seen that we reach the best result when r is set to 0.3. We further visualize some examples of pruning on feature maps when r is set to 0.3, shown in Fig 9. It shows the module effectively prunes the background feature.

Conclusion

This paper gives a new insight into the cause of poor generalization problems and the effectiveness of image features. Based on that, we propose an advanced 3DHPE framework that leverages effective image cues and improves the generalization ability. It comprises two stages: the first involves a broad query for valuable image features, and the second stage focuses on critical features. To accomplish this, we proposed a novel Pose-guided Transformer Layer to reduce the keypoints’ attention to background and an Adaptive Feature Selection Module to prune less significant image features. Extensive experiments show that our method achieves state-of-the-art performance on two widely used benchmark datasets and shows great generalization ability. We hope our exploration can provide insights for future 3DHPE research.

Acknowledgements

This work was supported partly by the National Natural Science Foundation of China (Grant No. 62173045), and the Natural Science Foundation of Hainan Province (Grant No. 622RC675).

References

- Cao, P.; Yang, L.; Liu, D.; Liu, Z.; Li, S.; and Song, Q. 2022. Lsap: Rethinking inversion fidelity, perception and editability in gan latent space. *arXiv preprint arXiv:2209.12746*.
- Cao, P.; Yang, L.; Liu, D.; Liu, Z.; Li, S.; and Song, Q. 2023a. What Decreases Editing Capability? Domain-Specific Hybrid Refinement for Improved GAN Inversion. *arXiv preprint arXiv:2301.12141*.
- Cao, P.; Yang, L.; Zhou, F.; Huang, T.; and Song, Q. 2023b. Concept-centric Personalization with Large-scale Diffusion Priors. *arXiv preprint arXiv:2312.08195*.
- Chen, T.; Fang, C.; Shen, X.; Zhu, Y.; Chen, Z.; and Luo, J. 2021. Anatomy-aware 3d human pose estimation with bone-based pose decomposition. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(1): 198–209.
- Chen, Y.; Wang, Z.; Peng, Y.; Zhang, Z.; Yu, G.; and Sun, J. 2018. Cascaded pyramid network for multi-person pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7103–7112.
- Dang, Y.; Yang, F.; and Yin, J. 2020. DWnet: Deep-wide network for 3D action recognition. *Robotics and Autonomous Systems*, 126: 103441.
- Dang, Y.; Yin, J.; Zhang, S.; Liu, J.; and Hu, Y. 2022. Learning Human Kinematics by Modeling Temporal Correlations between Joints for Video-based Human Pose Estimation. *arXiv preprint arXiv:2207.10971*.
- Ding, P.; and Yin, J. 2022. Towards more realistic human motion prediction with attention to motion coordination. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(9): 5846–5858.
- Fang, H.-S.; Xu, Y.; Wang, W.; Liu, X.; and Zhu, S.-C. 2018. Learning pose grammar to encode human body configuration for 3d pose estimation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Gong, J.; Foo, L. G.; Fan, Z.; Ke, Q.; Rahmani, H.; and Liu, J. 2023. Diffpose: Toward more reliable 3d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13041–13051.
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2013. 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.
- Li, C.; and Lee, G. H. 2019. Generating multiple hypotheses for 3d human pose estimation with mixture density network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9887–9895.
- Li, H.; Shi, B.; Dai, W.; Zheng, H.; Wang, B.; Sun, Y.; Guo, M.; Li, C.; Zou, J.; and Xiong, H. 2023. Pose-Oriented Transformer with Uncertainty-Guided Refinement for 2D-to-3D Human Pose Estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1296–1304.
- Li, S.; and Chan, A. B. 2015. 3d human pose estimation from monocular images with deep convolutional neural network. In *Computer Vision—ACCV 2014: 12th Asian Conference on Computer Vision, Singapore, Singapore, November 1–5, 2014, Revised Selected Papers, Part II 12*, 332–347. Springer.
- Li, S.; Ke, L.; Pratama, K.; Tai, Y.-W.; Tang, C.-K.; and Cheng, K.-T. 2020. Cascaded deep monocular 3d human pose estimation with evolutionary training data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6173–6183.
- Li, W.; Liu, H.; Tang, H.; Wang, P.; and Van Gool, L. 2022. Mhformer: Multi-hypothesis transformer for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13147–13156.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Liu, J.; Ding, H.; Shahroudy, A.; Duan, L.-Y.; Jiang, X.; Wang, G.; and Kot, A. C. 2019. Feature boosting network for 3D pose estimation. *IEEE transactions on pattern analysis and machine intelligence*, 42(2): 494–501.
- Liu, X.; Yin, J.; Liu, J.; Ding, P.; Liu, J.; and Liu, H. 2020. Trajectorycnn: a new spatio-temporal feature learning network for human motion prediction. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(6): 2133–2146.
- Martinez, J.; Hossain, R.; Romero, J.; and Little, J. J. 2017. A simple yet effective baseline for 3d human pose estimation. In *Proceedings of the IEEE international conference on computer vision*, 2640–2649.
- Mehta, D.; Rhodin, H.; Casas, D.; Fua, P.; Sotnychenko, O.; Xu, W.; and Theobalt, C. 2017. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *2017 international conference on 3D vision (3DV)*, 506–516. IEEE.
- Newell, A.; Yang, K.; and Deng, J. 2016. Stacked hourglass networks for human pose estimation. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, 483–499. Springer.
- Nie, B. X.; Wei, P.; and Zhu, S.-C. 2017. Monocular 3d human pose estimation by predicting depth on joints. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 3467–3475. IEEE.
- Pavlo, D.; Feichtenhofer, C.; Grangier, D.; and Auli, M. 2019. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7753–7762.

- Sun, K.; Xiao, B.; Liu, D.; and Wang, J. 2019. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5693–5703.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Xu, T.; and Takano, W. 2021. Graph stacked hourglass networks for 3d human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16105–16114.
- Xu, Y.; Wang, W.; Liu, T.; Liu, X.; Xie, J.; and Zhu, S.-C. 2021. Monocular 3d pose estimation via pose grammar and data augmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10): 6327–6344.
- Zhao, L.; Peng, X.; Tian, Y.; Kapadia, M.; and Metaxas, D. N. 2019. Semantic graph convolutional networks for 3d human pose regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3425–3435.
- Zhao, W.; Wang, W.; and Tian, Y. 2022. Graformer: Graph-oriented transformer for 3d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20438–20447.
- Zhou, K.; Han, X.; Jiang, N.; Jia, K.; and Lu, J. 2019. Hem-lets pose: Learning part-centric heatmap triplets for accurate 3d human pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2344–2353.
- Zou, Z.; and Tang, W. 2021. Modulated graph convolutional network for 3D human pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11477–11487.