

# Probabilistic Restoration with Adaptive Noise Sampling for 3D Human Pose Estimation

Xianzhou Zeng<sup>1</sup>, Hao Qin<sup>1</sup>, Ming Kong<sup>1,2</sup>, Luyuan Chen<sup>3</sup>, Qiang Zhu<sup>1§</sup>

<sup>1</sup>Zhejiang University, Hangzhou, China

<sup>2</sup>Hikvision Research Institute, Hangzhou, China

<sup>3</sup>Beijing Information Science and Technology University, Beijing, China

{xzhouzeng, haoqin, zjukongming, zhuq}@zju.edu.cn, chenly@bistu.edu.cn

**Abstract**—The accuracy and robustness of 3D human pose estimation (HPE) are limited by 2D pose detection errors and 2D to 3D ill-posed challenges, which have drawn great attention to Multi-Hypothesis HPE research. Most existing MH-HPE methods are based on generative models, which are computationally expensive and difficult to train. In this study, we propose a Probabilistic Restoration 3D Human Pose Estimation framework (PRPose) that can be integrated with any lightweight single-hypothesis model. Specifically, PRPose employs a weakly supervised approach to fit the hidden probability distribution of the 2D-to-3D lifting process in the Single-Hypothesis HPE model and then reverse-map the distribution to the 2D pose input through an adaptive noise sampling strategy to generate reasonable multi-hypothesis samples effectively. Extensive experiments on 3D HPE benchmarks (Human3.6M and MPI-INF-3DHP) highlight the effectiveness and efficiency of PRPose. Code is available at: <https://github.com/xzhouzeng/PRPose>.

**Index Terms**—3D Human Pose Estimation, 2D-to-3D Lifting, Multi-Hypothesis Generation, Adaptive Noise

## I. INTRODUCTION

3D Human Pose Estimation (HPE) is a crucial task for restoring the 3D position of human joints from a monocular image, which has broad applications in virtual reality, action recognition and human-computer interaction, etc [1]. Current mainstream 3D HPE solutions follow a two-stage approach: first to estimate the 2D human pose and then predict the 3D pose based on the estimated 2D results, i.e., 2D-to-3D lifting [2], [3]. Among them, 2D-to-3D lifting is the pivotal stage, which presents two critical challenges: Firstly, the mapping from 2D to 3D is ill-posed, as a single 2D pose can correspond to multiple plausible 3D poses, impacting model convergence. Secondly, accurate 2D estimation is essential, as errors amplify during lifting, leading to unpredictable discrepancies in final 3D estimation outcomes [4].

Early methods for 2D-to-3D lifting estimate a unique 3D pose from a given 2D pose, named Single-Hypothesis Human Pose Estimation (SH-HPE), as illustrated in Figure 1.a, which primarily employ discriminative models, such as regression networks. These methods have high computational efficiency but suffer from the ill-posed nature of the 2D-to-3D mapping and inaccurate 2D pose estimation [4]. Recent research [5] addresses these issues through Multi-Hypothesis Human Pose

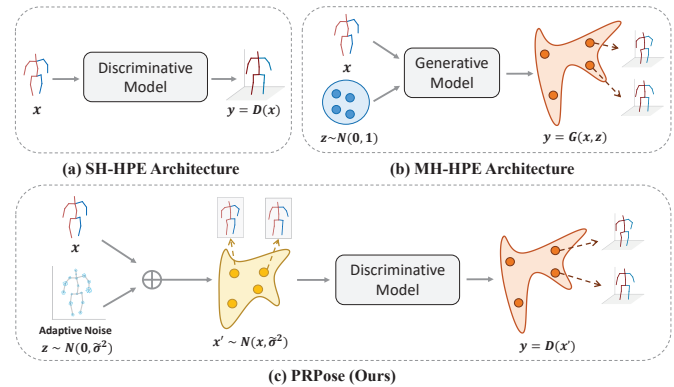


Fig. 1. (a) The classical SH-HPE architecture produces a unique 3D pose based on a single input 2D pose; (b) The traditional MH-HPE architecture based on the generative model, adding random noise to generate multi-hypothesis results; (c) We propose PRPose that can be extended from any SH-HPE method, generating the input distribution through adaptive noise to reconstruct the original probabilistic modeling process of SH-HPE, and producing multi-hypothesis outputs.

Estimation (MH-HPE), which leverages probabilistic modeling to accommodate estimation errors stemming from imprecise 2D pose detection or ill-posed mapping, generating multiple distinct 3D pose outputs aligned with the input 2D pose.

Most MH-HPE methods adopt generative models (e.g., Diffusion Model) to enhance 2D-3D mapping for 3D human pose estimation, as shown in Figure 1.b. These methods introduce fixed-distribution noise into the generative model to diversify the output [6], [7]. However, such processes are computationally demanding and challenging to converge. Moreover, existing methods employ uniform perturbations across all joints, disregarding variations in pose and joint complexities. Consequently, joints with higher uncertainty, such as those on occluded limbs, may suffer from insufficient hypothesis generation, while easily estimable joints, such as trunk joints, could experience unnecessary perturbations, leading to excessive disturbances and ineffective hypothesis sampling.

Recent studies have revealed that the SH-HPE method has potential probabilistic modeling capabilities in 2D-to-3D lifting [8]. Inspired by that, we propose a novel PRPose method (Fig. 1.c) to take full advantage of the probabilistic

This work is supported by the National Science and Technology Major Project (Grant No.2022ZD0115904).

<sup>§</sup> Corresponding author.

modeling capabilities of SH-HPE models. Specifically, we propose a weakly supervised adaptive noise sampling strategy that creates pseudo-labels from the single-hypothesis model to generate variable noise scales for error-prone joints. Through a weakly supervised learning process, PRPose recovers the probabilistic modeling process of the SH-HPE model, generating an adaptive variance matrix from the 2D pose, and remapping it to create adaptive noise for the 2D input joints. Finally, the SH-HPE model generates multiple 3D pose hypotheses based on different sampled 2D poses. Compared to previous methods, the scaling noise generated by PRPose adapts to 2D pose and joint topology, enhancing the realism of 3D pose assumptions.

Extensive experiments conducted on the Human3.6M and MPI-INF-3HP datasets show the comparable performance of the PRPose to the state-of-the-art approaches with a speed enhancement of two orders of magnitude. Comprehensive ablation experiments and visualizations demonstrate the reasoning for using adaptive noise addition and the effectiveness of a weakly supervised training strategy. Our main contributions can be summarized as follows:

- The proposed architecture enables simple extension of any SH-HPE model to the MH-HPE task. We preserve the lightweight and trainable advantages of the SH-HPE model, while suppressing the inherent flaws of the SH-HPE, such as 2D estimation error and ill-posed nature in the 2D-to-3D mapping process.
- We design a weakly supervised adaptive noise learning strategy that generates adaptive noise based on input 2D poses and joint positions. During the 2D-to-3D lifting process, we apply large noise corrections to joints with potentially large estimation errors, thereby improving the generated multi-hypothesis hit rate.
- Our method achieves a good balance between speed and accuracy. On the Human3.6M and MPI-INF-3DHP datasets, we achieve comparable accuracy to GFPose [7] (SOTA), but with over 100x faster pose estimation.

## II. RELATED WORK

Martinez *et al.* [4] first proposed a two-stage human pose estimation method that generates a unique 3D pose estimation from the 2D input, called Single-Hypothesis HPE (SH-HPE). They also proved that the error of the two-stage methods is mainly caused by the 2D pose estimation errors. Subsequent works tried to minimize the ambiguity in the 2D-to-3D lifting process [9], or propose data augmentation strategies to address the issue of limited available data and decrease the influence of data bias [10]. The information transmission between joints is another focus. Researchers tried to achieve the hidden-state fusion between joints with GCN [2] and Transformer [11], so as to enhance the representation ability and benefit highly-uncertainty joints prediction. Most SH-HPE methods employ discriminative models, which are computationally cheap and easy to train. However, limited by a single hypothesis output, SH-HPE methods still face challenges in solving the ill-posed problem of 2D to 3D lifting and 2D pose error transfer [12].

To address the limitations of SH-HPE, Jahangiri *et al.* [12] first proposed the idea of multiple hypotheses estimation (MH-HPE), with the main purpose of improving the accuracy and robustness adapted to various complexity and variability of scenes. Early MH-HPE methods proposed models with multiple prediction heads to achieve multi-hypothesis results [13], [14]. Thanks to the excellent capability of probabilistic modeling, most recent works applied generative models on the MH-HPE task, such as conditional variational autoencoder [5], [15] or normalizing flow [6]. With the enormous improvement in the generation capabilities of diffusion models [16], [7] propose GFPose and significantly improved the state-of-the-art of MH-HPE.

## III. METHOD

### A. Overview

Our work aims to extend lightweight single-hypothesis models to the multi-hypothesis method to simplify pose generation and enhance the efficiency of the 3D HPE task. Our proposed framework, shown in Figure 2, comprises two main modules: the Single-Hypothesis HPE (SH-HPE) and the Weakly Supervised Adaptive Noise Learning (WS-ANL). The SH-HPE module is the major task to train a discriminative single-hypothesis model, which is free to plug in existing frameworks (e.g., HTNet). The WS-ANL module is the auxiliary task to train an adaptive noise generator that is weakly supervised by pseudo-labels based on the difference between single-hypothesis prediction results and 3D pose ground truth. The inference process is performed in two steps. Firstly, we predict adaptive noise and add it to the original 2D pose to produce multiple augmented poses. Next, we input these 2D samples into a single-hypothesis module to generate multiple 3D poses.

Specifically, given the input 2D pose  $x \in \mathbb{R}^{2V}$  that contains  $V$  joints, our method can be expressed as:

$$Y = \Phi(\Psi(x, S)) \quad (1)$$

where  $Y = \{y_1, y_2, \dots, y_S | y_i \in \mathbb{R}^{3V}\}$  is the multi-hypothesis output of the model,  $\Psi$  is the sample amplification step that relies on adaptive noise, and  $S$  is the number of amplifications for 2D samples,  $\Phi$  is a single-hypothesis model for HPE task.

Next, we will briefly explain the working mechanism of the SH-HPE and the models adopted in our experiments and then introduce the design and implementation of the adaptive noise generation module in detail.

### B. SH-HPE Module

SH-HPE models are usually based on discriminative networks, and their working paradigm can be expressed as:

$$y = \Phi_\theta(x) \quad (2)$$

where  $y \in \mathbb{R}^{3V}$  is a single 3D skeleton prediction, and  $\theta$  is the trainable parameters in the SH-HPE model  $\Phi$ . Generally, the training process of  $\Phi$  is supervised by the ground truth 3D skeleton:

$$\mathcal{L} = |y - \hat{y}| \quad (3)$$

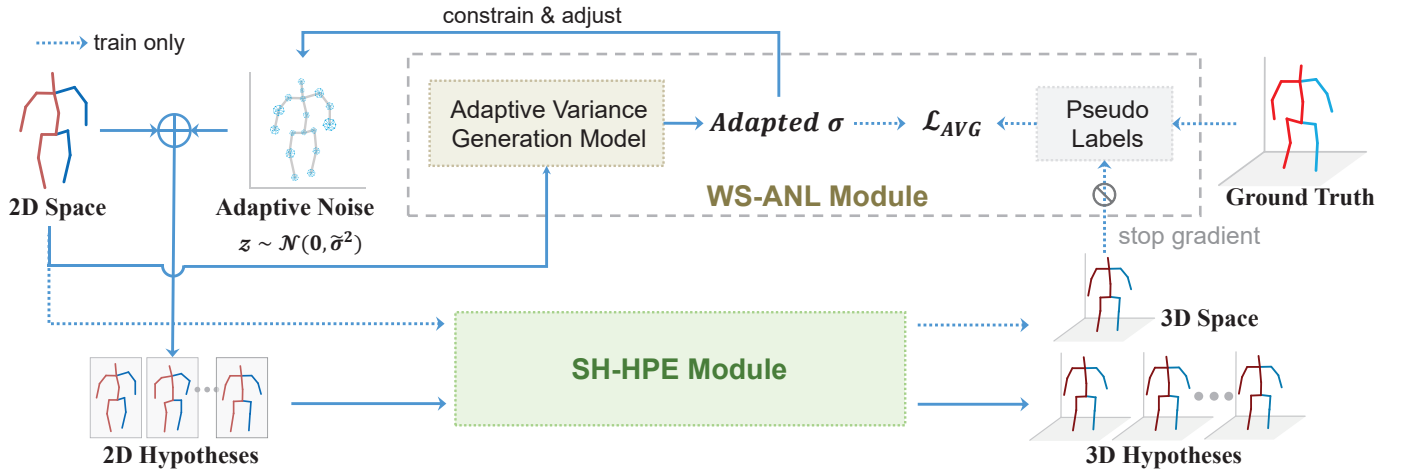


Fig. 2. Our proposed architecture PRPose consists of two main modules: a) Weakly Supervised Adaptive Noise Learning (WS-ANL): completes weakly supervised training of Adaptive Variance Generation Model; b) Single-Hypothesis Human Pose Estimation (SH-HPE): superimpose adaptive noise to the original 2D pose to generate multiple 2D hypotheses; map a single 2D pose to a 3D pose by an SH-HPE model (such as HTNet) during the inference process.

where  $\hat{y}$  is the ground truth 3D skeleton. SH-HPE often uses a model with a small amount of calculation and low training consumption (such as a regression model) as the backbone network of the framework. Our method can be adapted to any existing SH-HPE model and directly use the trained weights without minor adjustments.

Note that our framework can be seamlessly integrated with any SH-HPE model and can be customized to meet various scene-specific requirements, achieving an optimal balance of estimation accuracy and computational performance. In experiments, we compared two classic SH-HPE models HTNet [3] and MGCN [2]. The specific impact of different SH-HPE backbone models on the accuracy and speed of 3D HPE task is compared and analyzed.

### C. WS-ANL Module

The main purpose of adding noise to the MH-HPE method is to improve algorithm robustness. A strategy we often adopt is to generate more pose hypotheses by introducing random noise in the pose generation step in order to improve the hit rate of the 3D estimation results compared to the real pose. This technique can effectively improve the search range and space of the algorithm, reduce the overlap between hypotheses, and improve the robustness, accuracy, and generalization ability of the algorithm. However, the random noise strategy does not consider the variability brought by different poses and joints. It cannot fully utilize the prior information of 2D pose and the human body model itself. Therefore, we train an adaptive variance generative model ( $AVG$ ) with generalization ability through weakly supervised learning. The adaptive model can add customized noise according to pose and joints to achieve the best pose estimation effect:

$$A^\sigma = AVG_\varphi(x) \quad (4)$$

where  $AVG$  represents the Adaptive Variance Generation model,  $\varphi$  is its parameter, and  $A^\sigma = \{\sigma_1, \dots, \sigma_V\}$  is the set of output adaptive variance.

Without relying on the human-craft labels, we employ pseudo-labels for weakly supervised training to  $AVG$ . We hope that the output of  $AVG$  can be proportional to the magnitude of the error on each joint. Therefore, we first use the existing SH-HPE model to estimate 3D pose  $y$  corresponding to the current 2D input  $x$ , and then calculate the difference between the 3D estimate and the 3D ground truth as a supervisory signal:

$$\hat{A}^\sigma = \frac{D(y, \hat{y})}{C} \quad (5)$$

where  $\hat{A}^\sigma \in \mathbb{R}^V$  is treated as a pseudo-label to supervise the output of  $AVG$ ,  $D(\cdot, \cdot)$  represents the calculation of Euclidean distance, and  $C$  is a statistical normalization constant:

$$C = \frac{\sum_{p \in P} \sum_{v \in V} D(y_v^p, \hat{y}_v^p)}{PV} \quad (6)$$

where  $P$  is the number of training samples in the dataset.  $C$  counts the average single-hypothesis estimation error of each joint of all samples in the training set. We use it as a normalization constant to stabilize the training process. Finally, calculate the loss value of this process and optimize  $AVG$  by gradient descent method:

$$\mathcal{L}_{AVG} = \sum_x \frac{\sum_{v \in V} (A_v^\sigma - \hat{A}_v^\sigma)^2}{V} \quad (7)$$

During this process, the parameters of the SH-HPE model are frozen and do not participate in the update process.

In the test phase, for a specific 2D input  $x$ , the Gaussian noise distribution  $z_v \sim \mathcal{N}(0, \tilde{\sigma}_v^2)$  corresponding to each joint can be obtained after constrained and adjusted:

$$\tilde{\sigma}_v = \alpha * \max(\sigma_v, 1) \quad (8)$$

TABLE I

POSE ESTIMATION RESULTS ON THE HUMAN3.6M DATASET. WE REPORT THE MINMPJPE(MM) UNDER PROTOCOL #1 (NO RIGID ALIGNMENT) AND PROTOCOL #2 (WITH RIGID ALIGNMENT).  $S$  DENOTES THE NUMBER OF HYPOTHESES. BOLD AND UNDERLINED HIGHLIGHT THE OPTIMAL AND SUBOPTIMAL RESULTS, RESPECTIVELY, AND THE \* MARK HIGHLIGHTS THE OPTIMAL/SUBOPTIMAL RESULTS UNDER  $S = 10$ .

Protocol #1		Dire.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD	Smoke	Wait	WalkD	Walk	WalkT	Avg
GAN [17] ( $S = 10$ )	BMVC'20	62.0	69.7	64.3	73.6	75.1	84.8	68.7	75.0	81.2	104.3	70.2	72.0	75.0	67.0	69.0	73.9*
MDN [13] ( $S = 5$ )	CVPR'19	43.8	48.6	49.1	49.8	57.6	61.5	45.9	48.3	62.0	73.4	54.8	50.6	56.0	43.4	45.5	52.7
CVAE [5] ( $S = 200$ )	ICCV'19	37.8	43.2	43.0	44.3	51.1	57.0	39.7	43.0	56.3	64.0	48.1	45.4	50.4	37.9	39.9	46.8
GraphMDN [14] ( $S = 200$ )	IJCNN'21	40.0	43.2	41.0	43.4	50.0	53.6	40.1	41.4	52.6	67.3	48.1	44.2	44.9	39.5	40.2	46.2
NF [6] ( $S = 200$ )	ICCV'21	38.5	42.5	39.9	41.7	46.5	51.6	39.9	40.8	49.5	56.8	45.3	46.4	46.8	37.8	40.4	44.3
GFPose [7] ( $S = 10$ )	CVPR'23	39.9	44.6	40.2	41.3	46.7	53.6	41.9	40.4	52.1	67.1	45.7	42.9	46.1	36.5	38.0	45.1*
GFPose [7] ( $S = 200$ )	CVPR'23	31.7	35.4	31.7	32.3	36.4	42.4	32.7	31.5	41.2	52.7	36.5	34.0	36.2	29.5	30.2	<b>35.6</b>
<b>Ours</b> ( $S = 10$ )		40.1	44.3	39.7	43.1	45.5	50.3	41.8	40.8	50.5	58.6	44.4	41.7	47.6	35.9	36.9	<b>44.1*</b>
<b>Ours</b> ( $S = 200$ )		34.5	38.3	33.9	37.1	39.8	43.3	35.6	35.3	44.3	51.0	38.9	36.1	41.7	31.2	31.7	<b>38.2</b>
Protocol #2		Dire.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD	Smoke	Wait	WalkD	Walk	WalkT	Avg
CVAE [5] ( $S = 200$ )	ICCV'19	30.6	34.6	35.7	36.4	41.2	43.6	31.8	31.5	46.2	49.7	39.7	35.8	39.6	29.7	32.8	37.3
GraphMDN [14] ( $S = 200$ )	IJCNN'21	30.8	34.7	33.6	34.2	39.6	42.2	31.0	31.9	42.9	53.5	38.1	34.1	38.0	29.6	31.1	36.3
NF [6] ( $S = 200$ )	ICCV'21	27.9	31.4	29.7	30.2	34.9	37.1	27.3	28.2	39.0	46.1	34.2	32.3	33.6	26.1	27.5	32.4
GFPose [7] ( $S = 200$ )	CVPR'23	26.4	31.5	27.2	27.4	30.3	36.1	26.8	26.0	38.4	45.8	31.2	29.2	32.2	23.1	25.8	<b>30.5</b>
<b>Ours</b> ( $S = 200$ )		27.7	31.2	28.6	30.5	32.2	35.1	28.2	27.9	36.9	41.4	32.2	28.4	33.4	24.7	26.1	<b>31.0</b>

where  $\alpha$  is a coefficient for adjusting the noise distribution. This operation can increase the noise variance of joints with minor errors while maintaining the noise variance of joints with significant errors unchanged, ensuring the adaptability of sampling and avoiding weakening the solution ability of the MH-HPE method to the ill-posed problem due to the weak noise addition of some joints. A new 2D sample  $x'$  is obtained by randomly sampling a noise  $z_v \sim \mathcal{N}(0, \sigma_v^2)$  and adding it to the original coordinates of  $x$ :

$$x'_v = x_v + z_v \quad (9)$$

where the two coordinates of  $x_v$  are noised separately, and the two are independent. We repeat this process  $S$  times to generate  $S$  2D input samples that will be fed into an SH-HPE model to obtain multi-hypothesis predictions.

To further explain the rationality of our approach and explore the architecture, we provide the formula derivation of PRPose probabilistic modeling and the design and experiment of other paradigms in the supplementary material.

#### IV. EXPERIMENTS

##### A. Datasets and Evaluation Metrics

**Human3.6M** is currently the most widely used dataset for 3D human pose estimation [18]. We evaluate our method's performance on Human3.6M using two standard protocols, Protocol #1 (P1) and Protocol #2 (P2), similar to prior studies [7]. Samples of five subjects (S1, S5, S6, S7, and S8) are used to train the model, and samples of two subjects (S9 and S11) are used for testing.

**MPI-INF-3DHP** is a 3D human pose dataset that contains both indoor and complex outdoor scenes [19]. This dataset is commonly employed to evaluate a model's generalization capability. Therefore, we train our model on Human3.6M and evaluate its performance on MPI-INF-3DHP. Following prior research [7], we employ the metric of Percentage of Correct Keypoints (PCK) with a threshold of 150mm to evaluate the accuracy of the 3D human pose estimation.

TABLE II

COMPARISON OF POSE ESTIMATION RESULTS BASED ON SEQUENCES ( $T = 243$ ) ON THE HUMAN3.6M DATASET. THE P-BEST AND J-BEST SETTINGS FOLLOW D3DP [23].

Method		P1↓	P2↓
D3DP [23] ( $S = 20, P\text{-Best}$ )	ICCV'23	39.5	31.2
D3DP [23] ( $S = 20, J\text{-Best}$ )	ICCV'23	35.4	28.7
<b>Ours-Seq</b> ( $S = 20, P\text{-Best}$ )		<b>36.7</b>	<b>29.7</b>
<b>Ours-Seq</b> ( $S = 20, J\text{-Best}$ )		<b>25.5</b>	<b>20.3</b>

##### B. Implementation Details

For the experiments conducted on the Human3.6M dataset, we use the CPN [20] for 2D pose detection, following previous single-hypothesis works [3]. For generalization testing on the MPI-INF-3DHP dataset, we employ the pre-trained HRNet [21] as our 2D pose detector and then map to the skeletal definition of Human3.6M through interpolation. We apply the HTNet as the SH-HPE module and set the AVG as an MGCN framework with a channel dimension of 256. To align with existing multi-hypothesis works, our models have not utilized the Refinement Module [22] often mentioned in single-hypothesis works. The parameters of the single-hypothesis model are frozen, and only the parameters of the AVG model are trained. All experiments are conducted on a single NVIDIA RTX 3090 GPU. During the testing process, the coefficient  $\alpha$  for adjusting the noise distribution is set to 0.005 for the Human3.6M dataset and 0.01 for the MPI-INF-3DHP dataset.

##### C. Quantitative Results

**Effectiveness.** We validated the effectiveness of our method on the Human3.6M dataset. Following previous works [6, 7], we generated  $S$  3D pose estimates for each 2D pose and reported the minimum MPJPE (Mean Per Joint Position Error) between the ground truth and all estimates. As shown in Table I, with 10 hypotheses, our method outperforms the GFPose [7] accuracy in the Protocol #1 testing setting. With

TABLE III  
POSE ESTIMATION RESULTS ON THE MPI-INF-3DHP DATASET. FOLLOWING PRIOR RESEARCH [7], 200 SAMPLES ARE DRAWN. "GS" REPRESENTS THE "GREEN SCREEN".

Method	GS $\uparrow$	noGS $\uparrow$	Outdoor $\uparrow$	ALL PCK $\uparrow$
MDN [13] CVPR'19	70.1	68.2	66.6	67.9
NF [6] ICCV'21	86.6	82.8	82.5	84.3
GAN [17] BMVC'20	86.9	86.6	79.3	85.0
GFPose [7] CVPR'23	88.4	87.1	84.3	86.9
<b>Ours</b>	<b>93.1</b>	<b>87.6</b>	<b>89.3</b>	<b>90.2</b>

TABLE IV  
COMPARISON OF FPS RESULTS ON THE HUMAN3.6M DATASET. ALL DATA WERE TESTED ON THE SAME RTX 3090, WITH ONE FRAME TREATED AS ONE SAMPLE AND MULTIPLE HYPOTHESES GENERATED IN PARALLEL WITHIN A SINGLE SAMPLE.

Method	FPS ( $S = 10$ )	FPS ( $S = 200$ )	Param (M)	Protocol #1 ( $S = 10/200$ )
GraphMDN [14]	—	121.87	0.69	—/46.2
NF [6]	—	188.94	2.16	—/44.3
<b>Ours-L</b>	—	<b>201.51</b>	<b>0.17</b>	—/42.9
GFPose [7]	0.56	0.54	12.50	45.1/35.6
<b>Ours-B</b>	<b>66.62</b>	<b>59.29</b>	<b>4.14</b>	<b>44.1/38.2</b>

200 hypotheses, our method is slightly lower or comparable to the state-of-the-art accuracy [7] in Protocol #1 and Protocol #2 testing settings. Where  $S=200$  focuses on evaluating the upper limit of model performance. In practice, given the balance between efficiency and performance, the hypothetical number is usually a modest number, such as  $S=10$ . Considering the computational efficiency advantage (improved by two orders of magnitude compared to GFPose), the superiority of our method is still significant.

In addition, due to the universality of our proposed framework, we also extend PRPose to sequence-based 2D-to-3D lifting models (named *Ours-Seq*), and the results are also presented in Table II. Specifically, we use MixSTE [11] as the SH-HPE for experimentation, with video as a sample dimension, meaning that all frames in a video have the same noise added, thereby expanding the model to multi-hypothesis generation without changing the smoothness characteristics of the video. Finally, our method significantly exceeds the latest sequence based work D3DP [23], further verifying the universality and enormous potential of our architecture. For more details and discussion, please refer to Section S4.1 of the supplementary material.

**Generalization.** We evaluated the generalization ability of our model on the MPI-INF-3DHP dataset. All parameters of the model were trained on the Human3.6M dataset and were not fine-tuned on the MPI-INF-3DHP dataset. As shown in Table III, our method outperformed the state-of-the-art approaches [6], [7], and [17] by a significant margin. Even compared with two models that use ground-truth 2D pose [13] and are specifically designed for transfer learning [17], our method still achieved a remarkable improvement. Unlike other methods, our approach did not suffer a decrease in accuracy compared to the "noGS" condition in complex outdoor envi-

TABLE V  
ABLATION EXPERIMENT OF ADAPTIVE NOISE ON MPI-INF-3DHP DATASET. "NO ADAPTED" DENOTES CONSTANT VARIANCE WITHOUT ADJUSTMENT. "SAMPLE-JOINTS ADAPTED" FOLLOW THE MAIN EXPERIMENT SETTINGS.

Strategy	PCK ( $S = 10$ )	PCK ( $S = 200$ )
No adapted	82.47	86.11
Joints adapted	82.57	86.17
Sample adapted	84.40	89.12
Sample-Joints adapted	<b>84.61</b>	<b>90.19</b>

ronments. This is due to the regulation of the adaptive noise generation module, which allowed the correction of joints with large estimation errors.

**Efficiency.** To prove the advantages of our model on computational efficiency, we compare the frames per second (FPS) of our model with the state-of-the-art methods. In addition to our basic model defined above (hereby named *Ours-B* for distinguish), we propose a lightweight mode (named *Ours-L*), where the SH-HPE module is a lightweight version of MGCN [2] with a channel dimension of 128 that removes non-local modules and shares the same weight among different nodes, and *AVG* is implemented with a single-block Simple-Baseline3D [4].

The performance and efficiency comparison is shown in Table IV. The results illustrate that *Ours-B* achieved advanced estimation accuracy with a high inference speed. Especially, compared with GFPose, the state-of-the-art method of MH-HPE, *Ours-B* obtained more than 100 times speed gain at a comparable accuracy. Besides, *Ours-L* achieved the highest FPS while maintaining a competitive accuracy. The results indicate that our method can achieve a perfect balance between speed and accuracy.

**Ablation Study for Adaptive Noise.** We conducted ablation experiments to observe the impact of the adaptive noise generation module on the model's accuracy and generalization ability. The experimental results are shown in Table V. It can be found that compared with no adapted or single adapted, Sample-Joints adapted shows better performance, indicating the necessity of adaptive noise modules. With the increase in the number of hypotheses, the probability of sampling the expected sample increases, and the advantages of adaptive noise become more obvious.

#### D. Qualitative Results

To further demonstrate the effectiveness of our framework and adaptive noise addition strategy, we visualized module outputs and compared with No adapted strategy. Figure 3 displays results for Waiting, Smoking, Sitting Down actions. (During Smoking and Sitting Down action, some joints are obviously blocked.) Compared with No adapted strategy, our method can generate more diverse 3d human poses at joints with large estimation error and high depth uncertainty. This further suggests that adaptive noise is more effective in difficult samples. For those simple samples, fixed distribution noise can already complete the bias correction, making the

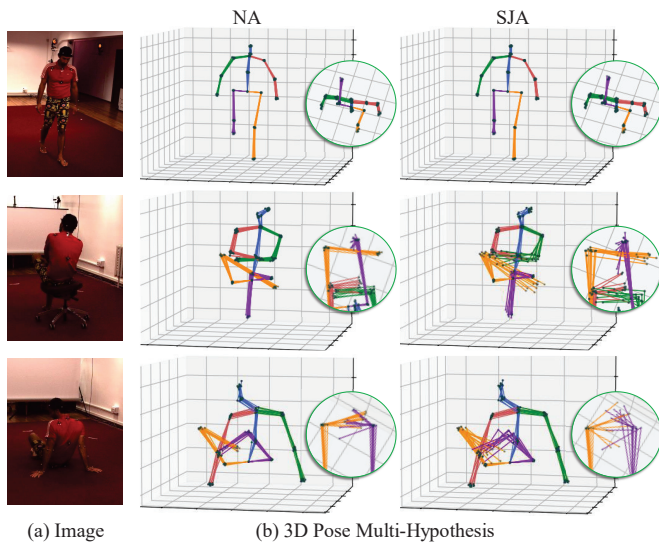


Fig. 3. Qualitative comparison of multi-hypothesis outputs for Sample-Joints adapted (SJA) and No adapted (NA) methods on the Human3.6M dataset.

advantage of adaptive noise less obvious. When the sample quality is low, fixed distribution noise is difficult to apply to various situations, and the advantages of adaptive noise are manifested.

## V. CONCLUSION

Aiming at the ill-posed problem of 2D-to-3D lifting in 3D HPE and the high computational cost of existing multi-hypothesis methods, we propose a general framework called PRPose for the probabilistic modeling process reconstruction of 3D human pose estimation. With adaptive noise sampling, it seamlessly extends lightweight single-hypothesis models to the multi-hypothesis method, significantly enhancing computational efficiency while maintaining state-of-the-art accuracy. Through extensive experiments, we demonstrate the effectiveness and rationality of our approach. Our work provides a new way of thinking for 3D HPE, enabling multi-hypothesis research to directly benefit from improving single-hypothesis models. Our work also reveals the great significance of researching adaptive noise generation strategies based on human topology and 2D pose, and we will explore these issues in the future.

## REFERENCES

- [1] W. Liu, Q. Bao, Y. Sun, and T. Mei, "Recent advances of monocular 2d and 3d human pose estimation: A deep learning perspective," *ACM Comput. Surv.*, vol. 55, no. 4, nov 2022. [Online]. Available: <https://doi.org/10.1145/3524497>
- [2] Z. Zou and W. Tang, "Modulated graph convolutional network for 3d human pose estimation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 11 477–11 487.
- [3] J. Cai, H. Liu, R. Ding, W. Li, J. Wu, and M. Ban, "Htnet: Human topology aware network for 3d human pose estimation," *arXiv preprint arXiv:2302.09790*, 2023.
- [4] J. Martinez, R. Hossain, J. Romero, and J. J. Little, "A simple yet effective baseline for 3d human pose estimation," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2640–2649.
- [5] S. Sharma, P. T. Varigonda, P. Bindal, A. Sharma, and A. Jain, "Monocular 3d human pose estimation by generation and ordinal ranking," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 2325–2334.
- [6] T. Wehrbein, M. Rudolph, B. Rosenhahn, and B. Wandt, "Probabilistic monocular 3d human pose estimation with normalizing flows," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 11 199–11 208.
- [7] H. Ci, M. Wu, W. Zhu, X. Ma, H. Dong, F. Zhong, and Y. Wang, "Gfpose: Learning 3d human pose prior with gradient fields," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 4800–4810.
- [8] H. Li, B. Shi, W. Dai, H. Zheng, B. Wang, Y. Sun, M. Guo, C. Li, J. Zou, and H. Xiong, "Pose-oriented transformer with uncertainty-guided refinement for 2d-to-3d human pose estimation," 2023.
- [9] K. Zhou, X. Han, N. Jiang, K. Jia, and J. Lu, "Hemlets push: learning part-centric heatmap triplets for 3d human pose and shape estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 6, pp. 3000–3014, 2021.
- [10] S. Li, L. Ke, K. Pratama, Y.-W. Tai, C.-K. Tang, and K.-T. Cheng, "Cascaded deep monocular 3d human pose estimation with evolutionary training data," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 6173–6183.
- [11] J. Zhang, Z. Tu, J. Yang, Y. Chen, and J. Yuan, "Mixste: Seq2seq mixed spatio-temporal encoder for 3d human pose estimation in video," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 13 232–13 242.
- [12] E. Jahangiri and A. L. Yuille, "Generating multiple diverse hypotheses for human 3d pose consistent with 2d joint detections," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2017, pp. 805–814.
- [13] C. Li and G. Lee, "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*, 2019, pp. 9887–9895.
- [14] T. Oikarinen, D. Hannah, and S. Kazerounian, "Graphmdn: Leveraging graph structure and deep learning to solve inverse problems," in *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2021, pp. 1–9.
- [15] D. P. Kingma, S. Mohamed, D. Jimenez Rezende, and M. Welling, "Semi-supervised learning with deep generative models," *Advances in neural information processing systems*, vol. 27, 2014.
- [16] Y. Song, C. Durkan, I. Murray, and S. Ermon, "Maximum likelihood training of score-based diffusion models," *Advances in Neural Information Processing Systems*, vol. 34, pp. 1415–1428, 2021.
- [17] C. Li and G. H. Lee, "Weakly supervised generative network for multiple 3d human pose hypotheses," in *BMVC*. BMVA Press, 2020. [Online]. Available: <https://www.bmvc2020-conference.com/assets/papers/0330.pdf>
- [18] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu, "Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments," *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 7, pp. 1325–1339, 2013.
- [19] D. Mehta, H. Rhodin, D. Casas, P. Fua, O. Sotnychenko, W. Xu, and C. Theobalt, "Monocular 3d human pose estimation in the wild using improved cnn supervision," in *2017 international conference on 3D vision (3DV)*. IEEE, 2017, pp. 506–516.
- [20] Y. Chen, Z. Wang, Y. Peng, Z. Zhang, G. Yu, and J. Sun, "Cascaded pyramid network for multi-person pose estimation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7103–7112.
- [21] K. Sun, B. Xiao, D. Liu, and J. Wang, "Deep high-resolution representation learning for human pose estimation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5693–5703.
- [22] Y. Cai, L. Ge, J. Liu, J. Cai, T.-J. Cham, J. Yuan, and N. M. Thalmann, "Exploiting spatial-temporal relationships for 3d pose estimation via graph convolutional networks," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 2272–2281.
- [23] W. Shan, Z. Liu, X. Zhang, Z. Wang, K. Han, S. Wang, S. Ma, and W. Gao, "Diffusion-based 3d human pose estimation with multi-hypothesis aggregation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 14 761–14 771.