

Back to Optimization: Diffusion-based Zero-Shot 3D Human Pose Estimation

Zhongyu Jiang^{1*} Zhuoran Zhou^{1*} Lei Li^{2†} Wenhao Chai¹
 Cheng-Yen Yang¹ Jenq-Neng Hwang¹

University of Washington¹ University of Copenhagen²

{zyjiang, zhouz47, wchai, cycyang, hwang}@uw.edu, lilei@di.ku.dk

Abstract

Learning-based methods have dominated the 3D human pose estimation (HPE) tasks with significantly better performance in most benchmarks than traditional optimization-based methods. Nonetheless, 3D HPE in the wild is still the biggest challenge for learning-based models, whether with 2D-3D lifting, image-to-3D, or diffusion-based methods, since the trained networks implicitly learn camera intrinsic parameters and domain-based 3D human pose distributions and estimate poses by statistical average. On the other hand, the optimization-based methods estimate results case-by-case, which can predict more diverse and sophisticated human poses in the wild. By combining the advantages of optimization-based and learning-based methods, we propose the **Zero-shot Diffusion-based Optimization (ZeDO)** pipeline for 3D HPE to solve the problem of cross-domain and in-the-wild 3D HPE. Our multi-hypothesis **ZeDO** achieves state-of-the-art (SOTA) performance on Human3.6M, with minMPJPE 51.4mm, without training with any 2D-3D or image-3D pairs. Moreover, our single-hypothesis **ZeDO** achieves SOTA performance on 3DPW dataset with PA-MPJPE 40.3mm on cross-dataset evaluation, which even outperforms learning-based methods trained on 3DPW. Our code is available here: <https://github.com/ipl-uw/ZeDO-Release>.

1. Introduction

As people become increasingly interested in Virtual Reality (VR), Augmented Reality (AR), Human-Computer Interaction and Sports Analysis, 3D Human Pose Estimation (HPE) becomes a crucial component for these applications. Compared with multi-view 3D HPE, monocular-based methods are easier to set up and have lower costs, which are more suitable for VR, AR, and mobile devices. Weng *et al.* [60], and Peng *et al.* [40] utilize 3D human poses with Neural Radiance Fields (NeRF) for 3D Human Reconstruction. Meanwhile, Bridgeman *et al.* [2] propose




Figure 1. ZeDO iteratively estimates 3D poses by minimizing the re-projection error via a diffusion-based method.

a 3D HPE and tracking pipeline for soccer analysis, and Jiang *et al.* [19] take advantage of 2D and 3D HPE to track the motion of golf players.

With the availability of more benchmark datasets, deep learning-based 3D HPE methods have been shown to outperform all traditional methods and dominate the areas. Combining 2D HPE with SMPL [30] model, Bogo *et al.* propose SMPLify [1] as an optimization-based 3D HPE pipeline. 2D-3D lifting [7, 39, 61, 64, 65] and diffusion-based 3D HPE [8, 10] networks leverage 3D human poses from the single-frame or multi-frame 2D poses. On the other hand, Image-to-3D networks [22, 24, 29, 52] estimate 3D human poses directly from images without intermediate 2D human poses. However, as mentioned in [9, 11], learning-based 3D HPE methods suffer from performance degradation with cross-domain or in-the-wild scenarios. During training, these methods implicitly learn camera intrinsic parameters, domain-based 3D human pose distributions, or image features in a certain domain. Although optimization-based methods can mitigate the impact of domain gaps by estimating 3D poses case-by-case, their performances are not comparable to learning-based methods at this moment.

To address this problem, Zhan *et al.* [62] decouple the camera intrinsic parameters from 2D-3D lifting networks

* Equal contribution. † Corresponding author.

learning by converting 2D keypoints to 2D rays. Gong *et al.* [11] and Gholami *et al.* [9] generate various 3D poses to bridge domain gaps. Furthermore, Chai *et al.* [5] propose a data augmentation pipeline to minimize the 3D human pose spatial distribution gap. However, the above methods still cannot outperform the learning-based 3D HPE. To decouple camera intrinsic and bridge 3D pose domain gaps simultaneously, we propose the Zero-shot Diffusion-based Optimization (ZeDO) pipeline for 3D human pose estimation, which combines a simple yet effective optimization pipeline with a pre-trained diffusion-based 3D human pose generation model.

Different from traditional optimization-based methods [1, 13, 37, 54], which include various kinematic constraints, the diffusion model denoises the output from the optimization pipeline iteratively to ensure optimized poses following human body constraints. Meanwhile, poses are optimized by minimizing 2D keypoint re-projection errors with a simple yet effective optimization pipeline. Thus, ZeDO is able to estimate 3D human poses without training on any 2D-3D or image-3D pairs. Our contributions are as follows:

- The proposed ZeDO pipeline is a Zero-Shot 3D Human Pose Estimation pipeline, which leverages a pre-trained diffusion-based 3D human pose generation model to optimize target 3D poses in the loop during the inference time.
- Compared with other generation and denoising tasks, we take the diffusion model as an optimization tool by combining a simple 3D HPE optimization pipeline with a pre-trained diffusion-based 3D human pose generation model.
- ZeDO achieves state-of-the-art zero-shot 3D HPE performance on Human3.6M, 3DHP, and 3DPW datasets, even on cross-dataset evaluation.

In Sec 2, we will discuss related works in 3D HPE, including optimization-based and learning-based methods. In Sec 3, details of our backbone architecture and optimization pipeline are addressed. Experimental results are presented in Sec 4, and the ablation studies will be discussed in Sec 5. At last, there are conclusions in Sec 6.

2. Related Works

2.1. Optimization-based Human Pose Estimation

Optimization-based methods, which estimate 3D poses frame-by-frame and case-by-case, are not handicapped by domain gaps or varying camera intrinsic parameters, but their performance so far is much worse than learning-based methods. Utilizing the SMPL [30] model, SMPLify [1] is capable of optimizing the 3D human poses by minimizing the 2D keypoint re-projection error and satisfying lots of

kinematic constraints. Furthermore, Müller *et al.* [37] propose SMPLify-XMC as an improved version of SMPLify with more constraints about the human body and more inputs including height and age. Zheng *et al.* [54] propose an optimization-based hierarchical 3D human pose estimation pipeline that can estimate both 3D human pose and locations at the same time. Recently, Song *et al.* [47] and Choutas *et al.* [6] focus on train an optimizer to fit the SMPL model to estimated 2D human poses.

2.2. 2D-3D Lifting Networks

2D-3D lifting networks employ either a single frame or a sequence of normalized 2D keypoints as input to generate corresponding 3D keypoints [33, 39, 61, 64]. Pavlo *et al.* [39] leverage dilated temporal convolutions with semi-supervised ways to improve 3D pose estimation in videos. Yang *et al.* [61] enhance 3D human pose estimation by leveraging in-the-wild 2D annotations and a novel refinement network module in a weakly-supervised framework. Li *et al.* [26] propose a scalable data augmentation technique that synthesizes unseen 3D human skeletons for training 2D-to-3D networks, effectively reducing dataset bias and improving model generalization to rare poses. These methods, despite the need of two-stage processing to obtain the 2D keypoints in advance, have demonstrated superior performance on several benchmark datasets and are highly efficient, especially when adapted for temporal considerations.

2.3. Image-to-3D Networks

End-to-end 3D HPE methods directly transform image data into 3D pose representations, such as those put forth by Guler *et al.* [14], Tung *et al.* [56], Tan *et al.* [53], and various other research teams including those behind SPIN [24], ROMP [51], BEV [52], and CLIFF [29], who successfully utilize scale and variable height information, effectively resolving issues of depth/height ambiguity. For instance, the methodology introduced by Sun *et al.* [51] is a one-stage process that allows for real-time, monocular 3D mesh recovery of multiple individuals. Further contributing to this field, Sun *et al.* [52] develop a single-shot method capable of simultaneously regressing the pose, shape, and relative depth of multiple people within a single image, utilizing the Bird’s-Eye-View representation for depth reasoning while accommodating variable heights. Within the area of 3D pose estimation from single images, these one-stage techniques consistently demonstrate robust performance, despite their comparatively streamlined architectural designs.

2.4. Diffusion Models in Human Pose Estimation.

Nowadays, Diffusion Probabilistic Model [46] and its descendants [15, 44, 48] have shown their outstanding capabilities of prior distribution learning in multiple areas,

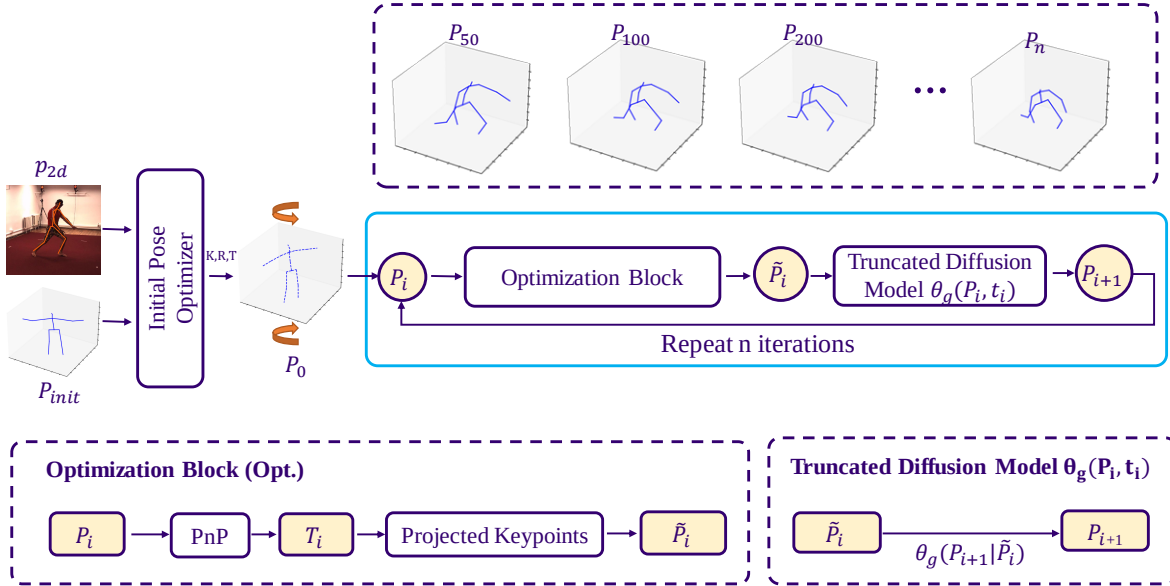


Figure 2. The pipeline of ZeDO, which takes an initial 3D pose, called a hypothesis, as input and estimates the pose by minimizing re-projection error with the target 2D pose. After 1000 iterations, ZeDO is able to generate the optimized 3D human poses.

such as image generation [42] and editing [3, 4], 3D-reconstruction [41], image inpainting [31], and human motion generation [55]. Intuitively, diffusion models are able to tackle depth/scale ambiguity and one-to-many mappings in 3D HPE tasks. Ci *et al.* [8] introduce a novel score-based generative framework to model plausible 3D human poses with a hierarchical condition masking strategy. Meanwhile, Gong *et al.* [10] propose a diffusion-based architecture including the initialization of 3D pose distribution, a GMM-based forward diffusion process, and a conditional reverse diffusion process.

Although learning-based methods achieve significant success in 3D HPE tasks, they are highly constrained by the training datasets. In other words, their performance would abruptly drop when tested on other datasets due to domain gaps. To resolve this issue, we integrate the generality of optimization methods into a robust pre-trained diffusion backbone and propose a Zero-shot Diffusion-based Optimization pipeline (ZeDO) for 3D HPE.

3. Method

As shown in Fig 2, ZeDO includes an initial pose optimizer for rotating initial poses and an optimizer in the loop for iteratively optimizing 3D poses.

Firstly, a randomly selected initial 3D pose (a hypothesis), $P_{init} \in R^{J \times 3}$, is rotated to an optimal pose, P_0 , by minimizing the re-projection error with detected or ground truth 2D keypoints $p_{2d} \in R^{J \times 2}$. Here, J stands for the number of keypoints. Then, in the i th optimization step,

P_i is optimized by the optimizer in the loop, and the pre-trained diffusion model is used to denoise it to P_{i+1} as input of the next iteration. After n iterations, P_n will be the estimated 3D human pose.

Different from other diffusion-based pose estimation methods [8, 10], our diffusion model θ_g is a pose generation model, which is only trained with 3D human poses, and during inference, our diffusion model only takes the optimized pose \tilde{P}_i and timestamp $t(i)$ as input without any additional pose condition information including 2D poses.

3.1. Pre-trained 3D Human Pose Generation Model

We apply the Score Matching [50] on the pre-trained backbone for our 3D human pose generation diffusion model, which rectifies the noisy poses generated after projection to get reasonable 3D poses. During pre-training, the model takes relative-to-pelvis 3D poses $x \in R^{J \times 3}$ as inputs and tries to recover them from recurrent Gaussian noise. In this case, we expect that the diffusion model learns the distribution of real 3D poses and reconstructs $\tilde{x} \in R^{J \times 3}$ to minimize the difference from the inputs. The perturbation strategy used in our Score Matching diffusion model expresses $p(x(t)|x(0))$ in the closed form as:

$$N\left(x(t); x(0)e^{-\frac{1}{2} \int_0^t \beta(s) ds}, [1 - e^{-\int_0^t \beta(s) ds}]^2 I\right). \quad (1)$$

Besides, built upon the learning strategy of the noise conditional score network (NCSN) [49], we formulate our loss

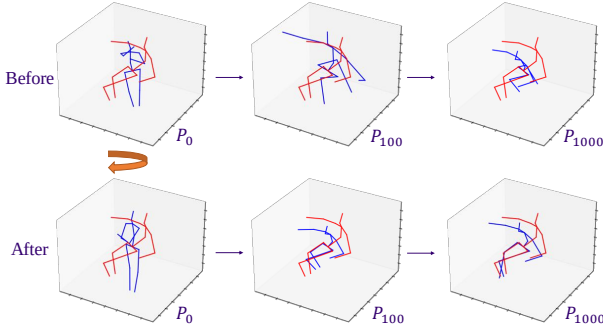


Figure 3. The rotation optimization in the initial pose optimizer finds the optimal initial pose and prevents the collapse of the estimated pose.

function as follows by choosing $\lambda(t) = \sigma(t)^2$:

$$L = E_{U(t;0,1)} \left[\lambda(t) \|s_\theta(x(t), t) + \frac{x(t) - \mu}{\sigma^2}\|_2^2 \right] \quad (2)$$

$$= E_{U(t;0,1)} \left[\|\sigma(t)s_\theta(x(t), t) + z\|_2^2 \right], \quad (3)$$

in which z stands for random noise vector $z \sim N(0, 1)$ and s_θ is the pre-trained score matching network. σ represents the variance mentioned in Eq.(1) as $[1 - e^{-\int_0^t \beta(s)ds}]$ and the timestamp or denoising time variable t is uniformly sampled 1000 times from $(0, 1]$. The 3D pose generation model is never trained with any 2D-3D or image-3D pairs. We report the results based on the Score Matching based model, similar to GFpose [8]. More results with different backbones, including DDIM [48] and DDPM [16], are shown in Table 5.

3.2. Initial Pose Optimizer

Similar to other optimization-based 3D HPE methods [1, 13, 54], our optimization pipeline starts from an initial pose, P_{init} . As depicted in Fig 3, the optimized pose may not suffice if the initial pose's orientation is significantly different from that of the target pose or even perpendicular to that of the target pose. Therefore, the initial pose optimizer is designed to find the optimal rotation matrix, $R_0 \in SO(3)$, and translation, T_0 , of P_{init} by minimizing the re-projection error with 2D keypoints p_{2d} .

$$\arg \min_{R_0, T_0} \|K(R_0 P_{init} + T_0) - p_{2d}\|_2 \quad (1)$$

$$\text{s.t. } T_{min} \leq T_0 \leq T_{max}, \quad (2)$$

where K is the camera intrinsic matrix. After the rotation, the $P_0 = R_0 P_{init}$ is the optimal pose sent to the iterative optimization pipeline. As shown in Fig 3, the rotation optimization aligns the initial poses with target 2D and 3D poses.

Algorithm 1 ZeDO pipeline

Require: Initial 3D pose P_{init} , Target 2D pose p_{2d} , 2D pose confidence scores C_{2d} , Camera intrinsic K , Diffusion timestamp t , Pre-trained diffusion model $\theta_g(P, t)$
 $R_0, T_0 \leftarrow \arg \min_{R_0, T_0} \|K(R_0 P_{init} + T_0) - p_{2d}\|_2$
// Initial Pose Optimization
 $P_0 \leftarrow R_0 P_{init}$
// Iterative Optimization and Denoising
 $r \leftarrow K^{-1} p_{2d}$
 $\hat{r} \leftarrow \frac{r}{\|r\|_2}$
for $i \leftarrow 0$ to $n - 1$ **do**
 if $i < \text{warmup}$ **then**
 $T_i \leftarrow T_0$
 else
 $T_i \leftarrow \arg \min_{T_i} \|C_{2d}(K(P_i + T_i) - p_{2d})\|_2$
 end if
 // Project 3D keypoints to rays
 $\tilde{P}_i \leftarrow ((P_i + T_i) \cdot \hat{r}) \hat{r} - T_i$
 $P_{i+1} \leftarrow \theta_g(\tilde{P}_i, t(i))$
end for
return P_n

3.3. Optimizer in the Loop

In previous works [1, 12, 37, 54], to optimize an accurate 3D human pose, it requires a lot of kinematic constraints, which call for a strong domain knowledge about human motion. Different from those previous works, as shown in Alg 1, our iterative optimizer utilizes the denoising capability and learned human pose prior of the pose generation diffusion model to estimate an accurate 3D human pose with a simple yet effective optimization pipeline without any explicit kinematic constraint.

With a camera intrinsic matrix, K , 2D keypoints, p_{2d} , can be converted to 3D rays, $r \in R^{J \times 3}$, based on perspective projection,

$$r = K^{-1} p_{2d}, \hat{r} = \frac{r}{\|r\|_2}. \quad (3)$$

Intuitively, as shown in Fig 44, projecting the 3D keypoints from P_0 to r will minimize the 2D re-projection error and provide the estimated 3D human poses,

$$\tilde{P}_0 \leftarrow ((P_0 + T) \cdot \hat{r}) \hat{r} - T. \quad (4)$$

However, there are two problems: 1) Simply projecting 3D keypoints from P_0 to r generates a noisy 3D human pose, which may not satisfy the kinematic constraints of the human body. 2) With different translations between P_0 and the camera, there are different projection results.

In order to solve these two problems, we need to ensure the estimated 3D poses, P_i , are valid poses and inherently




Figure 4. By projecting the P_i to r , we minimize the 2D re-projection error and find an optimized pose \tilde{P}_i .

follow the kinematic constraints to find the optimal translation, T_i . This calls for our use of pose prior.

Pose prior Although there is no re-projection error from the optimized poses, the optimized poses may not satisfy the kinematic constraints. Therefore, in previous works, kinematic constraints are added to the pose optimization pipeline, and a complex joint optimization problem is designed to find optimal poses. However, in our pipeline, we take advantage of a pre-trained diffusion-based pose generation model to ‘denoise’ our optimized poses. As mentioned in DDPM [16], DDIM [48] and Score Matching network [50], the diffusion model is trained by maximizing likelihood, which aims at finding the most possible valid pose based on the input noisy pose. As a result, we use the diffusion-based pose generation model to find the optimal P_{i+1} based on optimized pose \tilde{P}_i ,

$$P_{i+1} = \theta_g(\tilde{P}_i, t(i)), \quad (5)$$

which is different from other diffusion-based methods, like GFpose [8] and DiffPose [10], $P = \theta(x, c, t)$, where x is random noise, c is pose condition and t is timestamp.

However, during training, the reverse diffusion starts from the standard Gaussian noise, $\mathcal{N}(0, 1)$, but in our case, the generation model is utilized to denoise an optimized pose, which does not follow the standard Gaussian noise. Inspired by [36, 66], we adopt truncated diffusion model inference, whose timestamp is truncated from the training timestamp during inference. In our case, the timestamp, t , is truncated as $t \in (0, 0.1]$, instead of $(0, 1]$.

Find optimal translations Optimal translations are derived from T_o by minimizing the 2D re-projection error of P_i , depending on the current iteration number, since the optimized pose, P_i , is not reliable enough in the early iterations.

After certain iterations, the optimal translation is derived from the following:

$$T_i \leftarrow \arg \min_{T_i} \left\| C_{2d} \left(K(P_i + T_i) - p_{2d} \right) \right\|_2, \quad (6)$$

Methods	CE	Opt	PA-MPJPE ↓	MPJPE ↓
Kolotouros <i>et al.</i> [24]			59.2	96.9
Kocabas <i>et al.</i> [22]			51.9	82.9
Kocabas <i>et al.</i> [23]			46.4	74.7
Li <i>et al.</i> [25]			45.0	74.1
Ma <i>et al.</i> [32]			41.3	67.5
Li <i>et al.</i> [25]	✓		50.9	82.0
Kocabas <i>et al.</i> [22]	✓		56.5	93.5
Kocabas <i>et al.</i> [23]	✓		50.9	82.0
Gong <i>et al.</i> [11]	✓		58.5	94.1
Gholami <i>et al.</i> [9]	✓		46.5	81.2
Chai <i>et al.</i> [5]	✓		55.3	87.7
Song <i>et al.</i> [47]	✓		55.9	-
Choutas <i>et al.</i> [6]	✓		52.2	-
ZeDO ($S = 1, J = 17$)	✓	✓	40.3	69.7
ZeDO ($S = 1, J = 14$)	✓	✓	43.1	76.6

Table 1. Cross-domain evaluation results on 3DPW dataset. CE stands for cross-domain evaluation, and Opt means optimization-based method. Ground truth 2D poses are used.

where $C_{2d} \in R^J$ are the confidence scores of 2D keypoints, p_{2d} and K are the camera intrinsic matrices. Inspired by [18], C_{2d} are helpful to guide the translation optimization. There is a closed-form solution to the optimization problem. The details can be found in the supplementary material.

By incorporating optimal translations and kinematic constraints, we can optimize the 3D human poses by iteratively projecting P_i to r and denoising the projected P_{i+1} with the help of the diffusion model.

4. Experimental Results

In this section, we will introduce the experimental results of ZeDO on 3DPW [57], Human3.6M [17] and MPI-INF-3DHP [34] datasets. More results on Ski-Pose [43] datasets are included in the supplementary materials. Since ZeDO requires an initial pose as a starting point for optimization, and different initial poses may lead to different pose estimation results, we report our results in Mean Per Joint Position Error (MPJPE) with a single initial pose (single-hypothesis) or minimum MPJPE (minMPJPE) with multiple initial poses, in order to have fair comparison with previous multi-hypothesis 3D HPE methods [8, 45, 59].

4.1. Datasets

Human3.6M [17] is the most widely used single-person 3D pose benchmark with more than 3.6 million frames and corresponding 3D human poses. The dataset is collected within a $4 \text{ m} \times 3 \text{ m}$ indoor environment, with 11 professional actors (6 males and 5 females) performing 17 distinct actions such as discussion, smoking, capturing photographs, posing, greeting, and talking on the phone. Following the convention of previous works [26, 39, 64] for fair comparison,

Learning Methods	MPJPE ↓	PA-MPJPE ↓
Martinez <i>et al.</i> [33]	62.9	47.7
Zhao <i>et al.</i> [64]	57.6	-
Pavlo <i>et al.</i> [39] ($f = 1$)	52.7	40.9
Li <i>et al.</i> [29]	47.1	32.7
Gong <i>et al.</i> [11]	50.2	39.1
Gong <i>et al.</i> [10] ($f = 1$)	49.7	<u>31.6</u>
Ci <i>et al.</i> [8] ($S = 1$)	51.0	-
Ci <i>et al.</i> [8] ($S = 10$)	45.1	30.5
Optimization Methods	MPJPE ↓	PA-MPJPE ↓
Wang <i>et al.</i> [58]	88.0	-
Bogo <i>et al.</i> [1]	82.3	-
Li <i>et al.</i> [27]	78.6	-
Gu <i>et al.</i> [12]	77.2	-
Song <i>et al.</i> [47]	-	56.4
ZeDO ($S = 1$)	65.7	49.0
ZeDO ($S = 10$)	<u>57.3</u>	<u>45.1</u>
ZeDO ($S = 50$)	51.4	42.1

Table 2. 3D HPE quantitative results on Human3.6M dataset. S indicates the number of hypotheses. All results are reported in millimeters (mm). The pose generation model is trained on Human3.6M. Detected 2D poses by Stacked Hourglass are used.

we use the S1, 5, 6, 7, and 8 as the training dataset and evaluate the model on S9 and S11.

MPI-INF-3DHP [34] is a large 3D human pose dataset with more than 1.3 million frames captured indoors and outdoors. The 3DHP dataset captures poses of 8 actors, consisting of 4 males and 4 females with 8 different actions each, encompassing a range of activities from simple actions like walking and sitting to more complex exercise poses and dynamic movements. Following [11], we use a sampled 2929 frame test dataset.

3DPW [57] is the first dataset in-the-wild with accurate 3D poses for evaluation. Compared with Human3.6M and MPI-INF-3DHP, 3DPW focuses on outdoor scenarios and captures videos with static and moving cameras. There are 60 video sequences captured in the dataset with 18 different actors. Following [5], we test ZeDO on 3DPW only for cross-dataset evaluation.

4.2. Training and Inference Details

We pre-train our 3D pose generation model for 5000 epochs on one NVIDIA A100 with a batch size of 50k, a learning rate of $2e^{-4}$ with an Adam optimizer. The training schedule comes with warmup in the first 5k iterations and cosine learning rate decay in the following iterations. As [50], the timestamp, t , during the forward or reverse diffusion process, is uniformly sampled from $(0, 1]$. All 3D human poses are normalized to pelvis-related coordinates

Methods	CE	Opt	MPJPE ↓	PCK ↑	AUC ↑
Mehta <i>et al.</i> [35]			124.7	76.6	40.4
Martinez <i>et al.</i> [33]			84.3	85.0	52.0
Pavlo <i>et al.</i> [39] ($f = 1$)			86.6	-	-
Li <i>et al.</i> [28] ($f = 9$)			58.0	93.8	63.3
Zhang <i>et al.</i> [63] ($f = 1$)			57.9	94.2	63.8
ZeDO ($S = 1$)		✓	86.5	82.6	53.8
ZeDO ($S = 50$)		✓	55.2	93.0	65.6
Kanazawa <i>et al.</i> [21]	✓		113.2	77.1	40.7
Ci <i>et al.</i> [8]	✓		-	86.9	-
Gong <i>et al.</i> [11]	✓		73.0	88.6	57.3
Gholami <i>et al.</i> [9]	✓		68.3	90.2	59.0
Chai <i>et al.</i> [5]	✓		61.3	92.1	62.5
Müller <i>et al.</i> [37]	✓	✓	101.2	-	-
ZeDO ($S = 1$)	✓	✓	99.9	81.8	50.9
ZeDO ($S = 50$)	✓	✓	<u>69.9</u>	90.2	58.8

Table 3. 3D HPE quantitative results on 3DHP dataset. CE stands for cross-domain evaluation, and Opt means optimization-based method. Ground truth 2D poses are used.

during training and inference. To improve the robustness of the model, we apply flip and rotation data augmentation during training. For cross-domain evaluation, we pre-train the pose generation model in a different dataset and directly test the optimization pipeline without any fine-tuning.

During inference, the pipeline supports single or multiple initial poses. For the initial pose optimizer, we limit the rotation axis to the z-axis only for better performance. We set the number of warmup iterations as 200, and the number of total iterations as 1000. The timestamp, t , is uniformly sampled from $(0, 0.1]$. The initial poses are sampled from training sets of Human3.6M [17] or 3DHP [34] by the K-Means algorithm. For different numbers of hypotheses, we run K-Means with different numbers of clusters.

4.3. Results

Results on 3DPW. 3DPW is a challenging in-the-wild dataset, compared with Human3.6M and MPI-INF-3DHP datasets. 3DPW focuses on outdoor scenarios with both static and moving cameras. For cross-domain evaluation on the 3DPW dataset, we pre-train the pose generation model on the Human3.6M dataset and inference on the 3DPW dataset without any fine-tuning. On 3DPW, we find some of the previous works [23, 25] evaluate on 14 Leeds Sports Pose (LSP) [20] keypoints, while others [5, 9, 11] evaluate on 17 Human3.6M keypoints, and some other works [24] do not explain clearly which keypoints they use. In Table 1, we report the results of both 14 and 17 keypoints for a fair comparison. We achieve SOTA performance, PA-MPJPE 40.3mm. with a single hypothesis.

Results on Human3.6M. Although ZeDO is a Zero-shot Diffusion-based Optimization pipeline, ZeDO achieves

Dataset	Diff Model	RO	WU	RA	GT	$S = 1$		$S = 50$	
						MPJPE ↓	PA-MPJPE ↓	MPJPE ↓	PA-MPJPE ↓
H36M	H36M					75.0	52.7	53.4	42.7
H36M	H36M	✓				77.2	53.7	52.7	42.4
H36M	H36M	✓	✓			65.7 (9.3 ↓)	49.0 (3.7 ↓)	51.4 (2.0 ↓)	42.1 (0.6 ↓)
H36M	H36M	✓	✓	✓		69.5	51.4	52.9	42.5
H36M	H36M	✓	✓		✓	50.1	35.8	37.0	27.5
3DHP	H36M				✓	148.3	88.8	93.4	59.0
3DHP	H36M	✓	✓		✓	113.8	74.1	80.1	56.0
3DHP	H36M	✓	✓	✓	✓	99.9 (48.4 ↓)	67.9 (20.9 ↓)	69.9 (23.5 ↓)	49.0 (10.0 ↓)
3DHP	3DHP	✓	✓		✓	86.5	55.9	55.2	38.6

Table 4. The ablation study results of ZeDO. RO stands for rotation optimization as the initial pose optimization. WU denotes the warmup iterations. RA is the rotation data augmentation for training the pose generation model. The dataset name under the Dataset column is the testing dataset, and the name under the Diff Model column is the dataset used for diffusion model pre-training.

comparable results with learning-based. Following [8], we report the minMPJPE of multi-hypothesis, i.e., S number of initial poses, in inference and use the detected 2D poses as input. As shown in Table 2, on the Human3.6M dataset, we obtain 51.4mm in minMPJPE with $S = 50$, which is comparable with SOTA learning-based methods, while ZeDO does not train with any 2D-3D or image-3D pairs. Compared with other optimization-based methods, ZeDO outperforms previous works by a large margin, even with $S = 1$. In this experiment, we train the pose generation model on the training set of Human3.6M.

Results on 3DHP. Following previous works [5, 9, 11], we use ground truth 2D poses as input. As shown in Table 3, with $S = 50$, ZeDO even outperforms the learning-based methods by 2.7mm in minMPJPE. In the cross-domain evaluation, we achieve SOTA performance as 67.9mm in minMPJPE and outperform the optimization-based method by a large margin, with $S = 50$.

5. Discussion

5.1. Ablation Studies

Different diffusion-based pose generation models. As shown in table 5, we try to evaluate the cross-domain performance of our pipeline with different diffusion-based backbones on the 3DPW. We test the Score Matching Network [50], DDPM [16], and DDIM [48] models trained on the Human3.6M dataset and keep all other settings the same. It turns out that DDIM also achieves comparable performance in terms of PA-MPJPE and even lower MPJPE compared with the Score Matching Network we report above. The outcome validates the generality and viability of our idea, regardless of the specific structure of the diffusion backbone.

How does initial pose optimizer help ZeDO? The initial pose optimizer is designed to align the initial pose with the target 2D pose by rotation for better initialization. As

Diffusion Backbone	PA-MPJPE ↓	MPJPE ↓
Score Matching [50]	40.3	<u>69.7</u>
DDIM [48]	40.4	67.9
DDPM [38]	51.7	81.3

Table 5. Different diffusion backbone 3D HPE quantitative results on 3DPW dataset.

shown in Table 4, the combination of rotation optimization and warmup iterations reduces the MPJPE by 9.3 mm when $S = 1$ and the minMPJPE by 2.0mm when $S = 50$, on the Human3.6M dataset. When $S = 50$, hypotheses cover lots of different pose orientations, resulting in the relatively smaller improved performance from the initial pose optimizer when S is larger. On the 3DHP dataset, the initial optimization further improves the performance by 34.5 mm in MPJPE when $S = 1$ since 3DHP contains more complex 3D human poses in different orientations than Human3.6M. The initial pose optimization is able to generate a reliable optimized initial pose and an optimal translation as the warmup translation for the following iterative optimization pipeline.

Does data augmentation help the performance? In ZeDO, the diffusion model is pre-trained for pose genera-

Dataset	Diff Model	MPJPE ↓	PA-MPJPE ↓
H36M	H36M	37.0	27.5
H36M	mixed	35.7	26.5
3DHP	H36M	69.9	49.0
3DHP	3DHP	55.2	38.6
3DHP	mixed	52.4	37.7

Table 6. The pose generation models trained on mixed datasets (Human3.6M + 3DHP) achieves the best performance with $S = 50$, as well as better generalization. GT 2D keypoints are used.

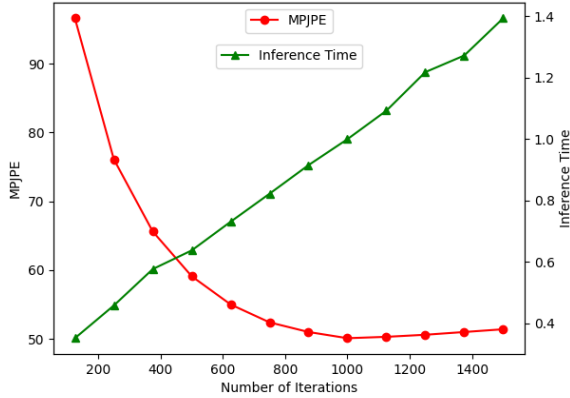


Figure 5. The MPJPE and inference time with different numbers of iterations on the Human3.6M dataset with ground truth 2D keypoints.

tion. However, the 3D pose distributions vary across different datasets. To ensure the pre-trained diffusion model can be adapted to different datasets, we utilize rotation and flip data augmentation during training. As expected, the data augmentation significantly improves the performance in cross-domain evaluation by 13.9 mm in MPJPE, shown in Table 4.

Boost the performance further by mixing dataset. According to Table 6, the pose generation model trained on the mixed datasets (Human3.6M + 3DHP) improves the performance of ZeDO by 1.3mm on Human3.6M and 2.8mm on 3DHP while using the same estimation algorithm.

How to pick the initial pose? We utilize K-Means in our experiments to select anchor poses from the training set as initial poses. K-Means effectively finds the most representative poses in the training set, making it superior to other sampling strategies, as shown in Table 7. Random Sampling randomly samples a pose from the training set, whereas Random Generation is generated by the pre-trained pose generation model.

What is the best number of optimization iterations? In ZeDO, the number of diffusion optimization iterations is set to 1000. Intuitively, increasing the number of iterations can enhance performance but may suffer the inference speed. In Fig 5, as expected, the inference time increases linearly with respect to the number of iterations. However, the fig-

Sampling	MPJPE ↓	PA-MPJPE ↓
Random Sampling	78.2	51.2
Random Generation	70.4	46.0
K-Means	50.1	35.8

Table 7. Results on Human3.6M dataset with different sampling strategies when $S = 1$. Ground truth 2D keypoints are used.

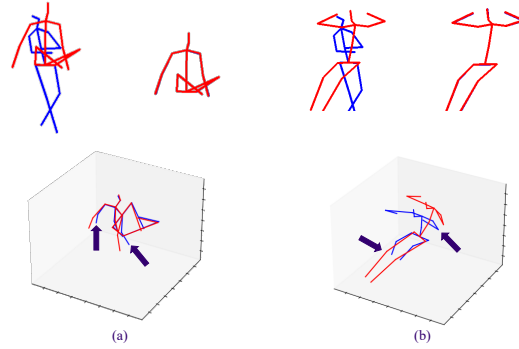


Figure 6. The failure cases of our method, because of the one-to-many issue in 3D HPE.

ure shows that the best performance is achieved when the number of iterations is around 1000. With the number of iterations exceeding 1000, there is no performance gain, and the inference speed decreases.

5.2. Limitations

Although ZeDO achieves state-of-the-art performance in various benchmarks and settings, several limitations still need to be further explored. 1) Similar to other optimization-based approaches, the optimizer in the loop requires camera intrinsic parameters. 2) Since our method is based on minimizing the 2D re-projection error, we are not able to solve the ambiguity of the depth and scale without additional information like bone length or height. 3) For an identical 2D human pose, there are multiple 3D human poses matched. Without image or temporal information, the 1-to-many mapping issue cannot be resolved by single frame lifting methods, as shown in Fig 6.

6. Conclusion

In this paper, we propose ZeDO, a Zero-shot Diffusion-based Optimization pipeline for 3D HPE. To the best of our knowledge, we are the first to introduce the diffusion model to the optimization-based method in the 3D HPE task. We leverage the pre-trained diffusion-based 3D human pose generation model and can optimize target 3D poses in the loop. To be specific, an optimizer that calculates the optimal translation is used iteratively with denoising steps in the diffusion model. Compared to other prior arts, ZeDO achieves state-of-the-art performance on Human3.6M, MPI-INF-3DHP, and 3DPW datasets, even with cross-dataset evaluation. In the future, we plan to further improve ZeDO by modifying the diffusion model and solving the limitations listed in Sec 5. It is our wish that this optimization method could become a common paradigm beyond end-to-end lifting networks in 3D human pose estimation tasks.

References

- [1] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 561–578. Springer, 2016. 1, 2, 4, 6
- [2] Lewis Bridgeman, Marco Volino, Jean-Yves Guillemaut, and Adrian Hilton. Multi-person 3d pose estimation and tracking in sports. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019. 1
- [3] Shidong Cao, Wenhao Chai, Shengyu Hao, and Gaoang Wang. Image reference-guided fashion design with structure-aware transfer by diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3524–3528, 2023. 3
- [4] Shidong Cao, Wenhao Chai, Shengyu Hao, Yanting Zhang, Hangyue Chen, and Gaoang Wang. Diffashion: Reference-based fashion design with structure-aware transfer by diffusion models. *arXiv preprint arXiv:2302.06826*, 2023. 3
- [5] Wenhao Chai, Zhongyu Jiang, Jenq-Neng Hwang, and Gaoang Wang. Global adaptation meets local generalization: Unsupervised domain adaptation for 3d human pose estimation. *arXiv preprint arXiv:2303.16456*, 2023. 2, 5, 6, 7
- [6] Vasileios Choutas, Federica Bogo, Jingjing Shen, and Julien Valentin. Learning to fit morphable models. In *European Conference on Computer Vision*, pages 160–179. Springer, 2022. 2, 5
- [7] Hai Ci, Xiaoxuan Ma, Chunyu Wang, and Yizhou Wang. Locally connected network for monocular 3d human pose estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(3):1429–1442, 2020. 1
- [8] Hai Ci, Mingdong Wu, Wentao Zhu, Xiaoxuan Ma, Hao Dong, Fangwei Zhong, and Yizhou Wang. Gfpose: Learning 3d human pose prior with gradient fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4800–4810, 2023. 1, 3, 4, 5, 6, 7
- [9] Mohsen Gholami, Bastian Wandt, Helge Rhodin, Rabab Ward, and Z Jane Wang. Adaptpose: Cross-dataset adaptation for 3d human pose estimation by learnable motion generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13075–13085, 2022. 1, 2, 5, 6, 7
- [10] Jia Gong, Lin Geng Foo, Zhipeng Fan, Qihong Ke, Hossein Rahmani, and Jun Liu. Diffpose: Toward more reliable 3d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13041–13051, 2023. 1, 3, 5, 6
- [11] Kehong Gong, Jianfeng Zhang, and Jiashi Feng. Poseaug: A differentiable pose augmentation framework for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8575–8584, 2021. 1, 2, 5, 6, 7
- [12] Renshu Gu. *Towards Multi-Person 3D Pose Estimation in Natural Videos*. University of Washington, 2020. 4, 6
- [13] Renshu Gu, Zhongyu Jiang, Gaoang Wang, Kevin McQuade, and Jenq-Neng Hwang. Unsupervised universal hierarchical multi-person 3d pose estimation for natural scenes. *Multimedia Tools and Applications*, 81(23):32883–32906, 2022. 2, 4
- [14] Riza Alp Guler and Iasonas Kokkinos. Holopose: Holistic 3d human reconstruction in-the-wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10884–10894, 2019. 2
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020. 4, 5, 7
- [17] 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 transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. 5, 6
- [18] Karim Isakov, Egor Burkov, Victor Lempitsky, and Yury Malkov. Learnable triangulation of human pose. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7718–7727, 2019. 5
- [19] Zhongyu Jiang, Haorui Ji, Samuel Menaker, and Jenq-Neng Hwang. Golfpose: Golf swing analyses with a monocular camera based human pose estimation. In *2022 IEEE International Conference on Multimedia and Expo Workshops (ICMEW)*, pages 1–6. IEEE, 2022. 1
- [20] Sam Johnson and Mark Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *bmvc*, volume 2, page 5. Aberystwyth, UK, 2010. 6
- [21] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7122–7131, 2018. 6
- [22] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5253–5263, 2020. 1, 5
- [23] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11127–11137, 2021. 5, 6
- [24] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2252–2261, 2019. 1, 2, 5, 6
- [25] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3383–3393, 2021. 5, 6
- [26] Shichao Li, Lei Ke, Kevin Pratama, Yu-Wing Tai, Chi-Keung Tang, and Kwang-Ting Cheng. Cascaded deep monocular 3d

- human pose estimation with evolutionary training data. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2, 5
- [27] Sijin Li, Weichen Zhang, and Antoni B Chan. Maximum-margin structured learning with deep networks for 3d human pose estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 2848–2856, 2015. 6
- [28] Wenhao Li, Hong Liu, Hao Tang, Pichao Wang, and Luc Van Gool. Mhformer: Multi-hypothesis transformer for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13147–13156, 2022. 6
- [29] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. Cliff: Carrying location information in full frames into human pose and shape estimation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V*, pages 590–606. Springer, 2022. 1, 2, 6
- [30] 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. 1, 2
- [31] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022. 3
- [32] Xiaoxuan Ma, Jiajun Su, Chunyu Wang, Wentao Zhu, and Yizhou Wang. 3d human mesh estimation from virtual markers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 534–543, 2023. 5
- [33] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 2640–2649, 2017. 2, 6
- [34] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *3D Vision (3DV), 2017 Fifth International Conference on*. IEEE, 2017. 5, 6
- [35] 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 transactions on graphics (tog)*, 36(4):1–14, 2017. 6
- [36] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2021. 5
- [37] Lea Muller, Ahmed AA Osman, Siyu Tang, Chun-Hao P Huang, and Michael J Black. On self-contact and human pose. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9990–9999, 2021. 2, 4, 6
- [38] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. 7
- [39] Dario Pavllo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 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*, pages 7753–7762, 2019. 1, 2, 5, 6
- [40] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021. 1
- [41] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 3
- [42] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. 3
- [43] Helge Rhodin, Jörg Spörr, Isinsu Katircioglu, Victor Constantin, Frédéric Meyer, Erich Müller, Mathieu Salzmann, and Pascal Fua. Learning monocular 3d human pose estimation from multi-view images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8437–8446, 2018. 5
- [44] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 2
- [45] Saurabh Sharma, Pavan Teja Varigonda, Prashast Bindal, Abhishek Sharma, and Arjun Jain. Monocular 3d human pose estimation by generation and ordinal ranking. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2325–2334, 2019. 5
- [46] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 2256–2265, 2015. 2
- [47] Jie Song, Xu Chen, and Otmar Hilliges. Human body model fitting by learned gradient descent. In *European Conference on Computer Vision*, pages 744–760. Springer, 2020. 2, 5, 6
- [48] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2, 4, 5, 7
- [49] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019. 3
- [50] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 3, 5, 6, 7

- [51] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Black Michael J., and Tao Mei. Monocular, one-stage, regression of multiple 3d people. In *ICCV*, 2021. 2
- [52] Yu Sun, Wu Liu, Qian Bao, Yili Fu, Tao Mei, and Michael J Black. Putting people in their place: Monocular regression of 3d people in depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13243–13252, 2022. 1, 2
- [53] Jun Kai Vince Tan, Ignas Budvytis, and Roberto Cipolla. Indirect deep structured learning for 3d human body shape and pose prediction. 2017. 2
- [54] Zheng Tang, Renshu Gu, and Jenq-Neng Hwang. Joint multi-view people tracking and pose estimation for 3d scene reconstruction. In *2018 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2018. 2, 4
- [55] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. 3
- [56] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. *Advances in Neural Information Processing Systems*, 30, 2017. 2
- [57] Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *European Conference on Computer Vision (ECCV)*, sep 2018. 5, 6
- [58] Chunyu Wang, Yizhou Wang, Zhouchen Lin, Alan L Yuille, and Wen Gao. Robust estimation of 3d human poses from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2361–2368, 2014. 6
- [59] Tom Wehrbein, Marco Rudolph, Bodo Rosenhahn, and Bastian Wandt. Probabilistic monocular 3d human pose estimation with normalizing flows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11199–11208, 2021. 5
- [60] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Humannerf: Free-viewpoint rendering of moving people from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16210–16220, 2022. 1
- [61] Cheng-Yen Yang, Jiajia Luo, Lu Xia, Yuyin Sun, Nan Qiao, Ke Zhang, Zhongyu Jiang, Jenq-Neng Hwang, and Cheng-Hao Kuo. Camerapose: Weakly-supervised monocular 3d human pose estimation by leveraging in-the-wild 2d annotations. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2924–2933, 2023. 1, 2
- [62] Yu Zhan, Fenghai Li, Renliang Weng, and Wongun Choi. Ray3d: ray-based 3d human pose estimation for monocular absolute 3d localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13116–13125, 2022. 1
- [63] Jinlu Zhang, Zhigang Tu, Jianyu Yang, Yujin Chen, and Junsong 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*, pages 13232–13242, 2022. 6
- [64] Long Zhao, Xi Peng, Yu Tian, Mubbasir Kapadia, and Dimitris N Metaxas. Semantic graph convolutional networks for 3d human pose regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3425–3435, 2019. 1, 2, 5, 6
- [65] Ce Zheng, Sijie Zhu, Matias Mendieta, Taojiannan Yang, Chen Chen, and Zhengming Ding. 3d human pose estimation with spatial and temporal transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11656–11665, 2021. 1
- [66] Huangjie Zheng, Pengcheng He, Weizhu Chen, and Mingyuan Zhou. Truncated diffusion probabilistic models and diffusion-based adversarial auto-encoders. *arXiv preprint arXiv:2202.09671*, 2022. 5