

# Pre-training a Density-Aware Pose Transformer for Robust LiDAR-based 3D Human Pose Estimation

Xiaoqi An<sup>1</sup>, Lin Zhao<sup>1\*</sup>, Chen Gong<sup>1</sup>, Jun Li<sup>1</sup>, Jian Yang<sup>1\*</sup>

<sup>1</sup>PCA Lab, Key Lab of Intelligent Perception and Systems for High-Dimensional Information of Ministry of Education  
School of Computer Science and Engineering, Nanjing University of Science and Technology  
{xiaoqi.an, linzhao, chen.gong, junli, csjyang}@njust.edu.cn,

## Abstract

With the rapid development of autonomous driving, LiDAR-based 3D Human Pose Estimation (3D HPE) is becoming a research focus. However, due to the noise and sparsity of LiDAR-captured point clouds, robust human pose estimation remains challenging. Most of the existing methods use temporal information, multi-modal fusion, or SMPL optimization to correct biased results. In this work, we try to obtain sufficient information for 3D HPE only by modeling the intrinsic properties of low-quality point clouds. Hence, a simple yet powerful method is proposed, which provides insights both on modeling and augmentation of point clouds. Specifically, we first propose a concise and effective density-aware pose transformer (DAPT) to get stable keypoint representations. By using a set of joint anchors and a carefully designed exchange module, valid information is extracted from point clouds with different densities. Then 1D heatmaps are utilized to represent the precise locations of the keypoints. Secondly, a comprehensive LiDAR human synthesis and augmentation method is proposed to pre-train the model, enabling it to acquire a better human body prior. We increase the diversity of point clouds by randomly sampling human positions and orientations and by simulating occlusions through the addition of laser-level masks. Extensive experiments have been conducted on multiple datasets, including IMU-annotated LidarHuman26M, SLOPER4D, and manually annotated Waymo Open Dataset v2.0 (Waymo), HumanM3. Our method demonstrates SOTA performance in all scenarios. In particular, compared with LPFormer on Waymo, we reduce the average MPJPE by 10.0mm. Compared with PRN on SLOPER4D, we notably reduce the average MPJPE by 20.7mm.

**Code** — <https://github.com/AnxQ/dapt>

## 1 Introduction

3D human pose estimation (3D HPE) is a fundamental computer vision task with a wide range of downstream usages such as human behavior understanding, trajectory prediction, autonomous driving (Cong et al. 2022; Lian et al. 2022), etc. To implement 3D HPE, a simple and direct way is to regress 3D keypoint coordinates directly from 2D HPE




Figure 1: We propose a novel framework for LiDAR-based 3D HPE. With comprehensive LiDAR human synthesis & augmentation for model pre-training as well as the learning of stable joint representations, our method produces robust results from low-quality point clouds.

results (Kang et al. 2023; Gong, Zhang, and Feng 2021). However, these methods have difficulty predicting world coordinates, and robust pose estimation in the real scenario remains challenging.

To get global 3D keypoints, most traditional 3D HPE methods are based on multi-view RGB images (Simon et al. 2017; Isakov et al. 2019; Tu, Wang, and Zeng 2020; Ye et al. 2022) or RGB-D images (Ying and Zhao 2021; Hong and Kim 2018), which require in-door laboratory environments with complex calibrations (Su et al. 2020; Zheng et al. 2022). Alternatively, LiDAR sensors can obtain accurate point-level depth in complex environments, which are more adaptable to long-range 3D HPE in open-world scenes.

However, compared to clear and dense point clouds generated by RGB-D images (Fan et al. 2018; Ionescu et al. 2014), the LiDAR-captured point clouds have various inconsistencies, making them difficult to learn directly with existing methods. As shown in Fig.2, the samples in Waymo Open Dataset (Sun et al. 2020) have various point densities, and the noisy points from the environment may lead to ambiguity and unstable predictions. Therefore, to achieve robust pose estimation, (Li et al. 2022a; Yan et al. 2024; Zhao et al. 2024) integrate ST-GCNs to get stabilized pose results from multiple LiDAR frames. (Cong et al. 2023; Fürst et al. 2021) perform multi-modal fusion to mitigate the lack of information in sparse point clouds, and (Li et al. 2022a; Yan et al. 2024; Zhang et al. 2024) introduce SMPL (Loper et al.

\*Corresponding authors.




Figure 2: Difficulties of LiDAR-based 3D HPE, mainly lie in samples with noisy, occluded, or sparse point clouds.




Figure 3: Results of the segment-regression-based methods. The color of the point indicates which joint it belongs to. The wrong segmentation leads to a biased joint location.

2015) as a priori to further align human mesh with point clouds. Despite the success of these methods, they inevitably introduce additional data acquisition or time-consuming optimization processes which complicate the entire framework and make it less suitable for practical applications.

In this paper, we propose a novel framework to learn stable representations for 3D HPE only using single-frame low-quality LiDAR point clouds. Our method is simple yet effective. Specifically, We provide insights into both model design for accurate pose estimation and model pre-training for effective body prior learning.

Firstly, we design a **Density-Aware Pose Transformer (DAPT)** that provides stable and explicit representations of joints. Most existing models attempt to regress the coordinates of joints from global or clustered point-wise features. This makes the model highly susceptible to the density and spatial distribution of the point cloud. As shown in Fig.3, when the point cloud near the joint is sparse or noisy, the network cannot correctly identify which body part the point belongs to, resulting in a biased joint location. To solve this problem, we introduce a set of learnable joint anchors. When extracting point cloud features, they can explicitly integrate the information across multiple density levels through a carefully designed exchange module. Then, we utilize 1D-heatmaps on the XYZ axes to represent joint positions, which allows the model to obtain a stable output.

Secondly, we introduce a comprehensive pre-training approach that conducts thorough LiDAR human synthesis and augmentation. Since annotating LiDAR data is expensive, inspired by (Weng et al. 2023a), we train on samples synthesized by SMPL mesh under ray casting. To tackle noisy and sparse point clouds in real scenarios, we add a square surface that can be varied in small magnitudes, such as the

ground surface, and put the human mesh into the scene with randomized poses and positions. On the other hand, since the light beams are highly susceptible to occlusion by foreground objects, we mask the range image with patches to simulate the occlusion. Therefore, the model will learn prior knowledge of the human body and mine important clues about the pose in low-quality point clouds.

With the synergy of pre-training and DAPT, our approach provides a robust way to understand human body configurations in outdoor environments. Although the proposed method is optimization-free and does not use information from any other modalities or time series, we still outperform SOTA methods on multiple datasets. Specifically, our method compared with LPFormer (Ye et al. 2024) on the manually annotated Waymo Open Dataset (Sun et al. 2020) reduces the mean per joint position error (MPJPE) by 10mm ( $\downarrow 16\%$ ). When compared with PRN (Fan et al. 2023b) on the IMU-annotated SLOPER4D dataset (Dai et al. 2023), it reduces the MPJPE by 20.7mm ( $\downarrow 58\%$ ).

In summary, our contribution lies in three main aspects:

- We propose a density-aware pose transformer that steadily mines pose cues from sparse and noisy point clouds.
- We thoroughly investigate the difficulties of estimating human pose by LiDAR data and design a comprehensive pre-training approach.
- Our method greatly improves the stability and accuracy of single-frame LiDAR-only human pose estimation, achieving SOTA performance in multiple scenarios.

## 2 Related works

### 2.1 LiDAR-based 3D human pose estimation

In recent years, many LiDAR point cloud-based 3D HPE methods have been proposed as the practical application value of LiDAR has been explored. (Li et al. 2022a) provides the first LiDAR HPE dataset and proposes the first fully supervised baseline for LiDAR-based motion capture. The keypoint coordinates are obtained by a temporal encoder and optimized with inverse kinematics and SMPL. (Yan et al. 2023) provides a climbing dataset with explicit scene interactions and attempts to perform scene-aware human pose estimation, followed by (Zhang et al. 2024) which uses environmental information of 3D neighbors sampled in the background to enhance the pose learning. (Ren et al. 2024b,a) achieves accurate motion tracking by exploiting temporal and spatial coherence. On the other hand, (Cong et al. 2023; Zheng et al. 2022; Fürst et al. 2021; Hu et al. 2024) perform multi-modal fusion, which utilizes the keypoint cues and geometric constraints provided by the 2D images for weakly-supervised 3D human pose learning. (Ye et al. 2024) proposes a multitasking architecture that augments model learning with segmentation and object detection tasks, and uses a keypoint transformer for multi-person 3D HPE.

Despite the success of these approaches, most of them inevitably use information from other modalities or utilize the priori of body structure provided by SMPL to compensate




Figure 4: Overall structure of our method. It mainly consists of (a) a comprehensive LiDAR human synthesis and augmentation framework to provide internal human priors, and (b) a Density Aware Pose Transformer that uses the multi-density exchange (MDE) module to extract stable joint representations from point cloud features.

for unstable results on sparse or noisy LiDAR point clouds. This can cause an increase in the complexity and latency of the system. Therefore, we think that an optimization-free approach based only on single-frame LiDAR is more practical. This leads us to learn a stable representation of joint for LiDAR-based 3D HPE which increases the reliability of the results and avoids cumbersome post-processing.

## 2.2 Pre-training for human pose estimation

Because pre-training allows models to learn generalized and dataset-independent representations, it has become an important technique to improve the performance of downstream tasks (Devlin et al. 2019; He et al. 2020; Wang et al. 2021; Li et al. 2021, 2022b). On one hand, there have been attempts to use existing data for augmented representation learning. (Xu et al. 2022; An et al. 2024) enhances the accuracy of Vision Transformer-based 2D pose estimation by masked self-supervised pre-training (He et al. 2022). (Qiu et al. 2023) utilizes physical constraints provided by computational photography to conduct weakly supervised pre-training of 3D poses and improve the model’s generalization. (Shan et al. 2022) brings an accuracy increase to 2D-to-3D lifting methods by applying a spatial-temporal mask of the skeleton.

On the other hand, some methods try to use synthetic data to make up for the lack of labeled pose data. (Lin et al. 2024) utilizes the multi-camera projection of SMPL models to generate enough labeled samples for pre-training. (Weng et al. 2023b) proposes a LiDAR scene generation approach based on human mesh under ray casting. (Ren et al. 2023, 2024b,a) craft large-scale synthetic datasets based on (Mahmood et al. 2019) to obtain rich human priors. However, this method is limited by not considering the relative positions of bodies and LiDAR sensors and ignoring the diversity of occlusions. In this paper, we perform a richer and more reasonable sampling of spatial positions so that the model can learn point cloud representations from more aspects. Moreover, we construct more complex occlusion and noise point clouds to simulate the real environment.

## 3 Methodology

Our approach consists of two parts: a) a density-aware pose transformer (DAPT) for stable joint representation learning, b) a comprehensive LiDAR human synthesis & augmentation for model pre-training. We adopt a two-stage training scheme, as shown in Fig.4, we first pre-train the model on synthetic samples, which are generated by ray casting with random augmentations and occlusions. When the samples are fed into the proposed model, they are first encoded as sparse point features. Then, the valid information is progressively extracted to a set of joint anchors through multi-density exchange (MDE) modules and decoded into 1-D heatmaps. Finally, we input real samples into the model for fine-tuning. Note that the network architecture is shared between pre-training and fine-tuning with only slight differences in the loss functions.

In this section, we present the proposed method module by module.

### 3.1 LiDAR human synthesis

The quality and diversity of samples used for pre-training will directly affect the effect of downstream fine-tuning. Inspired by (Weng et al. 2023b), we propose a more comprehensive strategy to sample and augment the scene.

**Scene sampling** Given shape parameters  $\vec{\beta} \in \mathbb{R}^{10}$  and pose parameters  $\vec{\theta} \in \mathbb{R}^{72}$  sampled from a real-world captured SMPL database (Li et al. 2022a), a human instance is simply generated by the SMPL model:

$$\{\mathbf{M}_h, \hat{\mathbf{J}}\} = \text{SMPL}(\vec{\beta}, \vec{\theta}), \quad (1)$$

where  $\mathbf{M}_h = \{\mathbf{V}_h \in \mathbb{R}^{N_V \times 3}, \mathbf{F}_h \in \mathbb{Z}^{N_F \times 3}\}$  is a human mesh with  $N_V$  vertices  $\mathbf{V}_h$  and  $N_F$  triangle faces  $\mathbf{F}_h$ .  $\hat{\mathbf{J}} \in \mathbb{R}^{K \times 3}$  are human joints.

To simulate real ground, a ground mesh  $\mathbf{M}_g = \{\mathbf{V}_g, \mathbf{F}_g\}$  is generated by a random normal vector  $\vec{n}$  with size  $s_g$ , locating at the point with a minimum value of Z axis in  $\mathbf{V}_h$ .

Then, we randomly sample a polar coordinate with distance  $r \in [4m, 20m]$  and azimuth  $\theta \in [-\pi, \pi]$ , and convert

it to Cartesian transition  $\vec{t}$ . Finally, the transformation is applied to  $\mathbf{M}_h$  and  $\mathbf{M}_g$  to get the scene mesh:

$$\mathbf{M} = \{\mathbf{V} = (\mathbf{V}_h \cup \mathbf{V}_g) + \vec{t}, \mathbf{F} = \mathbf{F}_h \cup \mathbf{F}_g\}. \quad (2)$$

**Ray casting** The LiDAR sensor obtains depth information through  $360^\circ$  scans at different elevation angles. Thus, the 3D points it captured can be represented by polar coordinates  $(r, \theta_l, \delta_l)$  within a laser grid  $\mathcal{G} = \{\theta_i\}_{i=1}^{N_\theta} \times \{\delta_j\}_{j=1}^{N_\delta}$ . Since the synthetic mesh only occupies a small angle range, for faster sample generation and to facilitate the subsequent application of laser-level masks, we only intercept the laser within valid azimuth and elevation angles:

$$\begin{aligned} \boldsymbol{\theta} &= [\theta_{\min}(\mathbf{V}), \theta_{\max}(\mathbf{V})] \cap \{\theta_i\}_{i=1}^{N_\theta}, \\ \boldsymbol{\delta} &= [\delta_{\min}(\mathbf{V}), \delta_{\max}(\mathbf{V})] \cap \{\delta_j\}_{j=1}^{N_\delta}, \\ \mathbf{R} &= \{(\theta_l, \delta_l) | \theta_l \in \boldsymbol{\theta}, \delta_l \in \boldsymbol{\delta}\}, \end{aligned} \quad (3)$$

where  $\theta_{\min}(\cdot), \theta_{\max}(\cdot)$  take the minimum and maximum azimuth angles of the vertices, and for the elevation angles, the corresponding values are taken by  $\delta_{\min}(\cdot), \delta_{\max}(\cdot)$ . Then, ray casting is performed with the mesh to get the point cloud  $\mathbf{P} \in \mathbb{R}^{N \times 3}$  and which faces the lasers hit  $\mathbf{H} \in \mathbb{Z}^N$ :

$$(\mathbf{P}, \mathbf{H}) = \text{RayCast}(\mathbf{M}; \mathbf{R}). \quad (4)$$

Note that a label map describing the joints to which triangle faces belong is given by the SMLP model. Hence, the ground truth segmentation  $\hat{\mathbf{S}} \in [0, 1]^{N \times (K+1)}$  of  $\mathbf{P}$  can be calculated by  $\mathbf{H}$ , where  $\hat{\mathbf{S}}_{i,j} = 1$  means the laser point  $p_i \in \mathbf{P}$  belongs to the  $j$  th joint.

**Laser-level masking** To better simulate the occlusion in the real environment, we divide the effective LiDAR laser grid obtained in Eq.3 into patches with an empirical size  $s_p = \lfloor \min\{\max \boldsymbol{\theta} - \min \boldsymbol{\theta}, \max \boldsymbol{\delta} - \min \boldsymbol{\delta}\} / 8 \rfloor$ . Then,  $1 - r_{\text{keep}}$  ratio of the total patches are being masked, which are denoted as  $\mathcal{M}$ . Finally, the point cloud is filtered by:

$$\mathbf{P}_{\text{syn}} = \{p_i | p_i \in \mathbf{P}, r_i \notin \mathcal{M}\}, \quad (5)$$

where  $r_i \in \mathbf{R}$  is the ray producing hit point  $p_i$ .

### 3.2 Density-aware pose transformer

We present a density-aware pose transformer based on UNet-like point transformers (Wu et al. 2024) to obtain stable joint representations of different densities. Specifically, the point cloud is firstly encoded into sparse point cloud features, which are subsequently fed into a decoder to be recovered to point-wise features. During the decoding process, we utilize an MDE module to export valid information from different pooling hierarchies and use 1D heatmaps to represent point locations, since heatmaps are generally more amenable to neural networks than coordinate regression.

**Point cloud feature deduction with MDE** Given an input point cloud  $\mathbf{P} \in \mathbb{R}^{N_0 \times 3}$ , it was encoded into latent point features  $f_M$  by an encoder  $\mathcal{E} : \mathbb{R}^{N_0 \times 3} \rightarrow \mathbb{R}^{N_M \times D_M}$ . Where  $N_0$  is the initial point numbers,  $N_M, D_M$  are the pooled point numbers and pooled point feature dimensions of  $M$  th pooling level.

As shown in Fig.4 right, it can be seen that at deeper pooling levels, the pooled point cloud shows a sparser spatial distribution, which inspires us to model joint-related features at these levels. Therefore, a set of learnable joint anchors  $\mathbf{A}_M \in \mathbb{R}^{K \times D_M}$  is introduced. As the spatial dimensions of the features are expanded through the point transformer blocks, MDE progressively exchanges information between point features  $f_m$  and joint anchors. Then, it updates the joint features by:

$$\mathbf{A}_{m-1} = \text{MDE}_m(\mathbf{A}_m, f_m), m = M, M-1, \dots, 1. \quad (6)$$

The MDE module first aligns the dimensions of joint features  $\mathbf{A}_m$  to the dimensions of  $f_m$  through an MLP, then applies self-attention on it while shortcutting it with padded  $f_m$  for cross-attention.

**Heatmap decoder for joints** Two 1D heatmaps have been proven to be effective in representing 2D keypoint coordinates (Li et al. 2022c). Inspired by this, we extend it to the representation of 3D coordinates of joints. Given the range and number of bins  $N_{\{x,y,z\}}$ , the heatmaps  $h_{\{x,y,z\};i}$  of  $i$  th joint can be predicted by MLPs on the corresponding axis. To get the target coordinate, we decode the 1D heatmaps by taking the peak location of  $h_{\{x,y,z\};i}$ , and map them back to the coordinate through the counted range and bins.

**Pre-training** The goal of pre-training is to allow the model to obtain human priors in the LiDAR-captured point clouds and understand the structure of the human body. Therefore, at this stage, to avoid learning preferences caused by differences in supervision intensity, following (Weng et al. 2023a), we still use coordinate regression on joint features and segmentation on point features for supervision. The joint regression loss  $\mathcal{L}_{\text{reg}}$  is defined by:

$$\mathcal{L}_{\text{reg}} = \sum_{i=1}^K \left\| \mathbf{J}_i - \hat{\mathbf{J}}_i \right\|_2 \cdot v_i / \sum_{i=1}^K v_i, \quad (7)$$

where  $\mathbf{J}$  are the decoded joints coordinate from  $\mathbf{A}_0$  through a shared decoder  $\mathcal{D}_{\text{reg}} : \mathbb{R}^{D_0} \rightarrow \mathbb{R}^3$ , and  $\hat{\mathbf{J}}$  are the corresponding ground truth,  $v_i$  is the joint visibility. The segmentation loss  $\mathcal{L}_{\text{seg}}$  is defined by:

$$\mathcal{L}_{\text{seg}} = - \sum_{i=1}^{N_0} \sum_{j=1}^{K+1} \mathbf{S}_{i,j} \log(\hat{\mathbf{S}}_{i,j}), \quad (8)$$

where  $\mathbf{S}$  are the decoded part segmentation of points from  $f_0$  through a shared decoder  $\mathcal{D}_{\text{seg}} : \mathbb{R}^{D_0} \rightarrow \mathbb{R}^{K+1}$ . Overall, we input synthetic data and minimize:

$$\mathcal{L}_{\text{pre}} = \lambda_{\text{reg}} \mathcal{L}_{\text{reg}} + \lambda_{\text{seg}} \mathcal{L}_{\text{seg}}, \quad (9)$$

where  $\lambda_{\text{reg}}, \lambda_{\text{seg}}$  are loss weights.

**Fine-tuning** The goal of the fine-tuning phase is to adapt the model to the joint annotations of different datasets and their specific distribution. Therefore, in this stage, we only enable the heatmap loss. We input real data and minimize:

$$\mathcal{L}_{\text{ft}} = \mathcal{L}_{\text{hm}} = \sum_{c \in \{x,y,z\}} \sum_{j=1}^{K+1} D_{\text{KL}}(h_{c;j} \| \hat{h}_{c;j}), \quad (10)$$

where  $D_{\text{KL}}$  is the Kullback–Leibler divergence between GT and predicted 1D heatmaps.

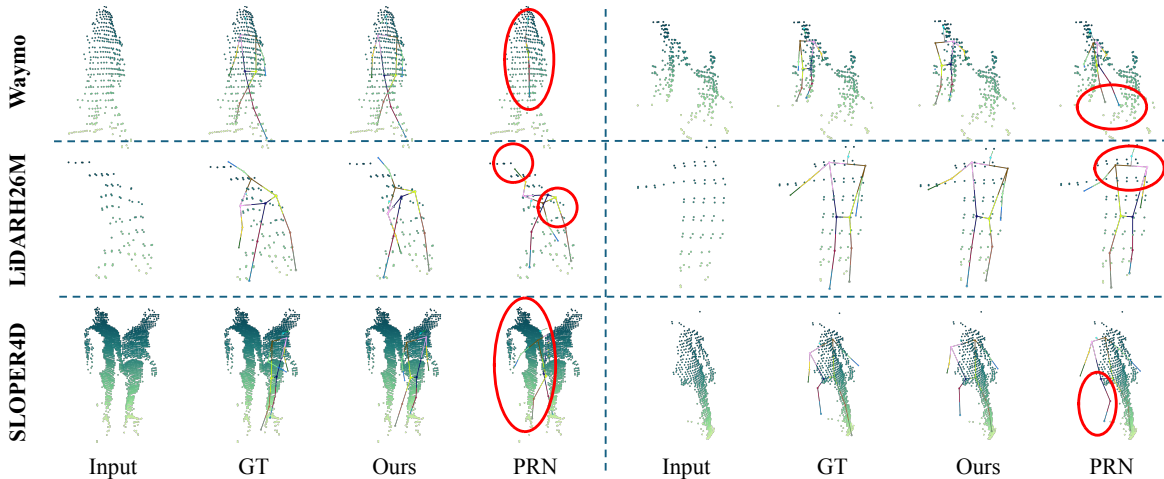


Figure 5: Visualization results on common datasets with challenging samples. Our method demonstrates strong stability on low-quality point clouds including occlusion, noise, and sparsity.

Dataset Metrics	Waymo			HumanM3			LiDARHuman26M				SLOPER4D			
	MPJPE	PCK3	PCK5	MPJPE	PCK3	PCK5	MPJPE	PA-	PCK3	PCK5	MPJPE	PA-	PCK3	PCK5
LidarCap	-	-	-	-	-	-	79.31	66.72	86.00	95.00	101.89	78.93	78.15	89.77
NE	-	-	-	-	-	-	<b>72.23</b>	61.67	87.94	95.79	96.80	76.70	79.22	90.51
LPFormer	61.60	94.52	98.04	83.69	89.54	96.85	95.72	79.03	84.38	94.87	49.31	38.92	97.22	99.44
PRN	68.48	93.60	97.87	70.56	93.18	97.80	80.01	63.67	88.77	96.54	48.76	39.42	97.38	99.45
DAPT(ours)	<b>51.59</b>	<b>97.34</b>	<b>98.98</b>	<b>59.76</b>	<b>95.39</b>	<b>98.30</b>	73.47	<b>58.40</b>	<b>90.54</b>	<b>97.06</b>	<b>28.01</b>	<b>21.52</b>	<b>99.30</b>	<b>99.87</b>

Table 1: Comparison with other methods on IMU-annotated LidarHuman26M, SLOPER4D, and manually annotated Waymo Open Dataset, HumanM3. PA- stands for PA-MPJPE, and the units of MPJPE and PA-MPJPE are *mm*. The best results are highlighted in bold. For the same dataset, the skeleton structures are aligned to ensure fair comparison.

## 4 Experiment

### 4.1 Implementation Details

For the model specification, we follow the typical configuration of PTv3 (Wu et al. 2024), with a voxelization grid size of 0.01. For data synthesis, the SMPL models are sampled from (Li et al. 2022a), and a simulated LiDAR sensor with 64 lines and 2650 angles is applied to perform ray casting, the  $r_{keep}$  is set to 0.6 for laser-level masking. For model training, we perform 50 epochs in both pre-training and fine-tuning with AdamW (Loshchilov and Hutter 2019) optimizer on 2 RTX 4090. We set the batch size to 64 and apply the cosine annealing decay strategy. For pre-training, we set the learning rate to  $3 \times 10^{-4}$  and set  $\lambda_{reg} = 0.5$ ,  $\lambda_{seg} = 1.0$ . For fine-tuning, the learning rate is set to  $5 \times 10^{-4}$ .

### 4.2 Datasets

We use four datasets with different scenarios and annotation methods to evaluate our method:

**LiDARHuman26M** (Li et al. 2022a) A multi-modal dataset uses inertial measurement units (IMUs) captured human poses in SMPL format. It uses a fixed LiDAR sensor to capture a variety of daily actions within a range of 12m to 24m. The scenes are clear and ideal.

**SLOPER4D** (Dai et al. 2023) An IMU annotated dataset captured within a more realistic environment. A mobile Li-

DAR sensor is utilized to track a walking person for capturing point cloud data. Since the official train-test split is unavailable, we utilize the same data split following (Zhang et al. 2024).

**HumanM3** (Fan et al. 2023a) A multi-person pose dataset utilizing automatic annotation and manual review for accurate ground truth poses. It mainly includes multi-person scenes on sports fields, with point clouds and images captured by multiple sets of RGB-LiDAR units. Due to its large size and the small differences between adjacent frames, we only use 20% of the data for training.

**Waymo Open Dataset v2** (Mei et al. 2022) A large-scale multi-task autonomous driving dataset with manual 3D pose annotations, contains 10K human instances.

### 4.3 Metrics

According to common practices, to evaluate model performance, we report **MPJPE**↓ (Mean Per Joint Position Error), **PA-MPJPE**↓ (Procrustes-Aligned Mean Per Joint Position Error), **PCK-3**↑ (Percentage of Correct Keypoints with distance to GT lower than 30% of torsal length), **PCK-5**↑ (Percentage of Correct Keypoints with distance to GT lower than 50% of torsal length). Note that we do not evaluate PA-MPJPE on Waymo and HumanM3, as their visibility labels of joints will interfere with the rigid body alignment process.

Pre-training dataset	Waymo	LiDARH26M
w/o	59.2	76.5
LIPD (Ren et al. 2023)	58.7	75.6
FreeMotion (Ren et al. 2024b)	57.9	74.3
NoiseMotion (Ren et al. 2024a)	59.0	76.3
FreeMotion <sup>†</sup>	54.6	<b>73.2</b>
LiDARH26M <sup>†</sup>	<b>51.6</b>	73.4

Table 2: The effectiveness of our enhanced pre-training method, <sup>†</sup> means applying our data synthesis pipeline rather than the point cloud provided by the dataset.

Anchor	MDE	Heatmaps	MPJPE (gain)
	PRN Baseline		68.5
✓			54.3 (-14.2)
✓	✓		52.3 (-16.2)
✓		✓	53.8 (-14.7)
✓	✓	✓	<b>51.7 (-16.8)</b>

Table 3: Ablation study of progressively enabling our proposed components.

#### 4.4 Comparison methods

We compare our method with several state-of-the-art approaches. Specifically, we first evaluate against advanced Transformer-based models, LPFormer (Ye et al. 2024) and PRN (Fan et al. 2023b), both of which are capable of inferring complete human body poses from single-frame LiDAR data. Additionally, we compare our method with LiDARCap (Li et al. 2022a), which leverages SMPL optimization and temporal information. We also include the Neighbor Enhanced (NE) 3HPE (Zhang et al. 2024), which incorporates information from the surrounding scene. To ensure a consistent comparison, we replace the point cloud backbone in PRN with PTv3. We note that LiDARCap and NE require consecutive frames and ground truth SMPL params, this limits the evaluation process on Waymo and HumanM3 datasets. Therefore, the results are not included.

#### 4.5 Quantitative Results

Tab.1 presents the comparison results, where our method consistently achieves outstanding results across four datasets. On LiDARHuman26M with clear scenes, our method obtains comparable MPJPE to the current best NE but with a reduced PA-MPJPE ( $-3.2mm$ ), suggesting closer alignment with GT. On the more challenging SLOPER4D, our method gets a remarkable  $28.01mm$  MPJPE, representing a significant improvement over the current best PRN ( $-20.7mm$ ). In the motion-focused dataset HumanM3, our method records an MPJPE of  $59.76mm$ , surpassing PRN by  $10.8mm$ . Finally, on the Waymo dataset, which centers on autonomous driving scenarios, our method achieves an MPJPE of  $51.59mm$ , improving by  $10.0mm$  over the SOTA method LPFormer. Overall, our method consistently achieves higher PCK-30 and PCK-50 scores, indicating greater robustness and value of practical usage.

Method	MPJPE (gain)	$r_{keep}$	MPJPE
scratch	59.2	0.5	53.0
+pre-training	52.9 (-6.3)	0.6	<b>51.7</b>
+scene sampling	52.4 (-6.8)	0.7	51.9
+laser masking	<b>51.7 (-7.5)</b>	0.8	52.5
		0.9	52.4

(a) Influence of augmentations

(b) Influence of  $r_{keep}$

Table 4: Ablation study of pre-training related components.

#### 4.6 Stability Evaluation

To evaluate the stability of our method, we introduce various types of disturbance to the point cloud and re-evaluate the model’s performance. Specifically, we apply clusters of noise points at different locations and introduce positional offsets to each point. Results are presented in Fig.7.

**Point jittering** We add point noise to each point coordinate of the input and use different thresholds for clipping. The results are shown in Fig.7a. It can be observed that compared with baseline methods, our method has less performance degradation, and the proposed MDE module also plays a positive role in dealing with disturbances.

**Noise clusters** We simulate extreme input conditions by adding clusters containing varying amounts of noise points to the input point cloud. As shown in Fig.7b, as the number of noise points per cluster increases, our method yields more stable results. Notably, our pre-training substantially enhances the model’s ability to handle noisy point clouds.

**Errors on end joints** We evaluate the errors on the most challenging end joints, and the results are shown in Fig.8. Our method demonstrates smaller average error and variance on these joints. Especially, the proposed pre-training method effectively stabilizes the prediction of ankles and wrists.

#### 4.7 Ablation study

**Model components** In Tab.3, we conduct ablation studies on the Waymo dataset to evaluate the contributions of the key components by incrementally enabling each proposed module. Specifically, row 1 represents the PRN baseline. Row 2 shows the results after introducing joint anchors and a 4-layer Transformer on top of the off-the-shelf PTv3. Row 3 demonstrates the effect of enabling the proposed MDE module on multi-scale point cloud features. In row 4, the coordinate-based decoding is replaced with heatmap-based decoding. The final row presents the results with all modules enabled. The combined effect of joint anchors and the MDE module leads to a substantial  $16.2mm$  improvement in MPJPE over the baseline. Replacing the decoder with a heatmap-based one results in an additional  $0.6mm$  improvement.

**Pre-Training** In Tab.4, we conduct ablation studies to assess the impact of our proposed pre-training strategies. Tab.4a evaluates the effectiveness of the enhancements: the first row represents training without any pre-training, the second row uses the same pre-training strategy as (Weng




Figure 6: Visualization of synthesis samples. For the same SMPL human body, our method can produce diverse point clouds, and the laser level mask introduces extreme cases, which equips the model with more robust human priors and the ability to recover human pose from low-quality point clouds.

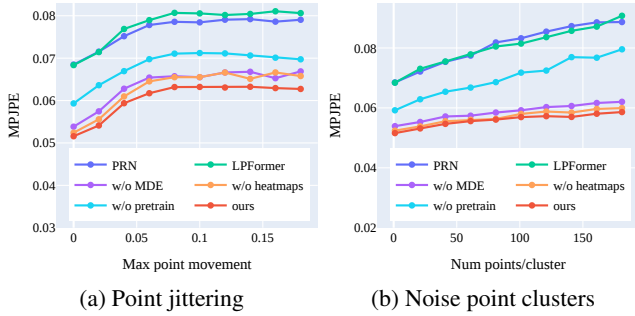


Figure 7: Stability evaluation of our method. We observe that the performance degradation of the proposed method is within a more acceptable range.




Figure 8: Errors comparison of the most challenging human end joints. The worst 600 instances of each method are included in the statistics.

et al. 2023a), the third row implements scene resampling, and the fourth row further incorporates the proposed laser level masking. Comparing rows 1 and 2, the inclusion of pre-training brings an MPJPE reduction of  $6.3mm$ . Further comparison of rows 1 and 4 shows that incorporating both scene resampling and laser-level masking leads to an even greater improvement of  $7.5mm$ . Moreover, as shown in Tab.4b, to strike a balance between providing sufficient information for human pose reconstruction and enabling the model to recover human structures from occluded point clouds, we vary the proportion of the laser level mask  $r_{keep}$  from 0.5 to 0.9 to determine the optimal parameter. In addition, as shown in Tab.2, we also compare our synthesis pipeline with LIPD (Ren et al. 2023), FreeMotion (Ren et al. 2024b), and NoiseMotion (Ren et al. 2024a), which also use synthetic data. Pre-training with data synthesized by the proposed method can bring more performance improvements.

## 4.8 Visualization

**Prediction results** Fig.5 presents the visualization results of our method. We select some challenging samples from the dataset, and our method consistently produces more stable results. Row 1 left shows a walking person from the Waymo dataset, but the left-right symmetry of the body is unclear. A segmentation-regression-based method produces wrong results, attempting to disregard the human body’s rigid structure to mitigate inference errors caused by left-right ambiguity. In contrast, our method penalizes such ambiguities in the heatmap, effectively avoiding this issue. Row 2 illustrates the detection results on sparse point clouds from the LiDARHuman26M dataset. In this scenario, our method effectively recognizes human body orientation from limited clues. Lastly, as depicted in row 1 right and row 3 left, when multiple human instances are unexpectedly introduced by the human detector, our approach does not confuse the two individuals. In summary, our method demonstrates strong capability in robust inference.

**Synthetic samples** We present the synthesis samples of our approach in Fig.6. It is evident that scene resampling produces diverse synthetic samples, effectively generating both sparse and dense point clouds of human instances. Additionally, the laser-level masks introduce more complex occlusions, which in turn equip the model with more robust human priors and point cloud representations.

## 5 Conclusion

We propose a novel method for robust LiDAR-based 3D human pose estimation, with two main contributions. First, we introduce a density-aware pose transformer, which employs joint anchors and special exchange modules to extract valid features from point clouds of varying densities, enabling explicit learning of stable keypoint representations. Second, we present a comprehensive LiDAR-based human synthesis and augmentation approach for model pre-training. By integrating more randomized position sampling, ground modeling, and laser-level masking, we generate highly realistic and challenging samples. Qualitative and quantitative evaluations demonstrate that with the synergy of the proposed model and pre-training strategies, our method achieves state-of-the-art performance across multiple datasets.

Overall, our approach provides a comprehensive solution for LiDAR-based 3D HPE. Future work can try to address left-right reversal and multi-frame jitter in timing or perform pose-based human behavior understanding tasks.

## Acknowledgments

This work was supported by the National Science Fund of China under Grant Nos.U24A20330, 62172222, 62361166670, and the National Key Research and Development Program of China (International Collaboration Special Project, No. SQ2023YFE0102775).

## References

- An, X.; Zhao, L.; Gong, C.; Wang, N.; Wang, D.; and Yang, J. 2024. SHaRPose: Sparse High-Resolution Representation for Human Pose Estimation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(2): 691–699.
- Cong, P.; Xu, Y.; Ren, Y.; Zhang, J.; Xu, L.; Wang, J.; Yu, J.; and Ma, Y. 2023. Weakly Supervised 3D Multi-Person Pose Estimation for Large-Scale Scenes Based on Monocular Camera and Single LiDAR. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(1): 461–469.
- Cong, P.; Zhu, X.; Qiao, F.; Ren, Y.; Peng, X.; Hou, Y.; Xu, L.; Yang, R.; Manocha, D.; and Ma, Y. 2022. STCrowd: A Multimodal Dataset for Pedestrian Perception in Crowded Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19608–19617.
- Dai, Y.; Lin, Y.; Lin, X.; Wen, C.; Xu, L.; Yi, H.; Shen, S.; Ma, Y.; and Wang, C. 2023. SLOPER4D: A Scene-Aware Dataset for Global 4D Human Pose Estimation in Urban Environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 682–692.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805.
- Fan, B.; Wang, S.; Guo, W.; Zheng, W.; Feng, J.; and Zhou, J. 2023a. Human-M3: A Multi-view Multi-modal Dataset for 3D Human Pose Estimation in Outdoor Scenes. arXiv:2308.00628.
- Fan, B.; Zheng, W.; Feng, J.; and Zhou, J. 2023b. LiDAR-HMR: 3D Human Mesh Recovery from LiDAR. arXiv:2311.11971.
- Fan, L.; Zhu, Y.; Zhu, J.; Liu, Z.; Zeng, O.; Gupta, A.; Creus-Costa, J.; Savarese, S.; and Fei-Fei, L. 2018. SUR-REAL: Open-Source Reinforcement Learning Framework and Robot Manipulation Benchmark. In *Proceedings of The 2nd Conference on Robot Learning*, 767–782.
- Fürst, M.; Gupta, S. T. P.; Schuster, R.; Wasenmüller, O.; and Stricker, D. 2021. HPERL: 3D Human Pose Estimation from RGB and LiDAR. In *Proceedings of International Conference on Pattern Recognition*, 7321–7327.
- Gong, K.; Zhang, J.; and Feng, J. 2021. PoseAug: A Differentiable Pose Augmentation Framework for 3D Human Pose Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8575–8584.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked Autoencoders Are Scalable Vision Learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15979–15988.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9729–9738.
- Hong, S.; and Kim, Y. 2018. Dynamic Pose Estimation Using Multiple RGB-D Cameras. *Sensors*, 18(11): 3865.
- Hu, X.; Zhong, B.; Liang, Q.; Zhang, S.; Li, N.; and Li, X. 2024. Towards Modalities Correlation for RGB-T Tracking. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2014. 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.
- Iskakov, K.; Burkov, E.; Lempitsky, V.; and Malkov, Y. 2019. Learnable Triangulation of Human Pose. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7718–7727.
- Kang, Y.; Liu, Y.; Yao, A.; Wang, S.; and Wu, E. 2023. 3D Human Pose Lifting with Grid Convolution. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(1): 1105–1113.
- Li, J.; Zhang, J.; Wang, Z.; Shen, S.; Wen, C.; Ma, Y.; Xu, L.; Yu, J.; and Wang, C. 2022a. LiDARCap: Long-range Markerless 3D Human Motion Capture with LiDAR Point Clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20470–20480.
- Li, X.; Wang, W.; Yang, L.; and Yang, J. 2022b. Uniform Masking: Enabling MAE Pre-training for Pyramid-based Vision Transformers with Locality. arXiv:2205.10063.
- Li, Y.; Yang, S.; Liu, P.; Zhang, S.; Wang, Y.; Wang, Z.; Yang, W.; and Xia, S.-T. 2022c. SimCC: A Simple Coordinate Classification Perspective for Human Pose Estimation. In *ECCV*, 89–106.
- Li, Z.; Ye, J.; Song, M.; Huang, Y.; and Pan, Z. 2021. Online Knowledge Distillation for Efficient Pose Estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11740–11750.
- Lian, J.; Wang, D.; Zhu, S.; Wu, Y.; and Li, C. 2022. Transformer-Based Attention Network for Vehicle Re-Identification. *Electronics*, 11(7).
- Lin, K.; Lin, C.-C.; Liang, L.; Liu, Z.; and Wang, L. 2024. MPT: Mesh Pre-Training With Transformers for Human Pose and Mesh Reconstruction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 3415–3425.
- Loper, M.; Mahmood, N.; Romero, J.; Pons-Moll, G.; and Black, M. J. 2015. SMPL: A Skinned Multi-Person Linear Model. *ACM Transactions on Graphics*, 34(6): 248:1–248:16.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. arXiv:1711.05101.
- Mahmood, N.; Ghorbani, N.; Troje, N. F.; Pons-Moll, G.; and Black, M. J. 2019. AMASS: Archive of Motion Capture As Surface Shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5442–5451.



- Mei, J.; Zhu, A. Z.; Yan, X.; Yan, H.; Qiao, S.; Zhu, Y.; Chen, L.-C.; Kretschmar, H.; and Anguelov, D. 2022. Waymo Open Dataset: Panoramic Video Panoptic Segmentation. *arXiv:2206.07704*.
- Qiu, Z.; Qiu, K.; Fu, J.; and Fu, D. 2023. Weakly-Supervised Pre-Training for 3D Human Pose Estimation via Perspective Knowledge. *Pattern Recognition*, 139: 109497.
- Ren, Y.; Han, X.; Yao, Y.; Long, X.; Sun, Y.; and Ma, Y. 2024a. LiveHPS++: Robust and Coherent Motion Capture in Dynamic Free Environment. In *ECCV*, 127–144.
- Ren, Y.; Han, X.; Zhao, C.; Wang, J.; Xu, L.; Yu, J.; and Ma, Y. 2024b. LiveHPS: LiDAR-based Scene-level Human Pose and Shape Estimation in Free Environment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1281–1291.
- Ren, Y.; Zhao, C.; He, Y.; Cong, P.; Liang, H.; Yu, J.; Xu, L.; and Ma, Y. 2023. LiDAR-aid Inertial Poser: Large-scale Human Motion Capture by Sparse Inertial and LiDAR Sensors. *IEEE Transactions on Visualization and Computer Graphics*, 29(5): 2337–2347.
- Shan, W.; Liu, Z.; Zhang, X.; Wang, S.; Ma, S.; and Gao, W. 2022. P-STMO: Pre-trained Spatial Temporal Many-to-One Model for 3D Human Pose Estimation. In *ECCV*, 461–478.
- Simon, T.; Joo, H.; Matthews, I.; and Sheikh, Y. 2017. Hand Keypoint Detection in Single Images Using Multiview Bootstrapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4645–4653.
- Su, Z.; Xu, L.; Zheng, Z.; Yu, T.; Liu, Y.; and Fang, L. 2020. RobustFusion: Human Volumetric Capture with Data-Driven Visual Cues Using a RGBD Camera. In *ECCV*, 246–264.
- Sun, P.; Kretschmar, H.; Dotiwalla, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; Vasudevan, V.; Han, W.; Ngiam, J.; Zhao, H.; Timofeev, A.; Ettinger, S.; Krivokon, M.; Gao, A.; Joshi, A.; Zhang, Y.; Shlens, J.; Chen, Z.; and Anguelov, D. 2020. Scalability in Perception for Autonomous Driving: Waymo Open Dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2446–2454.
- Tu, H.; Wang, C.; and Zeng, W. 2020. VoxelPose: Towards Multi-camera 3D Human Pose Estimation in Wild Environment. In *ECCV*, 197–212.
- Wang, H.; Liu, Q.; Yue, X.; Lasenby, J.; and Kusner, M. J. 2021. Unsupervised Point Cloud Pre-Training via Occlusion Completion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9782–9792.
- Weng, Z.; Gorban, A. S.; Ji, J.; Najibi, M.; Zhou, Y.; and Anguelov, D. 2023a. 3D Human Keypoints Estimation From Point Clouds in the Wild Without Human Labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1158–1167.
- Weng, Z.; Gorban, A. S.; Ji, J.; Najibi, M.; Zhou, Y.; and Anguelov, D. 2023b. 3D Human Keypoints Estimation From Point Clouds in the Wild Without Human Labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1158–1167.
- Wu, X.; Jiang, L.; Wang, P.-S.; Liu, Z.; Liu, X.; Qiao, Y.; Ouyang, W.; He, T.; and Zhao, H. 2024. Point Transformer V3: Simpler Faster Stronger. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4840–4851.
- Xu, Y.; Zhang, J.; ZHANG, Q.; and Tao, D. 2022. ViTPose: Simple Vision Transformer Baselines for Human Pose Estimation. In *Advances in Neural Information Processing Systems*, volume 35, 38571–38584.
- Yan, M.; Wang, X.; Dai, Y.; Shen, S.; Wen, C.; Xu, L.; Ma, Y.; and Wang, C. 2023. CIMI4D: A Large Multimodal Climbing Motion Dataset under Human-scene Interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12977–12988.
- Yan, M.; Zhang, Y.; Cai, S.; Fan, S.; Lin, X.; Dai, Y.; Shen, S.; Wen, C.; Xu, L.; Ma, Y.; and Wang, C. 2024. RELI1D: A Comprehensive Multimodal Human Motion Dataset and Method. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2250–2262.
- Ye, D.; Xie, Y.; Chen, W.; Zhou, Z.; Ge, L.; and Foroosh, H. 2024. LPFormer: LiDAR Pose Estimation Transformer with Multi-Task Network. In *IEEE International Conference on Robotics and Automation*, 16432–16438.
- Ye, H.; Zhu, W.; Wang, C.; Wu, R.; and Wang, Y. 2022. Faster VoxelPose: Real-time 3D Human Pose Estimation by Orthographic Projection. In *ECCV*, 142–159.
- Ying, J.; and Zhao, X. 2021. Rgb-D Fusion For Point-Cloud-Based 3d Human Pose Estimation. In *Proceedings of IEEE International Conference on Image Processing*, 3108–3112.
- Zhang, J.; Mao, Q.; Hu, G.; Shen, S.; and Wang, C. 2024. Neighborhood-Enhanced 3D Human Pose Estimation with Monocular LiDAR in Long-Range Outdoor Scenes. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(7): 7169–7177.
- Zhao, Z.; Luo, L.; Pan, S.; Zhang, C.; and Gong, C. 2024. Graph Stochastic Neural Process for Inductive Few-shot Knowledge Graph Completion. *arXiv preprint arXiv:2408.01784*.
- Zheng, J.; Shi, X.; Gorban, A.; Mao, J.; Song, Y.; Qi, C. R.; Liu, T.; Chari, V.; Cornman, A.; Zhou, Y.; Li, C.; and Anguelov, D. 2022. Multi-Modal 3D Human Pose Estimation with 2D Weak Supervision in Autonomous Driving. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 4477–4486.