

SkelSplat: Robust Multi-view 3D Human Pose Estimation with Differentiable Gaussian Rendering

Laura Bragagnolo*
University of Padova
bragagnolo@dei.unipd.it

Leonardo Barcellona
University of Amsterdam
l.barcellona@uva.nl

Stefano Ghidoni
University of Padova
ghidoni@dei.unipd.it

Abstract

Accurate 3D human pose estimation is fundamental for applications such as augmented reality and human-robot interaction. State-of-the-art multi-view methods learn to fuse predictions across views by training on large annotated datasets, leading to poor generalization when the test scenario differs. To overcome these limitations, we propose *SkelSplat*, a novel framework for multi-view 3D human pose estimation based on differentiable Gaussian rendering. Human pose is modeled as a skeleton of 3D Gaussians, one per joint, optimized via differentiable rendering to enable seamless fusion of arbitrary camera views without 3D ground-truth supervision. Since Gaussian Splatting was originally designed for dense scene reconstruction, we propose a novel one-hot encoding scheme that enables independent optimization of human joints. *SkelSplat* outperforms approaches that do not rely on 3D ground truth in *Human3.6M* and *CMU*, while reducing the cross-dataset error up to 47.8% compared to learning-based methods. Experiments on *Human3.6M-Occ* and *Occlusion-Person* demonstrate robustness to occlusions, without scenario-specific fine-tuning. Our project page is available here: <https://skelsplat.github.io>.

1. Introduction

Accurately estimating the three-dimensional human pose is a cornerstone for intelligent systems that perceive, interpret, and learn from human behavior. From autonomous driving [57] and human-robot collaboration [43] to animation [58] and robot learning [59], applications across computer vision and AI increasingly demand robust 3D pose estimation. Among existing solutions, multi-view systems of calibrated cameras remain highly effective, as they triangulate 2D predictions to recover 3D poses [16, 22, 46, 55]. However, current state-of-the-art methods typically learn to fuse multi-view predictions




Figure 1. *SkelSplat* 3D Gaussian joints are optimized (red) by aligning renderings with 2D detection heatmaps (green arrows).

directly from annotated datasets, making them tightly coupled to specific camera configurations [38], pose distributions [30], and occlusion patterns [4]. Consequently, performance often degrades sharply when these conditions change, requiring retraining or fine-tuning for each deployment scenario, reducing real-world applicability.

In this work, we argue that multi-view fusion for 3D human pose estimation should be inherently flexible: it should generalize across arbitrary camera configurations and handle occlusions and appearance variations without retraining. To this end, we introduce *SkelSplat*, a novel framework based on differentiable Gaussian rendering [25].

SkelSplat leverages the Gaussian Splatting representation [25] to model each human joint as a 3D Gaussian. Given an initial 3D pose guess and 2D predictions from multiple cameras, *SkelSplat* constructs a Gaussian representation by building an anisotropic Gaussian for each body joint and optimizes its position and shape by minimizing a rendering loss, as shown in Fig. 1. We tailor the rendering function to human pose estimation by introducing a one-hot encoding scheme for joints. This allows our framework to robustly integrate multi-view cues without requiring scenario-specific training.

We evaluate *SkelSplat* on standard benchmarks such as *Human3.6M* [21] and *CMU Panoptic Studio* [24], show-

*Corresponding author.

casing its adaptability to diverse camera configurations. To assess robustness to occlusions, we further test on Human3.6M-Occ [4], where *SkelSplat* achieves state-of-the-art performance, and on Occlusion-Person [55], where it demonstrates strong accuracy without any training. Notably, the proposed method removes the need for data collection and fine-tuning in new scenarios, thus facilitating scalable deployment.

In summary, our contributions are:

1. We propose *SkelSplat*, a novel framework for multi-view 3D human pose estimation, leveraging differentiable Gaussian rendering for view fusion;
2. We adapt Gaussian Splatting, primarily used for dense scene modeling, to skeleton-based 3D pose estimation;
3. We modify the original Gaussian Splatting rendering function to encode human joints using a one-hot representation, enabling pose-specific optimization;
4. We demonstrate that *SkelSplat* achieves accurate 3D pose estimation under challenging occlusions and varying camera setups, without requiring retraining or fine-tuning.

2. Related Work

Gaussian Splatting Gaussian Splatting [25] had a disruptive impact on graphics [14, 17, 20, 54], achieving state-of-the-art results in novel view synthesis. While subsequent works enhanced representations [20, 54], reconstruction losses [14], or optimization techniques [17], our work focuses on applying the original formulation to multi-view 3D human pose estimation. Gaussian Splatting has already been explored in diverse domains, including SLAM [29], real-to-sim transfer [23], open-vocabulary segmentation [35], robot learning [1], and human reconstruction [26], typically incorporating Gaussian Splatting without any fundamental changes to rendering or optimization pipelines. In contrast, we introduce a novel rendering function and supervision strategy, adapting Gaussian Splatting to the specific challenges of multi-view pose estimation.

Gaussian Splatting for 3D Human Reconstruction One of the closest applications of Gaussian Splatting for human perception is human reconstruction. Early works reconstruct accurate 3D human avatars from sparse images [18, 26], often incorporating parametric human models [42, 51, 52], such as SMPL [3], or skeletal priors [19, 33, 48] for articulated shapes and novel view generation. Other approaches integrate diffusion for occlusion handling [40] or extract human features from Gaussian reconstructions [11, 34], but primarily focus on high-quality offline reconstruction and subject-specific optimization. While these methods prioritize visual fidelity, our work instead focuses on accurate 3D joint localization.

Multi-view 3D Human Pose Estimation Multi-view human pose estimation aims to recover the 3D positions of each body joint by leveraging multiple synchronized cameras. Several works have explored learning-based fusion strategies to mitigate the reliance on accurate 2D detections, which often fail under occlusions [16, 27, 55]. Iskakov et al. [22] aggregate 2D features into a shared 3D volume. He et al. [16] use epipolar geometry to attend to geometrically consistent pixels. TransFusion [27] combines self-attention with positional encoding for occlusion robustness. AdaFuse [55] adaptively weights views to reduce the impact of low-quality detections. While these methods handle occlusions, they remain limited in out-of-domain scenarios because their fusion strategies are learned from few collected datasets, such as Human3.6M [21] and CMU Panoptic [24].

Recent works have therefore focused on robustness and generalization without relying on direct 3D supervision, using temporal consistency [10, 30], geometric consistency [4, 56], multi-modal fusion [5] or optimization pipelines considering body priors learned on datasets [8]. While these approaches focus on learning fusion strategies, *SkelSplat* employs differentiable Gaussian rendering to directly optimize 3D poses without relying on dataset-specific assumptions, demonstrating that a Gaussian representation, derived from Kerbl et al. [25], generalizes better, especially in occluded scenes.

One-hot Encoding for Human Joints In many vision tasks, such as classification [12, 15] or semantic segmentation [6, 9], class identity is represented via one-hot encoding, with each class assigned to a dedicated output index or channel. 3D and 2D human pose estimation adopt the same principle, predicting joint coordinates [28, 53] or joint-specific heatmaps [31, 41, 50]. This strategy has not yet been explored in Gaussian Splatting, where the rendered channels are typically RGB views of the reconstructed scene [25], not intended for accurate 3D points localization. To address this, we adapt Gaussian Splatting to operate on a one-hot encoding representation by redefining the rendering function and extending the output channels beyond RGB, enabling precise 3D joint localization.

3. Method

This section presents *SkelSplat*, a novel method for multi-view 3D human pose estimation that leverages differentiable Gaussian rendering via Gaussian Splatting. An overview of our framework is shown in Fig. 2. In *SkelSplat*, a skeleton $SK = \{sk_0, \dots, sk_N\}$ is represented as a set of anisotropic 3D Gaussians $GS = \{gs_0, \dots, gs_N\}$, in which each joint sk_j is associated with a Gaussian gs_j . The skeleton is optimized using the set of 2D keypoint detections $SK_i^{2D} = \{sk_{i0}^{2D}, \dots, sk_{iN}^{2D}\}$ computed by a 2D human pose estimator on each camera i of a set of M synchronized

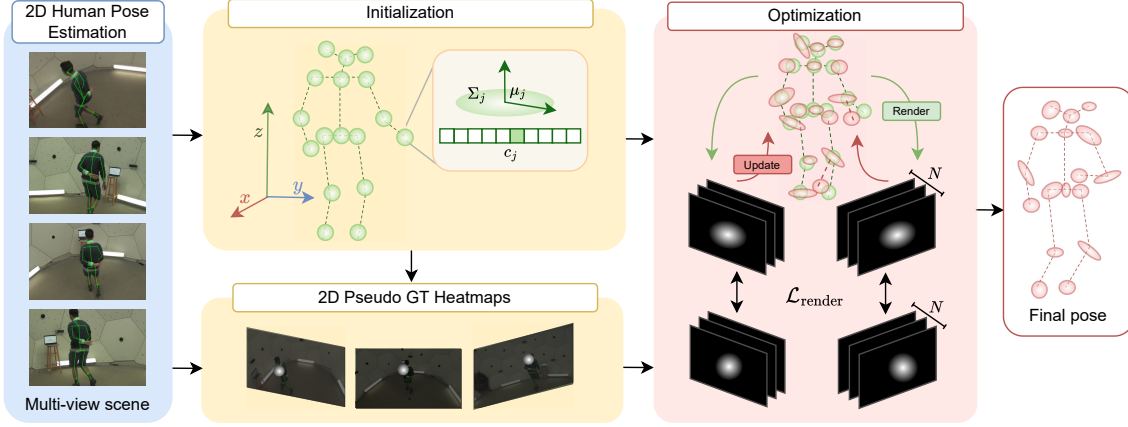


Figure 2. Overview of the *SkelSplat* framework. Given multi-view images and 2D pose detections, we initialize a skeleton of 3D Gaussians, one per human joint. Pseudo ground truth heatmaps are generated from the 2D detections and used to supervise the optimization, which refines the Gaussians by minimizing a differentiable loss between heatmaps and Gaussian renderings.

and calibrated views, with $i \in \{1, \dots, M\}$. *SkelSplat* uses the detections SK_i^{2D} to construct pseudo ground truth heatmaps $GS_i^{2D} = \{gs_{i0}^{2D}, \dots, gs_{iN}^{2D}\}$ needed to optimize GS via differentiable rendering. This formulation allows for direct supervision from 2D keypoints alone, without requiring RGB supervision or 3D ground truth, making the method adaptable to varying camera configurations and robust in diverse scenarios.

In the following, we first provide a brief overview of 3D Gaussian Splatting. We then present the details about the skeletal representation GS, the pseudo ground truth GS_i^{2D} , and the optimization process.

Background Before delving into the details of *SkelSplat*, we revise the classical definition of Gaussian Splatting [25]. The objective of Gaussian Splatting is representing a scene as a set G of anisotropic Gaussians to render novel views. Each Gaussian g_i is defined by a point center $p_i = (x, y, z) \in \mathbb{R}^3$ and a covariance matrix $\Sigma_i \in \mathbb{R}^{3 \times 3}$, which together describe g_i 's spatial extent and orientation. The covariance is parameterized using a scale vector $S_i \in \mathbb{R}^3$ and a rotation matrix $R_i \in SO(3)$, such that $\Sigma_i = R_i S_i S_i^T R_i^T$. Each Gaussian is also associated with an opacity $\alpha_i \in \mathbb{R}$ and a color representation c_i , modeled using spherical harmonics [37] to support view-dependent radiance. During rendering, each 3D Gaussian is projected onto the 2D image plane using an affine approximation of the projective transformation. The projected 2D covariance Σ_i^{2D} is computed as $\Sigma_i^{2D} = JW\Sigma_iW^TJ^T$, where W is the view transformation matrix and J is the Jacobian of the affine approximation. In the original implementation, the set G is initialized from Structure-from-Motion (SfM), while the parameters of the Gaussians are optimized to minimize the photometric loss between the rendered images and the ground truth

RGB images. During optimization, Gaussians densification and pruning are performed to better fit the entire scene.

3D Human Pose as Skeleton of Gaussians Given an initial estimation of the 3D pose of a person $\hat{SK} = \{\hat{sk}_0, \dots, \hat{sk}_N\}$, we represent the human pose as a skeleton of Gaussians GS, in which each $gs_j = (\mu_j, \Sigma_j)$ represents \hat{sk}_j , where $\mu_j \in \mathbb{R}^3$ is the 3D position of \hat{sk}_j , and $\Sigma_j \in \mathbb{R}^{3 \times 3}$ is the covariance, empirically initialized to $\Sigma_j = 3 \cdot \mathbf{I}_3$, with \mathbf{I}_3 denoting the 3×3 identity matrix. The approach is agnostic to the initialization method; therefore GS can be constructed by different approaches, such as triangulation or monocular pose estimation [4].

Joint encoding and rendering During the optimization process, we need to ensure that each gs_j is supervised according to the correct joint prediction sk_{ij}^{2D} . Thus, we modify the RGB-based appearance encoding c_j with a joint identity encoding that activates only the j -th channel:

$$c_j[k] = \begin{cases} 1 & \text{if } k = j \\ 0 & \text{otherwise} \end{cases} \quad \text{for } k = 1, \dots, N. \quad (1)$$

This identity vector is encoded using degree-zero spherical harmonics coefficients, repurposing the appearance field to represent joint identity rather than color. However, this change would not be effective without modifying the rendering function that splats the Gaussians onto the image plane. We modify this function to produce a N -channel tensor, where each channel corresponds to the splat of a single joint Gaussian gs_j on a camera view. This allows independent supervision of gs_j , even when multiple 2D joint detections sk_{ij}^{2D} are overlapping in the image plane.

Supervision from 2D Detections To supervise the optimization, we exploit geometric cues only, to enable flexible supervision across diverse datasets. To do this, we generate pseudo ground truth heatmaps from 2D keypoint detections obtained by a pre-trained model.

For each camera view $i \in \{1, \dots, M\}$, we generate a set of N pseudo ground truth images $\{I_{ij}\}_{j=1}^N$, one for each joint. Each image $I_{ij} \in \mathbb{R}^{H \times W}$ is constructed by rendering a 2D Gaussian gs_{ij}^{2D} centered at the detected 2D keypoint $sk_{ij}^{2D} \in \mathbb{R}^2$, with a covariance matrix $\Sigma_{ij}^{2D} \in \mathbb{R}^{2 \times 2}$. Covariance in 2D joint heatmaps must be geometrically consistent with the 3D Gaussian representation: 3D joints closer to the camera should exhibit larger projected covariances due to perspective scaling, while distant joints should appear smaller. To ensure multi-view geometric consistency, we compute Σ_{ij}^{2D} , 2D covariance of joint j in view i , by reprojecting the 3D covariance Σ_j of the initial 3D pose estimate:

$$\Sigma_{ij}^{2D} = J_i W_i \Sigma_j W_i^T J_i^T, \quad (2)$$

where $W_i \in \mathbb{R}^{4 \times 4}$ is the camera extrinsic transformation (world-to-camera), and $J_i \in \mathbb{R}^{2 \times 3}$ is the Jacobian matrix of the perspective projection at the joint position in camera coordinates.

This construction yields 2D joint heatmaps whose Gaussian shape accurately reflects the 3D joint Gaussians under perspective projection, enabling precise and geometrically consistent multi-view optimization.

Optimization The human pose, represented using GS, is splatted across all camera views and refined to minimize the masked \mathcal{L}_2 loss between the pseudo ground truth images and the rendered images:

$$\mathcal{L}_{\text{render}} = \sum_{j=1}^N \sum_{i=1}^M \left\| \mathcal{M}_{ij} \odot \left(I_{ij}^{\text{render}} - I_{ij}^{\text{pseudo}} \right) \right\|_2^2. \quad (3)$$

Here, N is the number of joints, M the number of camera views, and I_{ij}^{render} and I_{ij}^{pseudo} are the rendered and pseudo ground truth heatmaps for joint j in camera i , with the latter generated according to Eq. (2). \mathcal{M}_{ij} is a binary mask that selects only pixels with non-zero values in either I_{ij}^{render} or I_{ij}^{pseudo} . The choice of a masked \mathcal{L}_2 loss over the standard one is due to the high presence of background regions that can reduce convergence.

We introduce a 3D structural symmetry loss that regularizes the lengths of symmetric limbs, such as arms or legs, to encourage the recovered 3D structure to be anatomically coherent even in the presence of noisy observations, such as in the case of occlusions. For each pair of symmetric limbs, we penalize inconsistencies in their lengths:

$$\mathcal{L}_{\text{sym}} = \sum_{(l,r) \in \mathcal{S}} \left(\|\mathbf{p}_l^1 - \mathbf{p}_l^2\|_2 - \|\mathbf{p}_r^1 - \mathbf{p}_r^2\|_2 \right)^2, \quad (4)$$




Figure 3. One-hot encoding, each joint rendered to its channel.

where \mathcal{S} denotes the set of symmetric limb pairs; p_l^1 and p_l^2 are the 3D endpoints of the left limb, and p_r^1 and p_r^2 are the endpoints of the corresponding right limb. This term acts as a regularizer that complements the view-based loss by promoting globally consistent pose estimations.

The total loss used during optimization is a combination of the masked multi-view \mathcal{L}_2 loss and the 3D structural loss:

$$\mathcal{L}_{\text{opt}} = \mathcal{L}_{\text{render}} + \lambda_{\text{sym}} \mathcal{L}_{\text{sym}}, \quad (5)$$

where λ_{sym} is a weighting factor that balances the influence of the symmetry term, empirically set to 1e-5.

Unlike standard Gaussian Splatting, which updates the 3D Gaussians independently for each view after computing the loss, we accumulate gradients across all views, improving stability and cross-view coherence. In this way, cues from views are more effectively merged, allowing the optimization to better handle partial or occluded observations from individual views that might bias the optimization. Gaussian pruning and densification are removed to avoid changing the cardinality of GS; the optimization, based on Adam, is run for up to 125 iterations, while early stopping is triggered when the difference between the minimum loss in two consecutive windows of size M is less than 1e-6.

4. Experiments

We evaluate our approach on four multi-view datasets to assess its robustness under different levels of occlusion and to test generalization to novel scenarios.

4.1. Datasets

Human3.6M Human3.6M [21] is a widely used benchmark for 3D human pose estimation. It contains 3.6 million frames recorded from four synchronized cameras. The dataset includes 11 subjects performing everyday activities, such as walking and eating. Following standard protocol, subjects S9 and S11 are used for evaluation.

CMU Panoptic Studio The CMU Panoptic Studio [24] provides rich multi-view recordings of people engaged in various activities and social interactions. Its capture setup

includes 480 VGA cameras, 31 HD cameras, and 10 Kinect sensors. In this study, we use only sequences featuring a single person, following Xiang et al. [49].

Human3.6M-Occ To evaluate our approach on challenging occlusion scenarios, we consider Human3.6M-Occ [4], a modified version of Human3.6M that introduces occlusions by overlaying objects and animals from the Pascal VOC 2012 dataset [13] on the original images. We consider three configurations: *Human3.6M-Occ-2*, with 2 views occluded out of 4, *Human3.6M-Occ-3*, with 3 views occluded, and *Human3.6M-Occ-3-Hard*, which features larger occluders over the subject’s body.

Occlusion-Person Occlusion-Person [55] is a synthetic dataset comprising 13 human models animated with poses from the CMU Motion Capture database and rendered in 9 indoor scenes. Heavy occlusions are introduced using common objects like sofas and desks. Each scene is captured from 8 evenly spaced camera views arranged in a circle.

Evaluation We assess 3D pose estimation accuracy using the Mean Per Joint Position Error (MPJPE), which measures the average Euclidean distance in millimeters between predicted and ground-truth 3D joint coordinates. We take as the predicted 3D joint positions the means of the 3D Gaussians produced by *SkelSplat*. Unlike prior works, such as [55], that often report root-relative MPJPE, we consider absolute MPJPE, assessing both joint configuration and global position accuracy without aligning predicted and ground-truth poses at the root joint, since we are interested in the absolute pose in the 3D space. Lower MPJPE values indicate more precise 3D joint predictions.

4.2. Results

Results on H36M We evaluate *SkelSplat* on the Human3.6M dataset to compare against state-of-the-art methods, as shown in Tab. 1. We evaluate three variants of *SkelSplat*, each using different source of 2D keypoints: MeTRAbs [39], CPN [7] (as used in [32]), and ResNet-152 (as in [55]), with the latter two fine-tuned on Human3.6M. Using accurate, dataset-specific 2D keypoints ensures fair comparison with prior works, which fine-tune 2D pose estimators on the dataset of interest [10, 22, 27, 56], as body joints definitions differ across datasets [39]. MeTRAbs, trained on a large variety of datasets, is used off-the-shelf. While using 2D detected poses for optimization can indeed limit quality, this affects all multi-view methods, as they depend on 2D poses for fusion [27]. For the initial 3D pose estimate, we adopt the method proposed in [4], which combines monocular 3D predictions by minimizing 2D reprojection error. An ablation comparing initialization strategies is provided in Sec. 4.3. Among the variants,

the best performance is achieved using 2D keypoints from ResNet-152, which outperforms strong baselines such as [10, 27, 30, 55, 56]. Our method is surpassed only by Isakov et al. [22] which, however, relies on full 3D supervision and extensive training on Human3.6M. In contrast, *SkelSplat* does not require any 3D ground truth during training, optimizing multi-view scenes using only detected 2D keypoints, making it significantly more adaptable.

Cross dataset generalization We evaluate cross-dataset generalization in Tab. 2, where we compare our method to [2, 22, 30] trained on the CMU Panoptic dataset and tested on Human3.6M. Results clearly show that, while learning-based approaches retain excellent performance in the training domain, they significantly drop when applied to different datasets.

Results on CMU Panoptic Studio On this dataset we compare against baselines designed for cross-scene generalization, such as [2, 45], and with Algebraic Triangulation and RANSAC from [22]. We consider a 4-views setup, selecting cameras 1, 2, 10 and 13, following the configuration proposed in [22]. In Tab. 3 we report results obtained using 2D poses from MeTRAbs and considering the triangulation of 2D detections as initial guess. *SkelSplat* achieves an absolute MPJPE of 20.9 mm, outperforming other approaches aimed at cross-scene generalization. Notably, it surpasses both RANSAC and Algebraic Triangulation even when those methods use fine-tuned 2D poses: while these methods struggle with camera configurations that produce significant self-occlusions, *SkelSplat* shows a better ability in handling challenging visual conditions.

Results on Human3.6M-Occ We assess the robustness of *SkelSplat* to occlusions using the Human3.6M-Occ dataset, which introduces synthetic occlusions of varying severity. For joints that are more prone to being occluded, such as elbows, wrists, knees, and ankles, we slightly enlarge their initial covariance by applying a scaling factor of 1.25. This provides more flexibility during optimization, enabling the model to better cope with inaccurate 2D detections and to more effectively explore plausible joint configurations under occlusion. The effect of this covariance scaling is further analyzed in Sec. 4.3. We initialize our 3D pose with monocular predictions fusion as in [4]. As shown in Tab. 4, *SkelSplat* achieves state-of-the-art performance across all three Human3.6M-Occ benchmarks, demonstrating superior robustness compared to view-fusion methods specifically designed to handle occlusions [4, 27, 55]. The best results are obtained using 2D poses from a ResNet-152 model trained solely on the original Human3.6M dataset—the same model used in AdaFuse [55]. In the most challenging Human3.6M-Occ-3-Hard setting,

Absolute MPJPE, mm	3D GT	Dir.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg ↓
Tome et al. [44]	✓	43.3	49.6	42.0	48.8	51.1	64.3	40.3	43.3	66.0	95.2	50.2	52.2	51.1	43.9	45.3	52.8
Cross-view Fusion [36]	✓	28.9	32.5	26.6	28.1	28.3	29.3	28.0	36.8	41.0	30.5	35.6	30.0	28.3	30.0	30.5	31.2
Remelli et al. [38]	✓	27.3	32.1	25.0	26.5	29.3	35.4	28.8	31.6	36.4	31.7	31.2	29.9	26.9	33.7	30.4	30.2
Generalizable Triang. [2]	✓	27.5	28.4	29.3	27.5	30.1	28.1	27.9	30.8	32.9	32.5	30.8	29.4	28.5	30.5	30.1	29.1
Epipolar Transformer [16]	✓	25.7	27.7	23.7	24.8	26.9	31.4	24.9	26.5	28.8	31.7	28.2	26.4	23.6	28.3	23.5	26.9
TransFusion [27]	✓	24.4	26.4	23.4	21.1	25.2	23.2	24.7	33.8	29.8	26.4	26.8	24.2	23.2	26.1	23.3	25.8
AdaFuse [55]	✓	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	24.3
Geometry Transformer [30]	✓	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	22.7
Iskakov et al. - Vol [22]	✓	18.0	18.3	16.5	16.1	17.4	18.2	16.5	18.5	19.4	20.1	18.2	17.4	17.2	19.2	16.6	17.7
UPose3D [10]	✗	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	26.2
TRL [56]	✗	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	25.8
MV Pose Fusion [4]	✗	23.8	24.6	23.4	24.1	26.9	24.0	25.3	28.3	31.7	25.9	26.9	24.2	23.0	27.3	24.2	25.6
Iskakov et al. - Alg [22]	✗	21.7	23.7	22.2	20.4	26.7	24.2	19.9	22.6	31.2	35.6	26.8	21.2	20.9	24.6	21.1	24.5
RANSAC (as in [22])	✗	21.6	22.9	20.9	21.0	23.1	23.0	20.8	22.0	26.4	26.6	24.0	21.5	21.0	23.9	20.8	22.8
SkelSplat (MeTRAbs [39])	✗	24.4	25.9	24.7	25.0	28.2	24.4	25.1	28.2	30.3	26.9	28.7	25.0	26.9	28.1	28.7	26.7
SkelSplat (CPN)	✗	25.7	25.4	25.0	24.6	27.1	23.4	23.7	27.1	28.7	27.1	28.5	24.3	26.3	28.4	26.9	26.2
SkelSplat (ResNet-152)	✗	18.9	20.1	19.5	18.1	20.6	18.6	20.2	22.2	23.5	21.6	20.8	18.8	19.1	22.2	18.6	20.3

Table 1. Overall comparison with state-of-the-art methods on Human3.6M, where ‘✓’ indicates training with 3D GT.

Absolute MPJPE, mm	Train	Test	Avg ↓	Impr. (%) ↑
Geometry-biased transformer [30]	CMU	H36M	38.9	-
Iskakov et al. - Vol [22]	CMU	H36M	34.0	+ 12.6
Generalizable Triang. [2]	CMU	H36M	31.0	+ 20.3
SkelSplat (MeTRAbs [39])	-	H36M	26.0	+ 31.4
SkelSplat (ResNet-152)	-	H36M	20.3	+ 47.8

Table 2. Comparison between models trained on CMU and evaluated on H36M and *SkelSplat* tested on H36M without any training.

Absolute MPJPE, mm	Avg ↓
RANSAC (as in [22])	39.5
Voxelpose [45]	25.5
Generalizable Triangulation [2]	25.4
Iskakov et al. - Alg [22]	21.3
SkelSplat (MeTRAbs [39])	20.9

Table 3. Results on the CMU Panoptic Studio dataset, considering single person sequences in a 4-camera setup.

SkelSplat ranks second only to [4], which benefits from additional pose confidence weighting during fusion. In contrast, *SkelSplat* achieves competitive results without using any confidence scores or learned weights, relying only on the optimization using the 2D pose detections. Compared to standard non-learning baselines such as triangulation and RANSAC, which suffer significant performance degradation under heavy occlusion, *SkelSplat* maintains a high accuracy, showing robustness to incomplete and noisy 2D observations. Some qualitative results are shown in Fig. 4.

Results on Occlusion-Person We further evaluate the occlusion robustness of *SkelSplat* on the Occlusion-Person dataset, including baseline methods such as Algebraic Triangulation, RANSAC, and AdaFuse [55]. For this experiment, we report results using both 8-camera and 4-camera configurations, selecting cameras 1, 3, 5, and 7 for the lat-

ter. As initial pose, we use the triangulation of 2D keypoints predicted by a ResNet-50 model, following the same setup as AdaFuse [55]. *SkelSplat* outperforms all baselines that were not trained on the target dataset, improving upon the best-performing model by 5.3 mm and 5.6 mm, on the 4-view and 8-view settings, respectively. Note that for AdaFuse we report two variants: one trained on Human3.6M and one on Occlusion-Person. While the latter achieves better performance than *SkelSplat*, this comparison is less meaningful, as its performance significantly degrades when applied outside its training domain.

Overall, these experiments highlight *SkelSplat*’s ability to generalize across domains, camera configurations, and occluded conditions while maintaining high accuracy, making it easily applicable to novel environments.

4.3. Ablation Studies

To understand the contribution of each component, we conduct ablation studies analyzing the impact of our one-hot encoding, optimization design choices such as cross-view gradient accumulation and the 3D loss term, input conditions including camera selection, initialization




Figure 4. *SkelSplat* results (blue) on Human3.6M-Occ-3 with ground-truth poses (green). The rightmost column shows a failure case under occlusion, where the left knee is incorrectly predicted.

Absolute MPJPE, mm	H3.6M-Occ-2	H3.6M-Occ-3	H3.6M-Occ-3h
Alg. Triangulation (ResNet-152)	43.2	48.9	120.4
TransFusion [27]	40.8	76.3	96.5
RANSAC (as in [55])	33.7	38.6	80.7
Alg. Triangulation (MeTRAbs [39])	36.0	39.0	67.5
AdaFuse [55]	27.9	31.2	41.1
MV Pose Fusion [4]	33.4	36.7	37.8
SkelSplat (MeTRAbs [39])	29.6	31.1	38.1
SkelSplat (ResNet-152)	24.6	27.0	34.8

Table 4. Results on the three benchmarks of the Human3.6M-Occ dataset, which introduce occlusions on 2 views (Occ-2), 3 views (Occ-3), and severe occlusions on 3 views (Occ-3-Hard).

Absolute MPJPE, mm	4 views	8 views
Alg. Triangulation	59.1	49.2
RANSAC (as in [55])	71.4	39.2
Adafuse [55]†	19.5	12.6
Adafuse [55]	45.6	36.1
SkelSplat (ResNet-50)	40.3	30.5

Table 5. Results on the Occlusion-Person dataset, for 8-camera and 4-camera setups. ‘†’ indicates training on the target dataset.

strategies, and image resolution, as well as covariance scaling in pseudo ground-truth generation.

Impact of one-hot joint encoding On Human3.6M, we evaluate our one-hot encoding for joint Gaussians. We compare it against the standard RGB formulation of [25], which uses three shared channels, and a single-channel variant where all Gaussians are aggregated. As reported in Tab. 6, our one-hot encoding reduces absolute MPJPE by up to 10% relative to these alternatives.

Impact of different cameras We analyze the effect of the number of input views on *SkelSplat* on the CMU Panoptic Studio dataset. In Tab. 8, we report results for varying camera configurations using 4, 5, 6, 7, and 8 views. The 4-camera setup includes views 1, 2, 10, and 13, with additional views progressively added: camera 3 (5 views), camera 23 (6 views), camera 19 (7 views), and camera 30 (8 views). The approach clearly improves by increasing the number of views. We observe a consistent reduction in error as more cameras are added, with a slight deviation in the 7-view setup due to oblique angle and subject cropping in camera 19. Overall, this trend suggests that *SkelSplat* performance improves in dense camera scenarios.

Contribution of cross-view gradient accumulation On Human3.6M, we evaluate the impact of accumulating optimization gradients across multiple views. As shown in Tab. 7, we compare standard single-view gradient updates (as in the original Gaussian Splatting) with gradient accumulation over 2 views and 4 views, with the latter

Absolute MPJPE, mm	3 channels [25]	1 channel	One-hot enc.	Impr. (%) ↑
SkelSplat (MeTRAbs [39])	29.4	29.4	26.7	+ 10.1
SkelSplat (ResNet-152)	21.9	21.9	20.3	+ 7.9

Table 6. Ablation on joint encoding strategies for rendering.

Absolute MPJPE, mm	1 view	2 views	4 views
SkelSplat (MeTRAbs [39])	30.9	28.2	26.7
SkelSplat (ResNet-152)	23.0	21.5	20.3

Table 7. Results on Human3.6M accumulating gradient across 1, 2 or 4 views.

Absolute MPJPE, mm	4 views	5 views	6 views	7 views	8 views
SkelSplat (MeTRAbs [39])	20.9	17.7	14.7	15.7	15.6

Table 8. Impact of using different number of views for *SkelSplat* optimization. Results on the CMU Panoptic Studio dataset.

Init. method	2D poses	Guess	SkelSplat	Impr. (%) ↑
Triangulation	CPN	30.6	27.4	+ 10.5
	MeTRAbs [39]	29.0	26.7	+ 7.9
	ResNet-152	23.7	20.7	+ 12.7
3D Fusion [4]	CPN	40.2	26.2	+ 34.8
	MeTRAbs [39]	40.2	26.7	+ 33.6
	ResNet-152	40.2	20.3	+ 49.5

Table 9. Human3.6M results with different initialization strategies for 3D joint positions.

yielding the best performance. Notably, even accumulating over just 2 views results in a 2.7 mm improvement over the standard approach, highlighting that incorporating information from multiple views before parameter updates promotes optimization stability and cross-view consistency.

Impact of initialization and noise On Human3.6M, we analyze how initialization affects 3D joint optimization. In Tab. 9, we compare *SkelSplat* initialized with simple algebraic triangulation and with multi-view fusion of 3D monocular poses from an off-the-shelf estimator [4], using 2D detections from MeTRAbs [39], CPN and ResNet-152. Results show that the initialization strategy has little influence on the final accuracy, as different starting points converge to comparable MPJPE values. To further assess robustness, we perturb triangulated initializations with Gaussian noise of increasing standard deviation applied independently to each joint, as reported in Tab. 10. Performance remains stable up to 40 mm noise, degrades moderately at 60 mm, and drops more sharply beyond 80-100 mm. This shows that *SkelSplat* is resilient to realistic levels of initialization error, while reaching its limits under extreme perturbations.

Noise σ , mm	0	10	20	40	60	80	100
SkelSplat (MeTRAbs [39])	26.7	27.8	27.9	30.6	41.5	63.4	84.6
Variation (%)	-	+4.1	+4.5	+14.6	+55.4	+137.5	+216.8
SkelSplat (ResNet-152)	20.3	21.5	22.0	25.4	38.3	61.9	83.6
Variation (%)	-	+5.9	+8.4	+25.1	+88.7	+204.9	+311.3

Table 10. Ablation on robustness to poor initialization, absolute MPJPE and relative variation under increasing Gaussian noise.

Loss contribution to the final prediction We conduct an ablation on the effect of the 3D structural symmetry loss applied to different subsets of body limbs, considering Human3.6M and its occluded variants. The baseline setting (Symm-1), used in our proposed *SkelSplat*, applies the symmetry constraint to the lower arms and lower legs. We also evaluate extending the constraint to the upper arms and legs (Symm-2), and further to the torso limbs (Symm-3). On Human3.6M-Occ-2, Symm-1 yields the largest accuracy gain (+6.1%). Symm-2 and Symm-3 provide marginal improvements with respect to it (+3.3% and +3.7%) but at higher cost, with Symm-3 increasing optimization time by +43.5% over Symm-1. These results confirm Symm-1 as the best trade-off between performance and efficiency and is our proposed solution. More details are provided in the supplementary.

Effect of rendering quality We evaluate how rendering resolution affects pose accuracy by downscaling input images to 75%, 50%, and 25% of their original size. Higher resolutions (100-75%) yield the best performance with 20.3 and 20.4 mm errors on Human3.6M, while accuracy degrades slightly at lower scales, with errors of 20.7 and 21.7 mm at 50% and 25%. This demonstrates that *SkelSplat* benefits from high-resolution inputs but remains robust to reduced rendering quality. Further details are provided in the supplementary.

Impact of 2D covariance on pseudo ground truth We also evaluate the effect of enlarging the covariance for frequently occluded joints, such as elbows, hands, and knees, on Human3.6-Occ-2 and Human3.6-Occ-3. A moderate increase (1.25 \times) allows the model to better accommodate uncertainty in occluded regions, improving robustness. Larger covariances, however, reduce precision (+4.8% error on average at 1.5 \times and nearly doubled error at 2 \times), as optimization becomes overly tolerant to noisy inputs. This highlights the need to balance flexibility and spatial accuracy for optimal performance. Extended ablation results are provided in the supplementary.

4.4. Discussion

Building on the success of Gaussian Splatting in dense 3D reconstruction [19, 33, 47, 48], where continuous optimization over a differentiable representation has proven to be

more effective than direct geometric methods, we reinterpret multi-view pose estimation as a reconstruction problem. Unlike traditional Gaussian Splatting pipelines supervised by appearance, our supervision is driven by skeletal pose features. The presented experiments focus mainly on proving the accuracy of the skeleton represented by the means of the Gaussians. However, we believe covariance can potentially give useful information on the uncertainty of the prediction. For example, we measure the percentage of ground-truth 3D joints that lie within 1, 2, or 3 standard deviations from the corresponding Gaussian means on the Human3.6M dataset. We find that on average 98.0% of joints fall within 3 sigmas, 91.6% within 2 sigmas, and 57.2% within 1 sigma. Lower coverage is typically observed for joints more prone to occlusions and self-occlusion, such as hands, elbows, and ankles. Further details are provided in the supplementary. These results indicate that covariance can potentially include an estimate of the uncertainty of the joints. However, more experiments are needed to prove the correlation. While this work focuses on robust pose estimation, inference speed may be a limiting factor for real-time applications. The current implementation requires about 3 seconds to estimate a pose from 4 views, as each joint is rendered in its own channel to allow independent optimization. This motivates future work on more compact rendering schemes and alternative joint encodings that could significantly accelerate inference. Extending *SkelSplat* to multi-person scenarios by incorporating instance association across views to provide per-person 3D initialization, would further broaden its applicability in real-world environments.

5. Conclusions

We presented *SkelSplat*, a novel framework for multi-view 3D human pose estimation that departs from standard learning-based fusion strategies and instead leverages Gaussian Splatting for robust 3D human pose reconstruction. Unlike current methods, which rely heavily on training data and are tied to specific camera setups, pose distributions, or occlusion conditions, *SkelSplat* is inherently flexible and generalizes effectively to novel environments without retraining or fine-tuning. The experiments demonstrated that *SkelSplat* achieves state-of-the-art performance across multiple benchmarks, showing strong robustness to occlusions and adaptability to different camera configurations, including out-of-distribution scenarios. While our method offers accuracy and generalizability, it is not currently real-time as it operates by rendering a separate channel for each joint. Future work will explore more efficient formulations through compact joint encodings, while also investigating the combination of 3D human reconstruction with skeleton estimation, in particular using 3D reconstruction losses for 2D refinement to improve robustness. We further plan to extend *SkelSplat* to multi-person scenarios, broadening its applicability in real-world settings.

References

- [1] Leonardo Barcellona, Andrii Zadaianchuk, Davide Allegro, Samuele Papa, Stefano Ghidoni, and Efstratios Gavves. Dream to manipulate: Compositional world models empowering robot imitation learning with imagination. In *The Thirteenth International Conference on Learning Representations*, 2025. 2
- [2] Kristijan Bartol, David Bojanić, Tomislav Petković, and Tomislav Pribanić. Generalizable human pose triangulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11028–11037, 2022. 5, 6
- [3] 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. 2
- [4] Laura Bragagnolo, Matteo Terreran, Davide Allegro, and Stefano Ghidoni. Multi-view pose fusion for occlusion-aware 3d human pose estimation. In *European Conference on Computer Vision*, pages 117–133. Springer, 2025. 1, 2, 3, 5, 6, 7, 12
- [5] Anjun Chen, Xiangyu Wang, Zhi Xu, Kun Shi, Yan Qin, Yuchi Huo, Jiming Chen, and Qi Ye. Adaptivefusion: Adaptive multi-modal multi-view fusion for 3d human body reconstruction. *arXiv preprint arXiv:2409.04851*, 2024. 2
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 2
- [7] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7103–7112, 2018. 5
- [8] Zhuo Chen, Xu Zhao, and Xiaoyue Wan. Structural triangulation: A closed-form solution to constrained 3d human pose estimation. In *European Conference on Computer Vision*, pages 695–711. Springer, 2022. 2
- [9] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1290–1299, 2022. 2
- [10] Vanda Davoodnia, Saeed Ghorbani, Marc-André Carbonneau, Alexandre Messier, and Ali Etemad. Upose3d: Uncertainty-aware 3d human pose estimation with cross-view and temporal cues. In *European Conference on Computer Vision*, pages 19–38. Springer, 2024. 2, 5, 6
- [11] Arnab Dey, Cheng-You Lu, Andrew I Comport, Srinath Sridhar, Chin-Teng Lin, and Jean Martinet. Hfgaussian: Learning generalizable gaussian human with integrated human features. *arXiv preprint arXiv:2411.03086*, 2024. 2
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [13] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>. 5
- [14] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5354–5363, 2024. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2
- [16] Yihui He, Rui Yan, Katerina Fragkiadaki, and Shoou-I Yu. Epipolar transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7779–7788, 2020. 1, 2, 6
- [17] Lukas Höllein, Aljaž Božič, Michael Zollhöfer, and Matthias Nießner. 3dgs-lm: Faster gaussian-splatting optimization with levenberg-marquardt. *arXiv preprint arXiv:2409.12892*, 2024. 2
- [18] Liangxiao Hu, Hongwen Zhang, Yuxiang Zhang, Boyao Zhou, Boning Liu, Shengping Zhang, and Liqiang Nie. Gaussianavatar: Towards realistic human avatar modeling from a single video via animatable 3d gaussians. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 634–644, 2024. 2
- [19] Shoukang Hu, Tao Hu, and Ziwei Liu. Gauhuman: Articulated gaussian splatting from monocular human videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20418–20431, 2024. 2, 8
- [20] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers*, pages 1–11, 2024. 2
- [21] 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. 1, 2, 4, 12
- [22] Karim Iskakov, 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. 1, 2, 5, 6
- [23] Hanxiao Jiang, Hao-Yu Hsu, Kaifeng Zhang, Hsin-Ni Yu, Shenlong Wang, and Yunzhu Li. Phystwin: Physics-informed reconstruction and simulation of deformable objects from videos. *arXiv preprint arXiv:2503.17973*, 2025. 2

- [24] Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *Proceedings of the IEEE international conference on computer vision*, pages 3334–3342, 2015. 1, 2, 4
- [25] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 1, 2, 3, 7
- [26] Muhammed Kocabas, Jen-Hao Rick Chang, James Gabriel, Oncel Tuzel, and Anurag Ranjan. Hugs: Human gaussian splats. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 505–515, 2024. 2
- [27] Haoyu Ma, Liangjian Chen, Deying Kong, Zhe Wang, Xingwei Liu, Hao Tang, Xiangyi Yan, Yusheng Xie, Shih-Yao Lin, and Xiaohui Xie. Transfusion: Cross-view fusion with transformer for 3d human pose estimation. *arXiv preprint arXiv:2110.09554*, 2021. 2, 5, 6, 7
- [28] 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
- [29] Hidenobu Matsuki, Riku Murai, Paul HJ Kelly, and Andrew J Davison. Gaussian splatting slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18039–18048, 2024. 2
- [30] Olivier Moliner, Sangxia Huang, and Kalle Åström. Geometry-biased transformer for robust multi-view 3d human pose reconstruction. In *2024 IEEE 18th International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–8. IEEE, 2024. 1, 2, 5, 6
- [31] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016. 2
- [32] Dario Pavullo, 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. 5
- [33] Sida Peng, Chen Geng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Implicit neural representations with structured latent codes for human body modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8):9895–9907, 2023. 2, 8
- [34] Lorenza Prospero, Abdullah Hamdi, Joao F Henriques, and Christian Rupprecht. Gst: Precise 3d human body from a single image with gaussian splatting transformers. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6007–6017, 2025. 2
- [35] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20051–20060, 2024. 2
- [36] Haibo Qiu, Chunyu Wang, Jingdong Wang, Naiyan Wang, and Wenjun Zeng. Cross view fusion for 3d human pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4342–4351, 2019. 6
- [37] Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 497–500, 2001. 3
- [38] Edoardo Remelli, Shangchen Han, Sina Honari, Pascal Fua, and Robert Wang. Lightweight multi-view 3d pose estimation through camera-disentangled representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6040–6049, 2020. 1, 6
- [39] István Sárádi, Alexander Hermans, and Bastian Leibe. Learning 3d human pose estimation from dozens of datasets using a geometry-aware autoencoder to bridge between skeleton formats. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2956–2966, 2023. 5, 6, 7, 8, 12, 13
- [40] Adam Sun, Tiange Xiang, Scott Delp, Li Fei-Fei, and Ehsan Adeli. Occfusion: Rendering occluded humans with generative diffusion priors. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. 2
- [41] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5693–5703, 2019. 2
- [42] David Svitov, Pietro Morerio, Lourdes Agapito, and Alessio Del Bue. Haha: Highly articulated gaussian human avatars with textured mesh prior. In *Proceedings of the Asian Conference on Computer Vision*, pages 4051–4068, 2024. 2
- [43] Matteo Terreran, Leonardo Barcellona, and Stefano Ghidoni. A general skeleton-based action and gesture recognition framework for human–robot collaboration. *Robotics and Autonomous Systems*, 170:104523, 2023. 1
- [44] Denis Tome, Matteo Toso, Lourdes Agapito, and Chris Russell. Rethinking pose in 3d: Multi-stage refinement and recovery for markerless motion capture. In *2018 international conference on 3D vision (3DV)*, pages 474–483. IEEE, 2018. 6
- [45] Hanyue Tu, Chunyu Wang, and Wenjun Zeng. Voxelpose: Towards multi-camera 3d human pose estimation in wild environment. In *European conference on computer vision*, pages 197–212. Springer, 2020. 5, 6
- [46] Xiaoyue Wan, Zhuo Chen, and Xu Zhao. View consistency aware holistic triangulation for 3d human pose estimation. *Computer Vision and Image Understanding*, 236:103830, 2023. 1
- [47] Dongxu Wei, Zhiqi Li, and Peidong Liu. Omni-scene: Omni-gaussian representation for ego-centric sparse-view scene reconstruction. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 22317–22327, 2025. 8
- [48] Jing Wen, Xiaoming Zhao, Zhongzheng Ren, Alexander G Schwing, and Shenlong Wang. Gomavatar: Efficient animatable human modeling from monocular video using

- gaussians-on-mesh. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2059–2069, 2024. [2](#), [8](#)
- [49] Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. Monocular total capture: Posing face, body, and hands in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [5](#)
- [50] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision (ECCV)*, pages 466–481, 2018. [2](#)
- [51] Junjin Xiao, Qing Zhang, Yonewei Nie, Lei Zhu, and Wei-Shi Zheng. Rogsplat: Learning robust generalizable human gaussian splatting from sparse multi-view images. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 5980–5990, 2025. [2](#)
- [52] Junjin Xiao, Qing Zhang, Yonewei Nie, Lei Zhu, and Wei-Shi Zheng. Rogsplat: Learning robust generalizable human gaussian splatting from sparse multi-view images. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5980–5990, 2025. [2](#)
- [53] Suhang Ye, Yingyi Zhang, Jie Hu, Liujuan Cao, Shengchuan Zhang, Lei Shen, Jun Wang, Shouhong Ding, and Rongrong Ji. Distilpose: Tokenized pose regression with heatmap distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2163–2172, 2023. [2](#)
- [54] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19447–19456, 2024. [2](#)
- [55] Zhe Zhang, Chunyu Wang, Weichao Qiu, Wenhui Qin, and Wenjun Zeng. Adafuse: Adaptive multiview fusion for accurate human pose estimation in the wild. *International Journal of Computer Vision*, 129:703–718, 2021. [1](#), [2](#), [5](#), [6](#), [7](#)
- [56] Jiachen Zhao, Tao Yu, Liang An, Yipeng Huang, Fang Deng, and Qionghai Dai. Triangulation residual loss for data-efficient 3d pose estimation. *Advances in neural information processing systems*, 36:12721–12732, 2023. [2](#), [5](#), [6](#)
- [57] Jingxiao Zheng, Xinwei Shi, Alexander Gorban, Junhua Mao, Yang Song, Charles R Qi, Ting Liu, Visesh Chari, Andre Cornman, Yin Zhou, et al. Multi-modal 3d human pose estimation with 2d weak supervision in autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4478–4487, 2022. [1](#)
- [58] Shenhao Zhu, Junming Leo Chen, Zuozhuo Dai, Zilong Dong, Yinghui Xu, Xun Cao, Yao Yao, Hao Zhu, and Siyu Zhu. Champ: Controllable and consistent human image animation with 3d parametric guidance. In *European Conference on Computer Vision*, pages 145–162. Springer, 2024. [1](#)
- [59] Christian Zimmermann, Tim Welschhold, Christian Dornhege, Wolfram Burgard, and Thomas Brox. 3d human pose estimation in rgb-d images for robotic task learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1986–1992. IEEE, 2018. [1](#)

SkelSplat: Robust Multi-view 3D Human Pose Estimation with Differentiable Gaussian Rendering

Supplementary Materials

A. Pseudo Ground Truth Generation

This section provides additional details on how we generate the pseudo ground truth used during the optimization of *SkelSplat*. We begin by recalling that a human skeleton is defined as a set of joints $SK = \{sk_0, \dots, sk_N\}$, each of which is represented by an anisotropic 3D Gaussian. This yields a corresponding set of Gaussians $GS = \{gs_0, \dots, gs_N\}$, where each Gaussian gs_j encodes the spatial uncertainty around joint sk_j in 3D space. The optimization process leverages 2D keypoint detections obtained from a pre-trained 2D human pose estimator. Specifically, for each camera view i in a set of M synchronized and calibrated views (i.e., $i \in \{1, \dots, M\}$), we extract 2D keypoint locations $SK_i^{2D} = \{sk_{i0}^{2D}, \dots, sk_{iN}^{2D}\}$ corresponding to the projection of each joint in the image plane. To supervise the optimization with view-dependent supervision, we generate a set of pseudo ground truth heatmaps $\{I_{ij}\}_{j=1}^N$ for each camera view. Each heatmap $I_{ij} \in \mathbb{R}^{H \times W}$ represents a soft target for joint j in view i , and is constructed by rendering a 2D Gaussian gs_{ij}^{2D} centered at the detected 2D location $sk_{ij}^{2D} \in \mathbb{R}^2$. The shape of this Gaussian is determined by a covariance matrix $\Sigma_{ij}^{2D} \in \mathbb{R}^{2 \times 2}$, which is obtained by projecting the original 3D covariance Σ_j of joint j into the 2D image plane of camera i as follows:

$$\Sigma_{ij}^{2D} = J_i W_i \Sigma_j W_i^T J_i^T, \quad (6)$$

where $W_i \in \mathbb{R}^{4 \times 4}$ is the camera extrinsic transformation (world-to-camera), and $J_i \in \mathbb{R}^{2 \times 3}$ is the Jacobian matrix of the perspective projection at the joint position in camera coordinates. In detail, if $\mu_j \in \mathbb{R}^3$ is the 3D joint position, its homogeneous coordinate $\tilde{\mu}_j = [\mu_j^\top, 1]^\top$ is transformed into camera coordinates as:

$$\tilde{\mu}_{ij}^{\text{cam}} = W_i \tilde{\mu}_j = \begin{bmatrix} X_{ij} \\ Y_{ij} \\ Z_{ij} \\ 1 \end{bmatrix}. \quad (7)$$

The Jacobian J_i for the perspective projection with focal lengths $f_{x,i}, f_{y,i}$ is:

$$J_i = \begin{bmatrix} \frac{f_{x,i}}{Z_{ij}} & 0 & -\frac{f_{x,i} X_{ij}}{Z_{ij}^2} \\ 0 & \frac{f_{y,i}}{Z_{ij}} & -\frac{f_{y,i} Y_{ij}}{Z_{ij}^2} \end{bmatrix}. \quad (8)$$

A small constant h is added to the covariance to prevent numerical issues. To characterize the shape of the 2D covariance ellipse, we compute the eigenvalues λ_1 and λ_2 of Σ_{ij}^{2D} , which correspond to the principal axes variances:

$$\det = \Sigma_{ij}^{2D}(1, 1) \cdot \Sigma_{ij}^{2D}(2, 2) - (\Sigma_{ij}^{2D}(1, 2))^2, \quad (9)$$

$$m = \frac{\Sigma_{ij}^{2D}(1, 1) + \Sigma_{ij}^{2D}(2, 2)}{2}, \quad (10)$$

$$\lambda_1 = m + \sqrt{\max(\epsilon, m^2 - \det)}, \quad (11)$$

$$\lambda_2 = m - \sqrt{\max(\epsilon, m^2 - \det)}, \quad (12)$$

where ϵ is a small positive constant to ensure numerical stability, (s, q) indicates the element with s, q coordinates in the matrix.

B. Additional Details on Ablation Studies

In this section, we present further details on the ablation studies discussed in the main paper. Specifically, we include expanded analyses of noise robustness (as shown in Fig. 5), loss component contributions (see Tab. 12), covariance scaling behaviors (see Tab. 13), and the effects of rendering resolution (see Tab. 11).

Robustness to noisy initialization On Human3.6M [21], we perturb triangulation of MeTRAbs [39] and ResNet-152 poses with Gaussian noise of increasing standard deviation applied independently to each joint. We consider values for σ equal to 10, 20, 40, 60, 80 and 100 mm. From Fig. 5 we can observe how performance remains stable up to 40 mm noise and starts to degrade at 60 mm. With strong noise (80-100 mm) accuracy drops more sharply.

Contribution of 3D loss to optimization Results in Tab. 12 show results for *SkelSplat* using three variants of 3D symmetry loss during optimization: Symm-1 applies the symmetry constraint to the lower arms and lower legs, Symm-2 extends the constraint to the upper arms and legs and Symm-3 further adds constraints from the hip joints to the root and from the shoulders to the neck. Tab. 12 illustrates accuracy on Human3.6M [21], Human3.6M-Occ-2 and Human3.6M-Occ-3 [4].

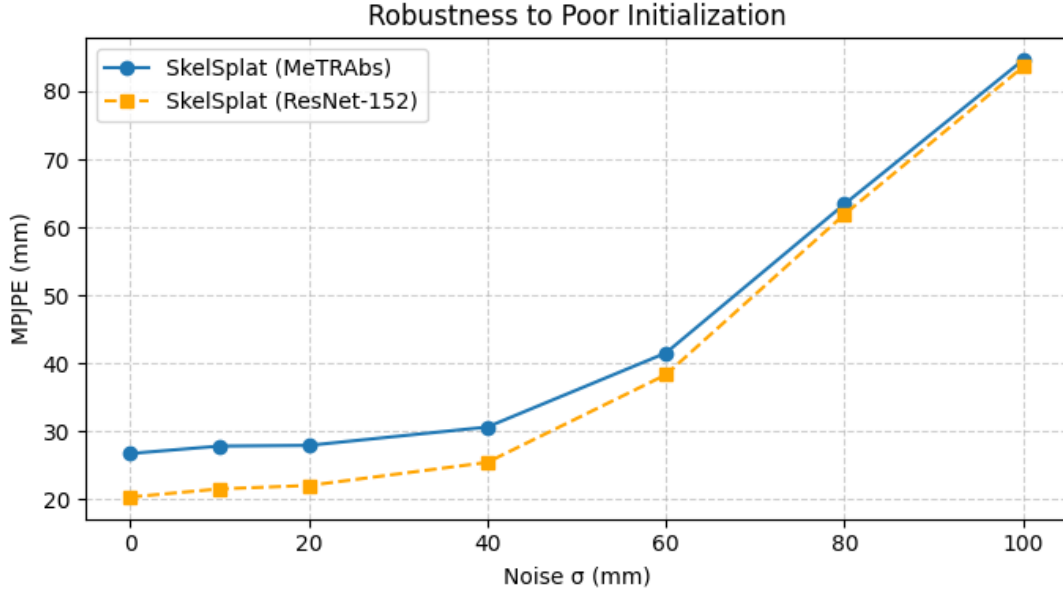


Figure 5. Ablation on robustness to poor initialization, adding Gaussian noise to triangulated joints.

Image resolution scale (%)	100	75	50	25
SkelSplat (MeTRAbs [39])	26.7	26.7	<u>26.8</u>	27.5
SkelSplat (ResNet-152)	20.3	<u>20.4</u>	20.7	21.7

Table 11. Resolution ablation. We downscale inputs to a percentage of the original size before rendering; higher resolution yields slightly lower MPJPE (mm).

Effect of rendering resolution We assess the impact of rendering resolution on pose accuracy by reducing the input images to 75%, 50%, and 25% of their original resolution. Tab. 11 reports full results for *SkelSplat* using 2D detections from MeTRAbs and ResNet-152.

Impact of covariance scaling for 2D pseudo ground truth generation We evaluate the effect of enlarging the covariance for frequently occluded joints, such as elbows, hands, and knees, on Human3.6-Occ-2 and Human3.6-Occ-3. Tab. 13 reports results using scaling factors of 1.25, 1.5, and 2 applied to the default covariance of these joints. Overly large covariances decrease reconstruction accuracy, since the optimization becomes too tolerant of noisy or imprecise 2D inputs.

C. Joint Covariance for Confidence Estimation

While our experiments primarily evaluate the 3D means of the Gaussians as joint predictions, the associated covariances also encode potentially useful information about prediction uncertainty. To assess this, we compute the percentage of ground-truth joints that fall within 1, 2, or 3 stan-

Absolute MPJPE, mm					
Human3.6M	MeTRAbs [39]	27.0	26.7	26.9	<u>26.8</u>
	ResNet-152	20.6	20.3	<u>20.4</u>	<u>20.4</u>
Human3.6M-Occ-2	MeTRAbs [39]	30.0	29.6	<u>29.5</u>	29.4
	ResNet-152	26.2	24.6	<u>23.8</u>	23.7
Human3.6M-Occ-3	MeTRAbs [39]	31.4	<u>31.1</u>	31.0	31.0
	ResNet-152	27.2	27.0	26.0	<u>26.1</u>
3D Symm-1	-	✓	✓	✓	✓
3D Symm-2	-	-	✓	✓	✓
3D Symm-3	-	-	-	-	✓
Sec/iter		0.028	<u>0.039</u>	0.045	0.056

Table 12. Ablation study on different 3D symmetry loss contributions on the Human3.6M dataset and its occluded versions.

Scaling factor		1.25	1.5	2.0
H3.6M-Occ-2	MeTRAbs [39]	29.6	<u>31.0</u>	60.5
	ResNet-152	24.6	<u>26.0</u>	54.3
H3.6M-Occ-3	MeTRAbs [39]	31.1	<u>32.6</u>	61.9
	ResNet-152	27.0	<u>28.9</u>	59.0

Table 13. Ablation study on different methods to initialize 3D joint positions, absolute MPJPE (mm).

dard deviations of the predicted Gaussian means on the Human3.6M dataset and on its occluded version Human3.6M-Occ-3. Here, in Fig. 6 and Fig. 7 we include results for joint-wise coverage in both occluded and non-occluded settings. Notably, in both cases, joints that are often occluded or self-occluded, such as hands, elbows, and ankles, tend to exhibit lower coverage.




Figure 6. Joint-wise coverage rates across different sigma thresholds for the Human3.6M dataset.




Figure 7. Joint-wise coverage rates across different sigma thresholds for the Human3.6M-Occ-3 dataset.

D. Supplementary Visualizations

In Fig. 8, we provide a visualization for a scene with four camera views from Human3.6M-Occ-3. We show the aggregated pseudo ground truth heatmap, obtained by summing the per-joint 2D heatmaps across all joints, together with a comparison between the initial and the final (optimized) 3D joint Gaussians. This highlights how the optimization step progressively refines the 3D pose, leading to improved alignment with the set of multi-view 2D detec-

tions and producing a more coherent reconstruction.

In Fig. 9, we report additional qualitative results on Human3.6M, Human3.6M-Occ-3, and the CMU Panoptic Studio datasets. For Human3.6M-Occ-3, we show only one of the occluded camera views. final row presents representative failure cases, in which our method struggles due to factors such as extreme occlusion or complex limb configurations, for which inconsistent multi-view evidence is common.

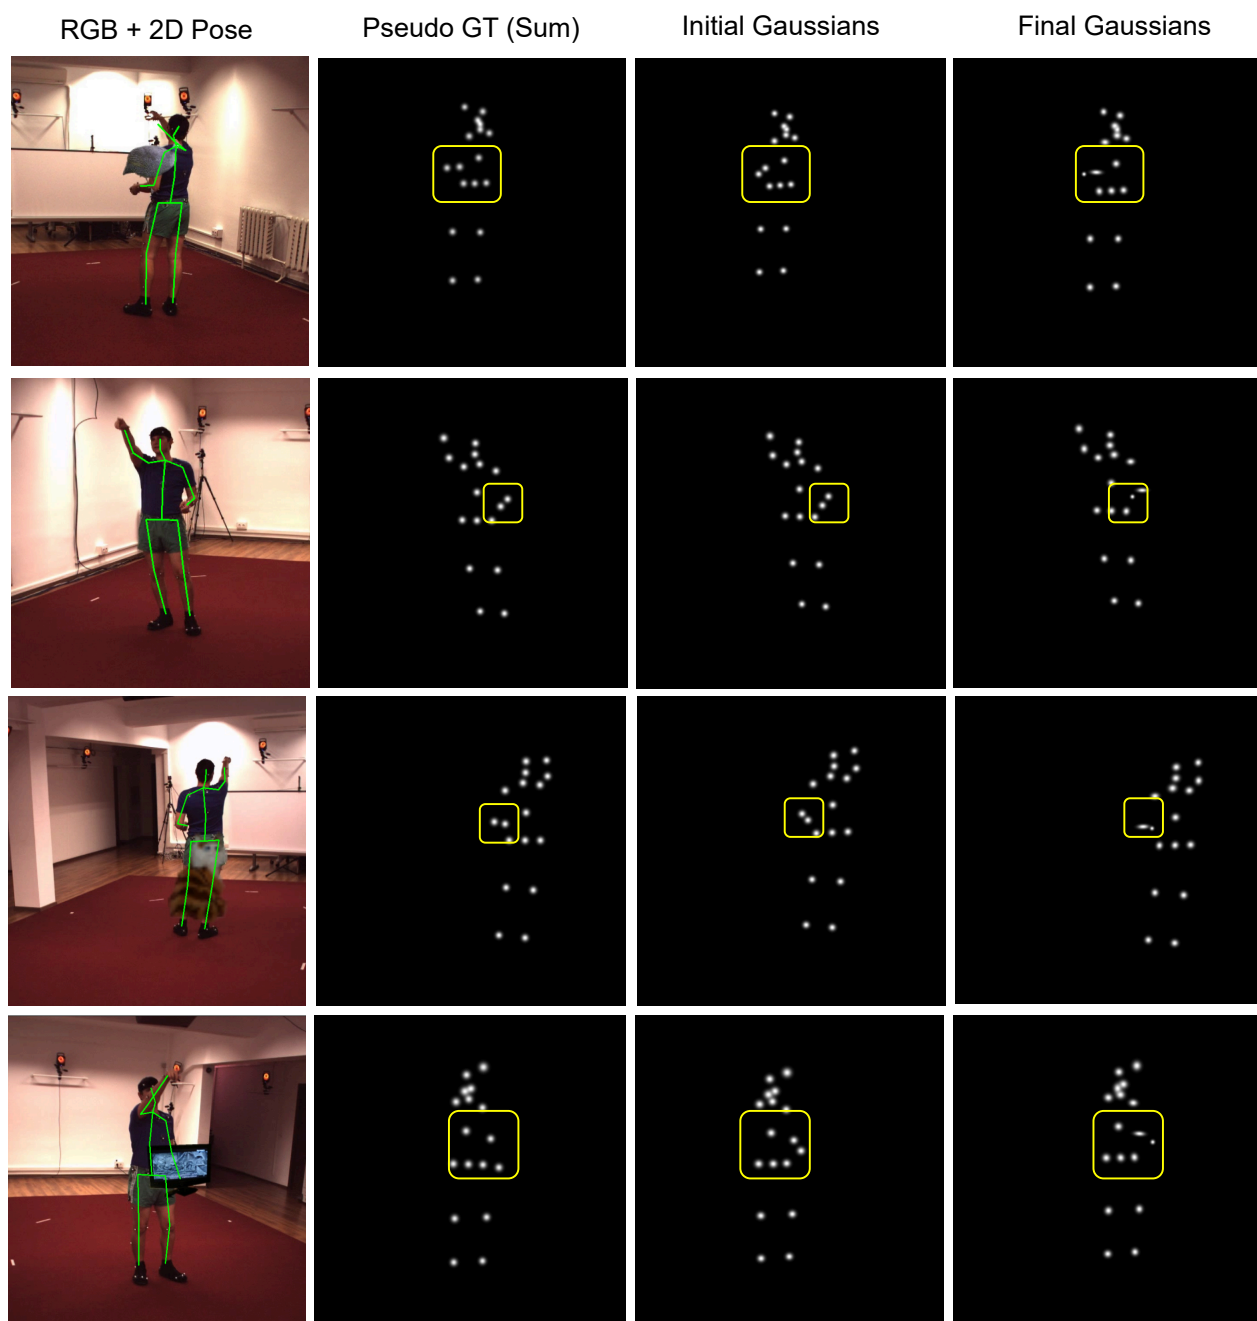


Figure 8. Visualization of pseudo ground-truth supervision and joint refinement. For 4 camera views, we show (left) the aggregated pseudo ground-truth heatmap obtained by summing the 2D Gaussian heatmaps of all joints, and (right) a 3D visualization of the joint Gaussians before and after optimization.




Figure 9. Qualitative results on Human3.6M (left), Human3.6M-Occ-3 (middle), and CMU Panoptic (right). For Human3.6M-Occ-3 we show one of the three occluded views. The last row shows some failure cases.