

DyCrowd: Towards Dynamic Crowd Reconstruction from a Large-scene Video

Hao Wen[†], Hongbo Kang[†], Jian Ma, Jing Huang, Yuanwang Yang, Haozhe Lin,
Yu-Kun Lai^{*} *Senior Member, IEEE*, Kun Li^{*} *Senior Member, IEEE*

Abstract—3D reconstruction of dynamic crowds in large scenes has become increasingly important for applications such as city surveillance and crowd analysis. However, current works attempt to reconstruct 3D crowds from a static image, causing a lack of temporal consistency and inability to alleviate the typical impact caused by occlusions. In this paper, we propose DyCrowd, the first framework for spatio-temporally consistent 3D reconstruction of hundreds of individuals’ poses, positions and shapes from a large-scene video. We design a coarse-to-fine group-guided motion optimization strategy for occlusion-robust crowd reconstruction in large scenes. To address temporal instability and severe occlusions, we further incorporate a VAE (Variational Autoencoder)-based human motion prior along with a segment-level group-guided optimization. The core of our strategy leverages collective crowd behavior to address long-term dynamic occlusions. By jointly optimizing the motion sequences of individuals with similar motion segments and combining this with the proposed Asynchronous Motion Consistency (AMC) loss, we enable high-quality unoccluded motion segments to guide the motion recovery of occluded ones, ensuring robust and plausible motion recovery even in the presence of temporal desynchronization and rhythmic inconsistencies. Additionally, in order to fill the gap of no existing well-annotated large-scene video dataset, we contribute a virtual benchmark dataset, *VirtualCrowd*, for evaluating dynamic crowd reconstruction from large-scene videos. Experimental results demonstrate that the proposed method achieves state-of-the-art performance in the large-scene dynamic crowd reconstruction task. The code and dataset will be available for research purposes.

Index Terms—3D reconstruction, dynamic crowd, large-scene video, multi-person pose and shape estimation.

I. INTRODUCTION

3D human pose and shape estimation in large scenes is essential for enabling effective behavioral modeling and crowd monitoring. This capability plays a pivotal role in a wide range of applications, from public safety and urban planning to large-scale event management. However, current efforts [1], [2] related to large-scene 3D pedestrians reconstruction are based on a single image, capturing static crowds’ 3D positions, postures and shapes. While these methods can provide useful

insights into the spatial distribution and individual characteristics of crowds, they fall short when it comes to representing temporal dynamics and continuous movement patterns, which are critical for understanding crowd behavior over time. This paper aims to reconstruct dynamic crowds within a global camera space from a large-scene video.

Most single-person reconstruction methods estimate temporally consistent and smooth pose and shape by using various temporal modeling techniques [3]–[6] or optimization paradigms [7], [8] from a single RGB video. To extend to large-scene crowd videos, an intuitive idea is to combine these methods with multi-object tracking techniques and independently reconstruct multiple people from cropped instance images. However, this approach ignores interactions among individuals and overlooks critical location and scale information within the global coordinate system of large scenes. For video-based multi-person reconstruction works, the regression-based approaches [9], [10] reconstruct multiple people at once from a video, but they fail to cope with large scenes because scaling the high-resolution images of large scenes down to input resolution would result in the loss of image features for most small and medium-sized individuals. Moreover, insufficient data availability hinders regression-based methods from effectively reconstructing crowds from large-scene videos. Additionally, existing optimization-based methods [11], [12], designed specifically for small scenes with handheld cameras, cannot be directly applied to large-scene videos due to unique challenges such as a large number of people, severe scale variation, wide spatial distribution, and frequent dynamic occlusions.

Crowd3D [1] and GroupRec [2] are two state-of-the-art methods for multi-person reconstruction in large scenes, predicting human body meshes with absolute positions for hundreds of pedestrians within a unified camera system from a single large-scene image. However, when applied to each frame of a large-scene video, these methods often yield temporally-unstable and unrealistic 3D motion and frequently lose objects due to severe mutual occlusion. Occlusions frequently occur and may last for a long time in large scene videos due to the high density of moving individuals, causing spatio-temporal discontinuity and thereby hampering dynamic human reconstruction in large scenes.

In this paper, we propose DyCrowd, the first framework capable of reconstructing the 3D positions, poses, and shapes of dynamic crowds from large-scene videos. To achieve robust and plausible crowd reconstruction in large scenes, we design a Coarse-to-fine Group-guided Motion (C.G.M.) optimization

[†] Equal contribution.

^{*} Corresponding author.

Hao Wen, Hongbo Kang, Jian Ma, Jing Huang, Yuanwang Yang and Kun Li are with the College of Intelligence and Computing, Tianjin University, Tianjin 300350, China. E-mail: {wenhao, hbkang, jianma, hj00, yyw, lik}@tju.edu.cn

Haozhe Lin is with Beijing National Research Center for Information Science and Technology, Tsinghua University, Beijing 10084, China. E-mail: linhz@tsinghua.edu.cn

Yu-Kun Lai is with the School of Computer Science and Informatics, Cardiff University, Cardiff CF24 4AG, United Kingdom. E-mail: Yukun.Lai@cs.cardiff.ac.uk




Fig. 1. Spatio-temporally consistent 3D crowd reconstruction from a large-scene video via the proposed DyCrowd. This approach effectively achieves dynamic 3D reconstruction in a global space and addresses frequent dynamic occlusions in large scenes. Even when certain subjects (e.g., subject 1 in the figure) remain segment-level occluded throughout the entire video, our method can still reconstruct reasonable motion trajectories and poses through optimization.

strategy, which includes VAE-based human motion prior optimization and segment-level group-guided optimization. This strategy fully leverages the collective behavior of crowds in large scenes [13]–[15]: (1) pedestrians with similar trajectories often exhibit similar motion postures, and (2) severely occluded pedestrians and unoccluded pedestrians frequently coexist in the same scene. To effectively leverage the collective motion within groups, unlike GroupRec [2], which reconstructs crowds through collective behavior in static images, we focus on the dynamic changes of the group across consecutive frames and the commonalities in motion among individuals. The goal is to utilize high-quality and unoccluded motion data to assist in recovering the motion of occluded pedestrians. Therefore, we first cluster individual motion segments based on trajectory similarity and introduce detection scores and joint importance to determine the confidence of motion segments, identifying those requiring correction and high-quality, unoccluded motion segments. It is important to note that a segment of an individual’s motion may be mixed with other segments from the same or different individuals in the same group. To address issues such as temporal misalignment and rhythm variations among motion sequences within the group, we propose an Asynchronous Motion Consistency (AMC) loss, ensuring the successful alignment and recovery of long-term occluded motion sequences under the guidance of unoccluded motion. To ensure the quality of group motion segments in C.G.M., we employed a transformer Variational

Autoencoder (VAE) and trained the model using occlusion data augmentation to develop a human motion prior, thereby improving the smoothness and authenticity of the motion.

Furthermore, to verify the proposed method, we create a synthetic large-scale dynamic crowd dataset *VirtualCrowd* by employing motion generation techniques [16] and datasets of virtual humans [17] and scenes [18]. This dataset provides eight 8K validation videos from four large-scale scenes, featuring varying crowd densities and perspectives (high-angle and low-angle), with detailed 2D/3D annotations, including 2D and 3D keypoints, global positions, and SMPL-X [19] parameters. *VirtualCrowd* is expected to facilitate future research on large-scene dynamic crowd reconstruction and analysis. Experimental results on both synthetic and real captured videos demonstrate that our method achieves spatio-temporally consistent 3D crowd reconstruction from a large-scene video, as shown in Fig. 1 with an example. The code and dataset will be publicly available for research purposes.

In summary, the main contributions of this paper are:

- This is the first framework for reconstructing the motion of crowds from a large-scene video, achieving spatial and temporal consistency, and continuous and natural reconstruction of 3D positions, poses and shapes of hundreds of people.
- We propose a coarse-to-fine group-guided motion optimization strategy to fit crowd motion, by introducing a VAE-based motion prior in optimization followed by a

segment-level group-guided optimization.

- We design an asynchronous motion consistency loss that measures the differences in motion segments within a group. This helps in recovering motion under long-term occlusions.
- We contribute a virtual dataset that includes four dynamic scenes with an average of hundreds of people, equipped with detailed 3D labels such as 3D joints, 3D positions, and SMPL-X parameters to evaluate large-scene dynamic crowd reconstruction.

II. RELATED WORK

A. Crowd Reconstruction in Large Scenes

In recent years, the field of imaging has made significant advancements, producing gigapixel-level images with a wide field of view and high resolution details. These images are typically captured by a multi-scale camera array [20], [21]. A notable dataset in this domain is PANDA [22] which is a human-centric large-scene video dataset, including high-resolution ($\sim 25,000 \times 15,000$) images of large-scale scenes with hundreds of people, and has been applied across a spectrum of visual tasks [23]–[27].

Particularly, building on the availability of large-scale scene data, Wen et al. [1] introduce the task of crowd reconstruction in a large scene, which aims to reconstruct 3D poses, shapes, and positions of hundreds of individuals within a unified camera space from a single large-scene image. They design a joint local and global inference framework called Crowd3D, which converts the complex crowd localization into pixel localization with the help of Human-scene Virtual Interaction Point (HVIP) and the parameters of pre-estimated scene-level camera and ground plane, directly reconstructing human body meshes in a large-scene camera system from cropped images. Moreover, GroupRec [2] solves the task by incorporating group features and positional information into network inference, accounting for collective behaviors and interactions to derive relatively accurate crowd positioning. However, these methods solely focus on a single-image reconstruction and as a result, they are unable to infer natural and smooth crowd motion from a large-scene video.

B. Video-based Single-person Reconstruction

Typically, reconstructing a single person from a video is a reasonable way to maintain the consistency of the pose and shape over time. Previous research [3], [4], [6], [28], [29] primarily focuses on exploring various temporal modeling architectures. VIBE [3] and TCMR [4] rely on recurrent neural networks (RNNs) for temporal information. VIBE builds a bidirectional gated recurrent unit (GRU) time encoder and enhances the plausibility of motion sequences through adversarial training. However, the reconstruction results are more heavily dependent on static features. To address it, TCMR aggregates temporal features of the past, future, and current frame to strongly constrain outputs with motion consistency. Compared to using an RNN, MPS-Net [28] captures motion continuity dependencies through attention mechanisms, enhancing temporal correlations. Due to the inability of attention

mechanisms to effectively capture local details, GLoT [6] utilizes transformer architectures to encode videos, capturing long-term temporal correlations while refining local details through attention on nearby frames. Recent approaches [7], [8], [30] focus on estimating the global human trajectory from a monocular dynamic camera. GLAMR [7] utilizes a transformer-based generative motion filler to fill missing motion from regression. It optimizes global trajectories between multiple frames of individuals in the scene but ignores the camera motion cues. Both WHAM [30] and TRAM [8] recover camera motion with SLAM (Simultaneous Localization and Mapping). WHAM recursively predicts the pose, shape, and global motion parameters without optimization for faster inference. However, it may have poor generalization for novel posture as it learns human motion from MoCap data. TRAM effectively captures more accurate human motion by composing camera trajectory in the world frame and a person's body motion in the camera frame, achieving better generalization to complex scenarios.

However, these methods struggle with large-scene videos featuring hundreds of people because they estimate the pose and shape of each person in local camera space, resulting in an incoherent spatial distribution of the reconstructed multi-person meshes. Additionally, they neglect interactions between multiple individuals and struggle to handle scenarios involving mutual occlusion.

C. Video-based Multi-person Reconstruction

Unlike methods that reconstruct humans on a person-by-person basis, some approaches focus on recovering multi-person meshes simultaneously. These methods demonstrate the ability to reconstruct globally consistent positions and scale relationships among individuals, offering greater robustness in handling mutual occlusions and preserving spatial coherence across the scene. Benefiting from the exploration of one-stage multi-person pose and shape estimation [31]–[33], some methods [9], [10] have achieved end-to-end prediction of multiple people in a video. Specifically, PSVT [10] employs a progressive video Transformer encoder to capture global feature dependencies among subjects, enabling end-to-end estimation of multiple people's poses and shapes. TRACE [9] utilizes optical flow as a motion cue and estimates multiple people by a novel 5D representation that represents the temporal motion of a person in both camera and global coordinates. Regrettably, these methods are unable to directly regress humans from an entire large-scene image because, compared to the image size, the scale of humans is relatively small and varies greatly.

Other methods [11], [12] utilize multi-stage spatio-temporal joint optimization, incorporating generative models [34], [35] as motion priors, to reconstruct spatially and temporally consistent multi-person poses and shapes from a video. PACE [12] integrates SLAM and human motion priors into an optimization framework, leveraging the strengths of both to jointly optimize the motions of humans and the camera. SLAHMR [11] decouples camera and human motion by integrating relative camera estimation and learned transition human motion priors to address the scene scale ambiguity, thereby inferring the

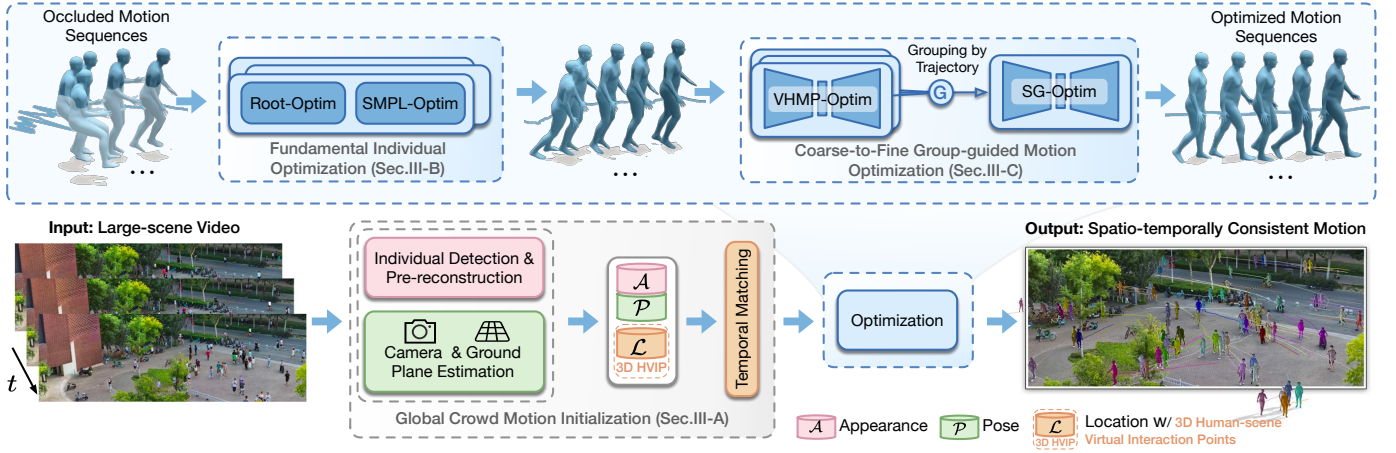


Fig. 2. Overview of DyCrowd framework. DyCrowd is a multi-stage optimization method that consists of three main components: global crowd motion initialization, fundamental individual optimization, and coarse-to-fine group-guided motion optimization. The coarse-to-fine group-guided motion optimization utilizes a VAE-based Human Motion Prior Optimization (VHMP-Optim) to enhance the temporal stability of crowd motion and employs a Segment-level Group-guided Optimization (SG-Optim) to address long-term occlusion issues in large-scene videos. This process results in a spatio-temporally consistent 3D reconstruction of crowd motion within a unified global space.

motion of multiple individuals in a shared coordinate system. Although multi-stage optimization processes may appear suitable for large-scene videos, these methods are typically tailored for small scenes captured by hand-held cameras and are not equipped to address the unique challenges of large scenes, such as significant scale variations, complex spatial extents, and frequent mutual occlusions.

In this paper, we introduce DyCrowd, the first framework for reconstructing spatio-temporally consistent 3D positions, poses, and shapes of hundreds of people in a unified global space from a large-scene video. DyCrowd produces natural crowd motion and addresses the problem of motion recovery under long-term occlusion in large-scene videos through our coarse-to-fine group-guided motion optimization strategy and asynchronous motion consistency loss. We also contribute a virtual dataset called *VirtualCrowd* to support the evaluation of large-scene dynamic crowd reconstruction.

III. METHODS

The proposed DyCrowd aims to achieve precise and continuous 3D reconstruction of the global positions, poses, and shapes of hundreds of people from a large-scene video. Fig. 2 illustrates the framework of our proposed approach. The core of our approach is a coarse-to-fine group-guided motion optimization strategy that addresses temporal instability through a human motion prior model and restores long-term occluded postures by leveraging the collective behavior of the crowd. DyCrowd employs a multi-stage optimization process to achieve spatial-temporal consistent global crowd dynamics reconstruction: 1) integrating estimations from data-driven models to obtain the global crowd motion initialization (Sec. III-A); 2) employing a fundamental individual optimization for reliable positions, poses and shapes of individuals (Sec. III-B); 3) addressing motion instability and dynamic occlusions by our coarse-to-fine group-guided motion optimization strategy (Sec. III-C), which progressively optimizes natural and occlusion-robust dynamic crowd motion.

We take as input a large-scene video with T frames of a scene with N individuals and represent their motion sequence in a global large-scene coordinate system with $\{Q_n\}_{n=1}^N$. We use SMPL [36] model and $Q_n = \{Q_{n,t} = \{\theta_{n,t}, \beta_{n,t}, \gamma_{n,t}, \tau_{n,t}\}\}_{t=1}^T$. For each individual, the SMPL model transforms body pose parameters $\theta_{n,t} \in \mathbb{R}^{23 \times 3}$, shape parameters $\beta_{n,t} \in \mathbb{R}^{10}$, rotation parameters $\gamma_{n,t} \in \mathbb{R}^3$ and translation $\tau_{n,t} \in \mathbb{R}^3$ into the human body mesh vertices $V_{n,t} \in \mathbb{R}^{6890 \times 3}$ and joint positions $J_{n,t}^{\text{SMPL}} \in \mathbb{R}^{24 \times 3}$ by applying a differentiable function \mathcal{M} :

$$[V_{n,t}, J_{n,t}^{\text{SMPL}}] = \mathcal{M}(\gamma_{n,t}, \theta_{n,t}, \beta_{n,t}) + \tau_{n,t}. \quad (1)$$

A. Global Crowd Motion Initialization

In large-scene crowd motion reconstruction, we begin by obtaining an initial estimate of the global crowd motion to facilitate subsequent optimization. This process involves individual detection and pre-reconstruction, camera and ground plane estimation, and temporal matching of individuals to generate preliminary global motion results, which serve as inputs for subsequent optimization.

In the process of individual detection and pre-reconstruction, extracting individual features presents significant challenges due to variations in pedestrian scale caused by differences in their distance from the camera, especially when individuals appear small. To address this issue, we first employ an adaptive human-centric cropping method [1], which divides each video frame into multiple regions. Within each region, we apply the top-down detection method, VitDet [37], to identify individuals. The detection results from all regions are then merged to accurately obtain the bounding boxes and masks of each individual. Based on these predicted bounding boxes, we extract the 2D keypoints $J_{n,t}^{2D} \in \mathbb{R}^{17 \times 2}$ for each individual using DWPose [38]. Additionally, we use the human mesh reconstruction method, HMR2.0 [39], to obtain the locally reconstructed SMPL

model's shape parameters $\bar{\beta}_{n,t} \in \mathbb{R}^{10}$ and pose parameters $\bar{\theta}_{n,t} \in \mathbb{R}^{23 \times 3}$ for each individual.

To estimate the ground plane $\bar{\mathbf{G}} : \{\bar{N}_g, \bar{D}_g\}$, and the camera's intrinsic parameters $\bar{\mathcal{K}}$, similar to Crowd3D [1], we use walking and standing human poses as priors for calibration. Here, \bar{N}_g represents the ground normal, and \bar{D}_g is a constant. Furthermore, we also infer the 2D and 3D Human-scene Virtual Interaction Points (HVIP) [1] for each individual by a single-person HVIP regression network and estimated camera parameters, which are denoted as: $\bar{\mathbf{p}}_{n,t}^{\text{hvip2D}} \in \mathbb{R}^2$ and $\bar{\mathbf{p}}_{n,t}^{\text{hvip3D}} \in \mathbb{R}^3$. HVIP represents the projection point of a person's 3D torso center on the ground plane, and we use it to locate the person's initial position in the global space.

We employ the PHALP [40] to achieve inter-frame individual matching. Unlike the original PHALP, we employ 3D HVIP as the position representation to overcome the limitations of the PHALP framework in large-scale scene tracking. Specifically, PHALP matches individuals frame by frame based on appearance (\mathcal{A}), pose (\mathcal{P}), and position information (\mathcal{L}), and compares these with the current frame's detection results to update trajectories, enabling continuous tracking. However, due to the significant scale differences between small and large scenes, the tracking performance of PHALP in large-scenes is suboptimal, and the position information provided by the framework underperforms in camera setups with a global perspective. In contrast, 3D HVIP accurately locates individuals on the ground, offering a more effective positional feature. Furthermore, we confine the matching process to spatially adjacent individuals based on the 3D HVIP positional representation, enhancing computational efficiency.

After the global crowd motion initialization, the motion sequence $\bar{\mathbf{Q}}_n$ for each individual is obtained and used as the initial input for the optimization process.

B. Fundamental Individual Optimization

As $\bar{\mathbf{Q}}_n$ is obtained from local reconstruction inference, its global position, including both rotation and translation, along with the pose of the SMPL model, remain ambiguous. To address this, the fundamental individual optimization process focuses on root optimization and SMPL optimization for each frame to ensure robustness and rationality in both position and pose. Below, we provide a detailed explanation of the optimization formulas and their corresponding objective functions.

The primary objective of root optimization is to obtain a reasonable positional representation by fitting the root joint of the human body. The first set of variables we optimize includes the global root translation $\tau_{n,t}$ and root rotation $\gamma_{n,t}$ for each individual. The objective function for this phase is as follows:

$$\begin{aligned} & \arg \min_{\tau_{n,t}, \gamma_{n,t} : \forall n, t} E_{\text{root}}, \quad \text{with} \\ E_{\text{root}} &= E_{2D} + E_{\text{hvip2D}} + E_{\text{contact}} + E_{\text{t-trans}}. \end{aligned} \quad (2)$$

The error term E_{2D} is used to ensure the reconstructed human motion is spatially reasonable and consistent with the image evidence. E_{2D} is measured by the 2D reprojection error

between the 3D motion estimation and the 2D body joint observations:

$$E_{2D} = \lambda_{2D} \sum_{n,t,j} \sigma_{n,t,j} \delta_{n,t} \left\| \Pi_{\mathcal{K}}(\tilde{\mathbf{J}}_{n,t,j}^{3D}) - \bar{\mathbf{J}}_{n,t,j}^{2D} \right\|_2^2, \quad (3)$$

where λ_{2D} represents the contribution weight of this term, and λ_* for the other terms mentioned later also indicates their respective contribution weights. $\sigma_{n,t,j} \in \{0, 1\}$ indicates whether the j -th keypoint of the n -th individual in the t -th frame is valid, $\tilde{\mathbf{J}}_{n,t,j}^{3D} \in \tilde{\mathbf{J}}_{n,t}^{\text{SMPL}}$ is the 3D keypoint coordinates to be optimized, $\delta_{n,t}$ is the inverse of scale (bounding box size) of the n -th individual in the t -th frame, and $\Pi_{\mathcal{K}}$ is the perspective projection with the camera intrinsic matrix \mathcal{K} . To address the limitation where the re-projection error based solely on 2D human keypoints fails to capture the interaction between the human body and the scene, we introduce a 2D HVIP loss term, E_{hvip2D} . This loss term optimizes the spatial position of the human body by penalizing unrealistic placements of keypoints in relation to the scene, thereby better describing the interaction between the human body and the environment. The specific formula is as follows:

$$E_{\text{hvip2D}} = \lambda_h \sum_{n,t} \zeta_{n,t} \delta_{n,t} \left\| \Pi_{\mathcal{K}}(\mathcal{H}(\tilde{\mathbf{J}}_{n,t}^{3D}, \bar{\mathbf{G}})) - \bar{\mathbf{p}}_{n,t}^{\text{hvip2D}} \right\|_2^2, \quad (4)$$

where $\zeta_{n,t} \in \{0, 1\}$ indicates whether the n -th individual in the t -th frame is valid. The function $\mathcal{H}(\cdot, \bar{\mathbf{G}})$ is employed to compute the 3D human virtual interaction point on the ground plane $\bar{\mathbf{G}}$. This loss function primarily focuses on the vertical projection point between the human body's center of mass and the ground plane, but it neglects the actual distance between the human body and the ground. To address this, we introduce E_{contact} to further determine a reasonable position by minimizing the distance between the lowest point of the mesh and the ground. The specific definition is as follows:

$$E_{\text{contact}} = \lambda_c \sum_{n,t} \zeta_{n,t} \left\| \min \left\{ \frac{|\bar{N}_g^T \tilde{\mathbf{V}}_{n,t}^{(i)} + \bar{D}_g|}{\|\bar{N}_g\|_2}, \forall i \right\} \right\|_1, \quad (5)$$

where $\|\cdot\|_1$ denotes the $L1$ norm. $\tilde{\mathbf{V}}_{n,t}^{(i)}$ denotes the i -th vertex point of the mesh vertices of person n at time t . We also use a penalty term $E_{\text{t-trans}}$ to penalize its movement in the time domain (Δ_t), thereby better fitting the root translation $\tilde{\tau}_{n,t} \in \mathbb{R}^3$. The specific definition is as follows:

$$E_{\text{t-trans}} = \lambda_{\text{t-trans}} \sum_{n,t} \zeta_{n,t} \|\Delta_t(\tilde{\tau}_{n,t})\|_2^2, \quad (6)$$

where $\Delta_t(\tau_{n,t}) = \tau_{n,t} - \tau_{n,t-1}$. This term allows us to obtain temporally more stable root translation with significantly less jitter.

The SMPL optimization obtains a plausible pose of human body by fitting the parameters of SMPL model. In addition to the root translation $\tau_{n,t}$ and root rotation $\gamma_{n,t}$ variables, we also optimize the shape parameters $\beta_{n,t}$ and pose parameters $\theta_{n,t}$. Due to the inherent ambiguity in mapping 2D features to 3D poses, similar to [11], [19], we employ the VPoser prior [41] to avoid generating impossible human poses. We

optimize the latent variable $z_{n,t}^\varphi$, which is obtained by transforming pose parameters through the VPoser encoder, and use its decoder to generate new pose parameters. The objective function for this stage is defined as:

$$\begin{aligned} \arg \min_{\tau_{n,t}, \gamma_{n,t}, \beta_{n,t}, z_{n,t}^\varphi : \forall n, t} E_{\text{smpl}}, \quad \text{with} \\ E_{\text{smpl}} = E_{\text{root}} + E_{\text{pose}} + E_{\text{shape}} + E_{\text{t-pose}}. \end{aligned} \quad (7)$$

The pose regularization term E_{pose} computes the loss of the latent variable based on the prior assumption of a Gaussian distribution, serving as a penalty term:

$$E_{\text{pose}} = -\lambda_\varphi \sum_{n,t} \log \mathcal{N}(\tilde{z}_{n,t}^\varphi; 0, I), \quad (8)$$

where $\tilde{z}_{n,t}^\varphi \in \mathbb{R}^{32}$ is the latent variable to be optimized. The shape regularization term E_{shape} penalizes overly large shape parameters $\tilde{\beta}_{n,t}$ to maintain the model's shape integrity and prevent the generation of unnatural shapes, which is defined as:

$$E_{\text{shape}} = \lambda_{\text{shape}} \sum_{n,t} \zeta_{n,t} \left\| \tilde{\beta}_{n,t} \right\|_2^2. \quad (9)$$

To minimize pose jitter in the time series and achieve smoother motion, we introduce a temporal pose penalty:

$$E_{\text{t-pose}} = \lambda_{\text{t-pose}} \sum_{n,t} \zeta_{n,t} \left\| \Delta_t(\tilde{\mathbf{J}}_{n,t}^{3D}) \right\|_2^2. \quad (10)$$

This term captures the essence of maintaining stability across temporal poses by penalizing excessive changes in the 3D joint configurations.

In the fundamental individual optimization process, we implement a two-stage optimization strategy to progressively enhance the accuracy of crowd motion in terms of position and posture. This involves applying various position-related constraints, such as 2D HVIP constraints, contact constraints, and 2D pose loss of the human, thereby providing reasonable initial values for subsequent optimization steps.

C. Coarse-to-fine Group-guided Motion Optimization

Our goal is to enhance the temporal stability of crowd motion and to recover poses that have been occluded for an extended period. To achieve this, we design a coarse-to-fine group-guided motion optimization strategy. This strategy initiates the optimization process with a Variational Autoencoder (VAE)-based human motion prior and continues with a segment-level group-guided optimization, progressively reconstructing the motion of the crowd.

Optimization with VAE-based Human Motion Prior. To generate smooth and realistic human motion while addressing short-term occlusions, we integrate prior knowledge of human motion by developing a transformer-based VAE [42] as our motion prior model, denoted as $\mathcal{V} = \mathcal{E}_p, \mathcal{E}_r, \mathcal{D}$. Unlike transition-based motion priors [34], our approach enables faster inference, which is crucial for handling large groups of people in extensive scenes. We decompose human motion into local and global components, where local motion refers to body posture features, and global motion captures the root

trajectory. We segment an individual's entire motion sequence into multiple segments, with each segment delineated as a distinct motion segment, each lasting 64 frames at a frame rate of 30 frames per second (so approximately 2 seconds). For an individual n in the s -th temporal segment, the local transformer encoder \mathcal{E}_p and the global MLP encoder \mathcal{E}_r generate the combined latent motion representation $z_{n,s}^\vartheta = (z_{n,s}^p, z_{n,s}^r)$, where $z_{n,s}^\vartheta \in \mathbb{R}^{256+128}$. These representations serve as optimization variables in the coarse-to-fine group-guided motion optimization, allowing the decoder \mathcal{D} to effectively recover motion sequences. The state of a moving person is represented as $\mathbf{x} = (\mathbf{x}_p, \mathbf{x}_r)$, $\mathbf{x} \in \mathbb{R}^{330}$ at time t . The local motion state $\mathbf{x}_p = (\hat{\mathbf{J}}, \hat{\mathbf{v}}, \hat{\boldsymbol{\theta}})$ is composed of joint positions $\hat{\mathbf{J}} \in \mathbb{R}^{21 \times 3}$, joints velocity $\hat{\mathbf{v}} \in \mathbb{R}^{21 \times 3}$ and joint rotations $\hat{\boldsymbol{\theta}} \in \mathbb{R}^{21 \times 9}$, where the joints are in local coordinate system that the local joints are relative to the current frame pelvis joint and retains the position coordinates of the ground direction. The global motion state $\mathbf{x}_r = (\hat{\mathbf{r}}, \hat{\boldsymbol{\tau}}, \hat{\boldsymbol{\gamma}})$ includes root positions $\hat{\mathbf{r}} \in \mathbb{R}^3$, the global translation $\hat{\boldsymbol{\tau}} \in \mathbb{R}^3$ and rotation $\hat{\boldsymbol{\gamma}} \in \mathbb{R}^9$ of SMPL [36]. An additional output is the contact probability c , which represents the likelihood of each joint being in contact with the ground. The contacts enable ground constraints in optimization, promoting the coherence of interaction with the ground. Consistent with prior approaches [43]–[46], we employ KL divergence, and L2 loss of each state to achieve training from AMASS dataset [47]. To enhance the occlusion completion ability of our prior model to the occlusion commonly existing in large scenes, we define the visibility masks $\mathbf{m}_p \in \{0, 1\}^{330}$ of input motion state. During the training phase, we synthesize various occlusion masks on the AMASS dataset, including random joint occlusion and frame occlusion, partial body occlusion, and random consecutive occlusion for global or local input states. In the optimization phase, we set the masks based on the confidence of the detected 2D joints.

We apply our motion prior model to optimization to obtain reasonable motion trajectories and postures. The optimization variables consist of the SMPL shape parameters for each segment of the motion sequence, denoted as $\beta_{n,s}$, and the latent variables $z_{n,s}^\vartheta$. The objective function for this stage is defined as:

$$\begin{aligned} \arg \min_{\beta_{n,s}, z_{n,s}^\vartheta : \forall n, s} E_{\text{motion}}, \quad \text{with} \\ E_{\text{motion}} = E_{\text{smpl}} + E_{\text{VAE}} + E_{\text{env}} + E_{\text{connect}}. \end{aligned} \quad (11)$$

In order to make the motion reasonable, E_{VAE} optimizes the latent variable based on the prior assumption of the Gaussian distribution as a penalty term:

$$E_{\text{VAE}} = -\lambda_{\text{vae}} \sum_{n,s} \log \mathcal{N}(\tilde{z}_{n,s}^\vartheta; 0, I), \quad (12)$$

where $\tilde{z}_{n,s}^\vartheta$ represents the latent variables to be optimized. The environmental energy E_{env} ensures that each individual maintains a specified height range from the ground upon contact, with a velocity of zero at the moment of contact.

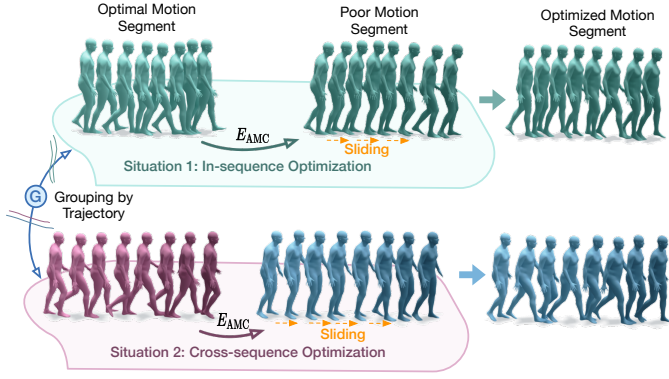


Fig. 3. Segment-level group-guided optimization. Motion segments are grouped based on similar relative trajectories, with the optimal motion segments in each group assisting in the optimization of heavily occluded segments. It includes both in-sequence and cross-sequence optimization situations, depending on whether the optimal motion sequence belongs to the same individual.

This is expressed as:

$$E_{\text{env}} = \sum_{n, t_s, j} \lambda_{\text{ch}} c_{n, t_s, j} \max(0, |\tilde{\mathbf{J}}_{n, t_s, j}^{3D, z} - h|) + \lambda_{\text{cv}} c_{n, t_s, j} \left\| \Delta_t(\tilde{\mathbf{J}}_{n, t_s, j}^{3D}) \right\|_2^2, \quad (13)$$

where t_s is the time frame of the s -th segment, $c_{n, t_s, j}$ is the contact probability of the j -th joint of the n -th individual in the t_s -th frame of the model output. The contact height term consists of the z -coordinate component of the contacting joint being no more than h above the ground in the standard frame. E_{connect} serves as a penalty term to ensure motion consistency between the end of one segment and the beginning of the adjacent segment within the same sequence:

$$E_{\text{connect}} = \lambda_{\text{con}} \sum_{n, s} \zeta_{n, s_l} \zeta_{n, s_l^+}^+ \left\| \tilde{\mathbf{J}}_{n, s_l^+}^{3D} - \tilde{\mathbf{J}}_{n, s_l}^{3D} \right\|_2^2, \quad (14)$$

where s_l is the last frame of the current segment, and s_l^+ is the first frame of the next segment. This constraint term ensures continuous motion from segment to segment.

Segment-level Group-guided Optimization. The motion prior model helps obtain temporally stable human motion and can recover from short-term and sporadic occlusions, but struggles to handle severe long-term occlusions. To address this limitation, we propose a segment-level group-guided optimization method that enhances the motion recovery of long-term occluded sequences by leveraging the collectivity of crowd behavior. As shown in Fig. 3, unlike the GroupRec [2] method, which groups based on collective behavior in static images, we first group the segments according to similar relative motion trajectories. Following this, we jointly optimize their motion sequences using our Asynchronous Motion Consistency (AMC) loss (see below). Within each group, the unoccluded segments are designated as the optimal motion segments, serving to guide those with poor motion segments. Based on whether the guiding and guided motion segments belong to the same individual, we can categorize the scenarios into two situations: in-sequence and cross-sequence. In the in-

sequence scenario, the sequence containing occlusions is recovered through guidance from its own unobstructed segments.

During the grouping process, the symmetric segment path distance [48] is adopted as the distance metric between motion trajectories. This approach allows us to focus more on the overall shape of the trajectories while accommodating variations in trajectory lengths. For clustering, we employ the affinity propagation clustering [49] algorithm, which eliminates the need to predefine the number of clusters.

To achieve the guidance optimization between motion segments within a group, we design an asynchronous motion consistency loss to quantify the differences in intra-group motion. It includes a motion segment confidence to identify the optimal and poor segments within a group. We calculate the motion segment confidence $c_{n, s}^{\text{seg}}$ for each motion segment in a group based on detection scores and joint importance:

$$c_{n, s}^{\text{seg}} = v_{n, s} \cdot \left(1 - \mathbb{1} \left(\sum_j w_j \mathcal{F}_j(\tilde{\mathbf{J}}_{n, s}^{2D}) > a \right) \right), \quad (15)$$

where $v_{n, s}$ represents the proportion of visible joints (those with high detection scores) in the motion segment, w_j denotes the predefined joint weight based on the importance of joint occlusion, and $\mathcal{F}_j(\cdot)$ is a decision function used to determine whether the j -th joint is continuously occluded. The parameter a is a threshold that signifies the need for repair. Based on the scores of motion segment confidence, we adaptively determine the segments that require repair and identify high-quality segments in a group:

$$\xi_{n, s} = \begin{cases} 1 & \text{if } c_{n, s}^{\text{seg}} = \max(\{c_{n, s}^{\text{seg}}\}_{s=1}^{S_g}) \\ 0 & \text{if } c_{n, s}^{\text{seg}} = 0 \\ -1 & \text{otherwise} \end{cases}, \quad (16)$$

where S_g is the number of segments within the group. When the value of $\xi_{n, s}$ is 1, it indicates the optimal motion segment; a value of 0 denotes a poor motion segment; and -1 signifies a segment that does not require optimization. The collective nature of intra-group motion is utilized for motion recovery, using unobstructed segments (i.e., those where $\xi_{n, s} = 1$) to restore the motion of severely obscured segments (i.e., those where $\xi_{n, s} = 0$). If there exists more than one $\xi_{n, s} = 1$, we prioritize using the optimal segment in the in-sequence scenario.

While the motions within a group appear similar in general, they often exhibit temporal misalignments and rhythmic variations in their posture sequences. Innovatively, we treat each individual's motion sequence as a time series and propose our asynchronous motion consistency loss E_{AMC} as follows:

$$E_{\text{AMC}} = \lambda_{\text{AMC}} \sum_{n, s} w \mathcal{L}_s(\tilde{\boldsymbol{\theta}}_{n, s}, \tilde{\boldsymbol{\theta}}_{n', s'}) \cdot \mathbb{1}(\xi_{n, s} = 0, \xi_{n', s'} = 1). \quad (17)$$

Here, $\tilde{\boldsymbol{\theta}}_{n, s}$ and $\tilde{\boldsymbol{\theta}}_{n', s'}$ represent the pose parameters of the poor segment and the optimal segment, respectively. The term w denotes the predefined joint weights. To preserve the flexibility of individual motions, we exclude joints associated with expressive details, such as those in the head and hands. The




Fig. 4. Eight validation videos in the *VirtualCrowd* dataset and the corresponding crowd distributions using the SMPL-X model. The first row is the high-angle view, and the second row is the low-angle view.

function \mathcal{L}_s corresponds to the Soft Dynamic Time Warping (sDTW) [50] method, a differentiable variant of Dynamic Time Warping (DTW) [51]. DTW is a dynamic programming algorithm designed to measure the similarity between two time series through nonlinear alignment, enabling it to handle sequences with differing speeds or lengths. In our approach, we employ sDTW to temporally stretch and align the joint rotations of motion segments within the group. The proposed asynchronous motion consistency loss measures the variance in motion within a group, facilitating the transition from occluded poses to high-quality poses, which is beneficial for motion recovery during periods of long-term occlusion.

We set $\beta_{n,s}$, and $z_{n,s}^\theta$ as optimization variables, and the objective function for optimization is:

$$\begin{aligned} \arg \min_{\beta_{n,s}, z_{n,s}^\theta} : \forall n, s \quad E_{\text{group}}, \quad \text{with} \\ E_{\text{group}} = E_{\text{motion}} + E_{\text{AMC}}. \end{aligned} \quad (18)$$

In our coarse-to-fine group-guided motion optimization process, we first leverage prior knowledge of human motion to address the issue of instability in human motion. Subsequently, we employ a segment-level group-guided optimization with asynchronous motion consistency loss to tackle the problem of pose loss due to long-term occlusions, progressively reconstructing natural and complete human motion. As a result, the optimized individuals, featuring spatio-temporally consistent locations and accurate human poses, constitute the reconstruction output of the DyCrowd framework.

IV. EXPERIMENTS

A. Datasets and Evaluation Metrics

Virtual Large-scene Dataset. To validate our proposed method, we create a comprehensive large-scale dynamic crowd dataset named *VirtualCrowd*. This dataset utilizes Blender’s ICity plugin [18] to construct diverse scenes and extracts ground plane information to generate human motion trajectories. By employing SynBody [52] human models and the DIMOS [17] algorithm, we generate a variety of motion sequences. Subsequently, we render the dataset using Blender, resulting in a large-scale crowd motion video dataset, as shown

in Fig. 4. The dataset consists of four scenes, each covering an area larger than 2500 square meters. For each scene, two configurations are provided, incorporating both high-angle and low-angle perspectives, with varying numbers of people in different configurations. The high-angle perspective offers a broader field of view, while the low-angle perspective has more severe occlusions and greater scale differences between individuals. The dataset ultimately includes eight validation videos, each with a resolution of 7680×4320 (8K) and a frame rate of 30 frames per second, totaling 1600 frames. The scenes exhibit varying crowd densities, with the number of people ranging from 60 to 200. In total, there are 931 motion sequences and 186,200 poses. Furthermore, the dataset is equipped with detailed 2D and 3D annotations, including 2D joints, MOT annotations, 3D joints, 3D positions, and SMPL-X parameters, for evaluating dynamic crowd reconstruction in large scenes. *VirtualCrowd* provides essential foundational data support for crowd reconstruction, behavior analysis, and related research.

Furthermore, to further enrich the dataset and increase the diversity of scenes, we have extended the dataset with two videos of a sloped scene featuring different perspectives and crowd distributions, based on simulation techniques from [53], thus providing support for subsequent research. For more details, please refer to the supplementary material.

Real Large-scene Dataset. PANDA [22] is a gigapixel-level human-centric video dataset captured by a gigapixel camera and covering real-world scenes. As PANDA does not contain the poses of persons, we conduct qualitative comparisons on the PANDA dataset.

Evaluation Metrics. To verify the consistency of the reconstructed crowd in global space, we follow [1] and use the Procrustes-aligned pair-wise percentual distance similarity (PA-PPDS) to evaluate the location distribution of the crowd based on the distances among people. We also use the percentage of correct ordinal depth (PCOD) [54] to measure the correctness of ordinal depth relation. To evaluate 3D pose estimation performance, we employ mean per joint position error (MPJPE) and Procrustes-aligned MPJPE (PA-MPJPE). We focus on the accuracy of pose sequences in the global space and report two metrics: WA-MPJPE and W-MPJPE.

TABLE I
QUANTITATIVE COMPARISON ON VIRTUALCROWD DATASET.

Method	PA-PPDS \uparrow	PCOD \uparrow	MPJPE \downarrow	PA-MPJPE \downarrow	WA-MPJPE \downarrow	W-MPJPE \downarrow	ACCEL \downarrow
Crowd3D [1]	87.98	91.11	122.99	73.70	-	-	-
GroupRec [2]	75.04	86.22	89.04	58.98	82.20	94.34	165.50
DyCrowd	89.10	92.20	69.74	48.57	68.99	83.39	15.72
SLAHMR-Large* [11]	87.95	90.34	106.35	69.41	90.10	108.40	12.25
DyCrowd*	91.23	95.38	68.81	45.34	65.91	80.34	15.53

The symbol \star denotes the use of ground-truth object tracking, while “-” means unavailable results. SLAHMR-Large is a variant of SLAHMR modified for large scenes, and we compare it in the ground-truth object tracking setting.

TABLE II
ABLATION STUDY RESULTS ON VIRTUALCROWD DATASET.

Method	PA-PPDS \uparrow	PCOD \uparrow	MPJPE \downarrow	PA-MPJPE \downarrow	WA-MPJPE \downarrow	W-MPJPE \downarrow	ACCEL \downarrow
DyCrowd w/o C.G.M.	89.15	92.18	86.87	61.27	80.23	88.57	73.33
DyCrowd w/o AMC	89.12	92.19	70.22	49.02	69.41	83.35	15.08
DyCrowd	89.10	92.20	69.74	48.57	68.99	83.39	15.72

TABLE III
ABLATION STUDY RESULTS FOR OCCLUDED INSTANCES IN THE VIRTUALCROWD DATASET.

Method	MPJPE \downarrow	PA-MPJPE \downarrow
DyCrowd w/o C.G.M.	73.01	60.20
DyCrowd w/o AMC	67.17	56.77
DyCrowd	66.09	56.21

These metrics [11] align each output sequence with the ground-truth data using the entire sequence and the first two frames, respectively. To measure the motion smoothness, we compute acceleration error (ACCEL) against the ground-truth acceleration.

B. Implementation Details

We execute experiments using PyTorch [55], employing the RMSprop [56] optimizer with a learning rate 0.01. We progressively carry out root optimization, SMPL optimization, optimization with a VAE-based human motion prior, and segment-level group-guided optimization, iterating 100, 150, 200, and 200 times respectively. The weights λ used in the experiments are empirically defined to balance the magnitude of individual energy terms and remain constant throughout all experiments. We use $\lambda_{2D} = 100$, $\lambda_h = 10000$, $\lambda_c = 50$, $\lambda_{t-trans} = 50$, $\lambda_\varphi = 0.5$, $\lambda_{shape} = 1$, $\lambda_{t-pose} = 10$, $\lambda_{vae} = 0.2$, $\lambda_{ch} = 500$, $\lambda_{cv} = 1000$, $\lambda_{con} = 100$, $\lambda_{AMC} = 0.03$. Optimizing a large-scene video with 100 people and 200 frames takes approximately 4 hours on a workstation equipped with an NVIDIA 3090 GPU and 128G memory.

C. Comparison

Currently available methods are not designed for spatio-temporally consistent reconstruction of crowd motion in large-scene videos. Thus, we compare our approach with two state-of-the-art image-based large-scene methods: Crowd3D [1]

and GroupRec [2]. Additionally, we modify the state-of-the-art video-based method SLAHMR [11], originally designed for multi-person reconstruction in small scenes, to extend its capabilities to large scenes, dubbing it SLAHMR-Large. For fair comparison, we provide the scene-level camera focal length and ground plane parameters estimated by our method to them. For Crowd3D and GroupRec, we process each image in the large-scene videos separately. To ensure consistency of the human sequence during evaluation, we assign the results of our object tracking to GroupRec, replacing its original detection input. Crowd3D does not participate in sequence-level evaluation because it lacks effective temporal matching between multiple targets across frames. We modify SLAHMR by replacing its initialization with our global crowd motion initialization and using camera extrinsic parameters derived from our estimated ground plane. We perform the optimization of individuals in the scene in batches, due to the large scene and the high video memory requirements of SLAHMR. Since SLAHMR exhibits low tolerance for data noise, its optimization process may be disrupted when dealing with sequences that have poor detections or tracking. Additionally, due to the high computational cost—requiring approximately 8 hours to process every 25 people—we restrict our evaluations to scenarios with ground-truth object tracking.

As demonstrated in Table I, our method outperforms existing methods across all evaluation metrics except for ACCEL. Our method achieves the best PA-PPDS and PCOD scores, indicating a superior spatial arrangement. Moreover, benefiting from our coarse-to-fine group-guided motion optimization strategy, our method substantially surpasses existing methods in terms of pose estimation accuracy. The small ACCEL value obtained by our method suggests smooth and coherent human motion. We observe that SLAHMR-Large yields even smoother predictions because it utilizes the motion prior HuMoR [34], which favors smooth transitions between consecutive frames. However, when processing long or occluded sequences, it often generates nonsensical sliding rather than accurate pose predictions. Fig. 5 and Fig. 6 show

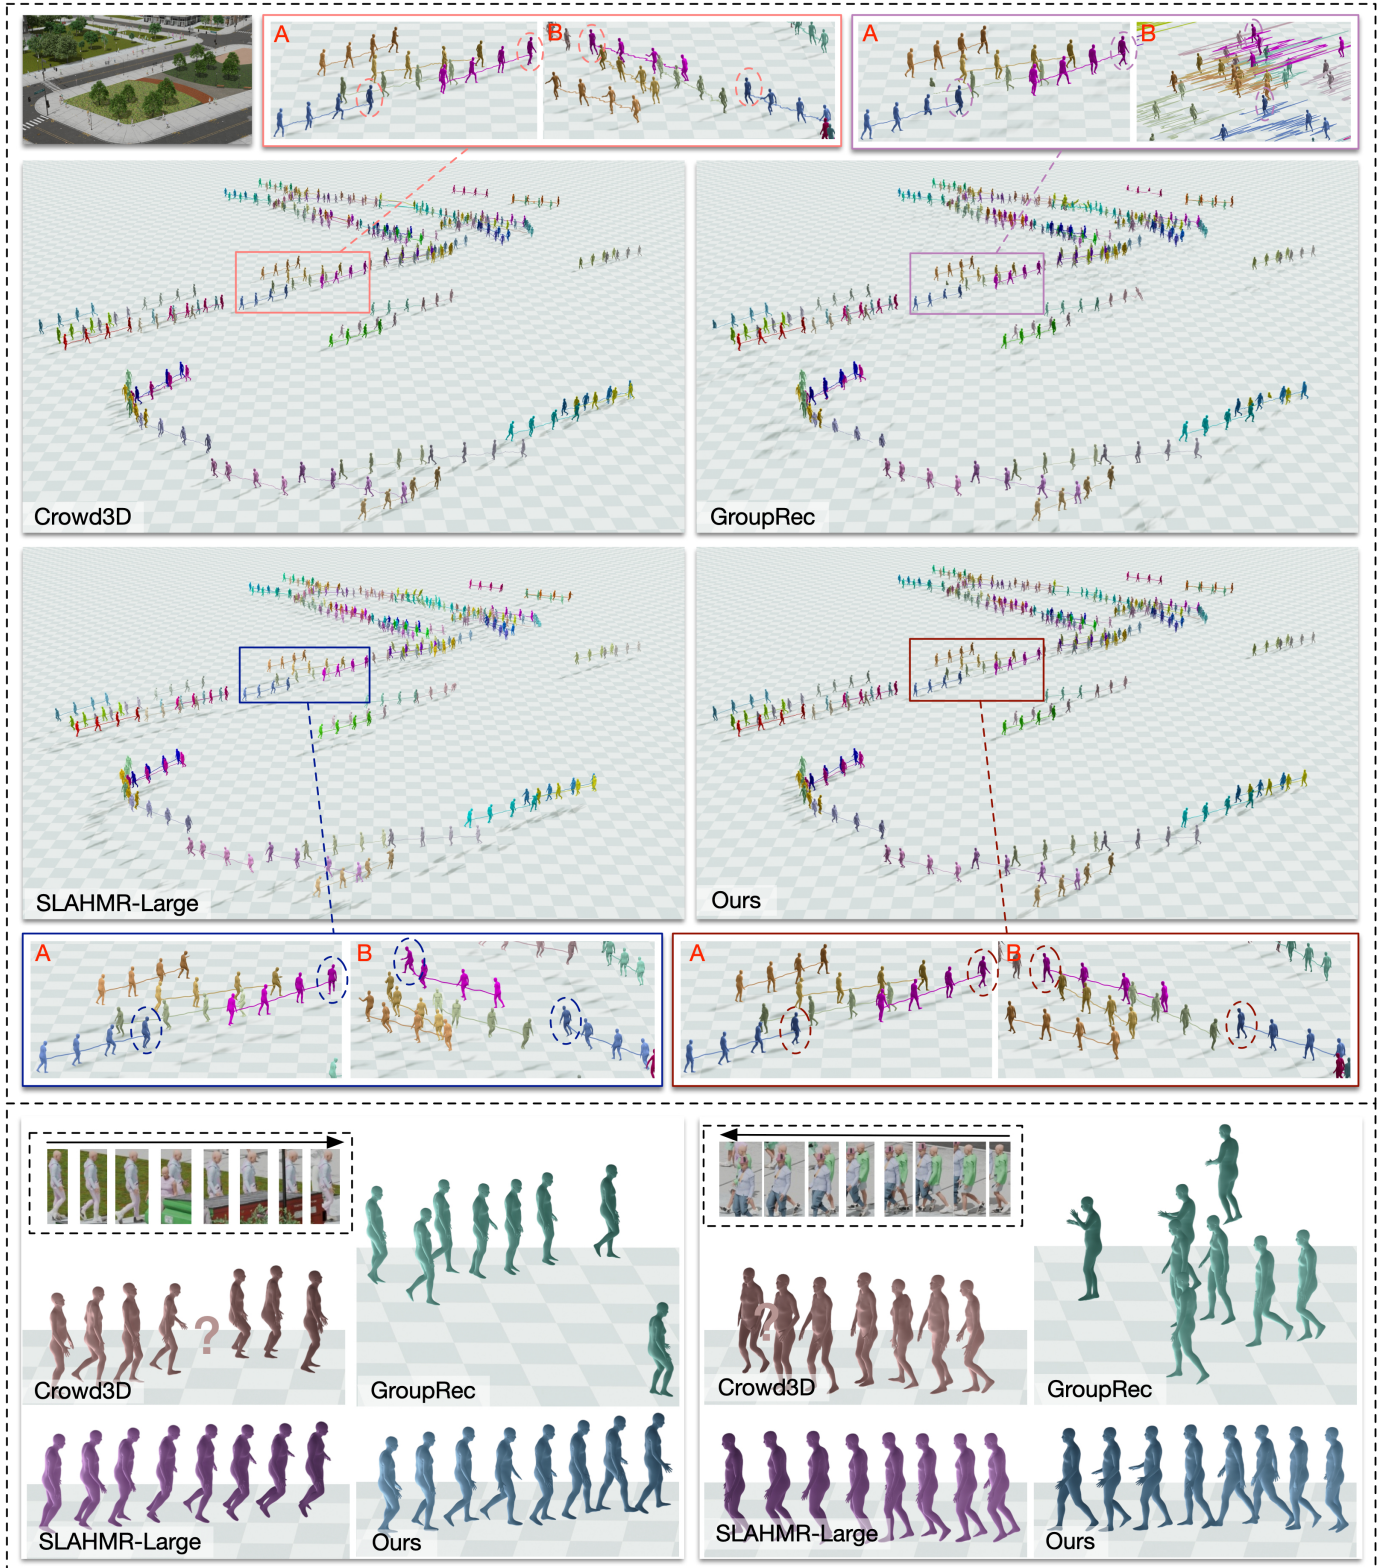


Fig. 5. Qualitative comparison on the *VirtualCrowd* dataset. The highlighted areas, including the camera perspective (A) and side perspective (B), demonstrate that our method achieves spatio-temporally consistent crowd reconstruction with plausible motion. Individual instance images demonstrate that our method is capable of restoring human motion even under severe occlusions.

the qualitative comparison on *VirtualCrowd* and *PANDA* [22] datasets. Our method achieves spatio-temporal consistency in the reconstruction of crowd motion, with reconstructed

individuals exhibiting more natural and plausible postures. Particularly, by leveraging the collective behavior of crowds in large scenes, our method is the only one capable of accurately

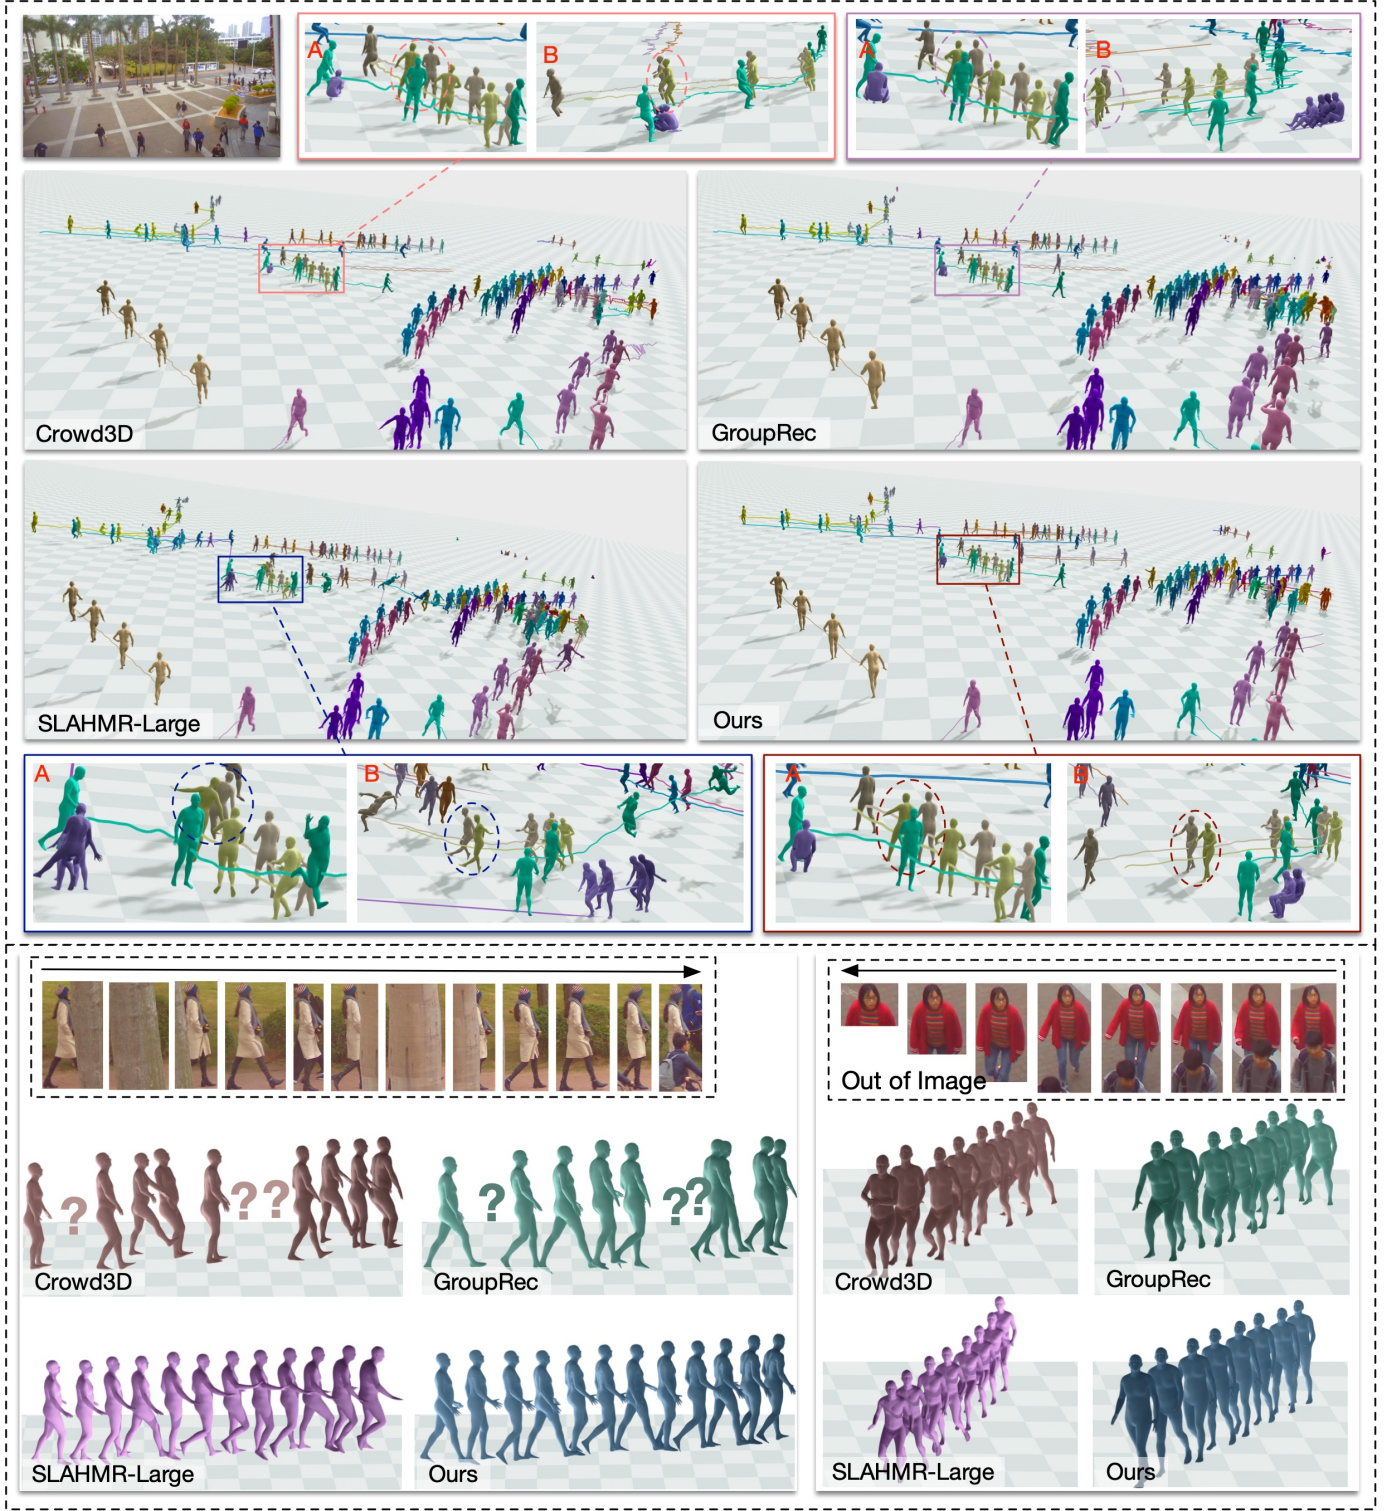


Fig. 6. Qualitative comparison on the *PANDA* dataset.

restoring human motion in instances of long-term occlusion.

D. Ablation Study

We conduct ablation studies to investigate the effects of key designed elements. For ease of reference, we abbreviate DyCrowd without Coarse-to-fine Group-guided Motion optimization to DyCrowd w/o C.G.M. and DyCrowd without

Asynchronous Motion Consistency loss to DyCrowd w/o AMC.

Coarse-to-fine Group-guided Motion Optimization. By comparing DyCrowd w/o C.G.M. and DyCrowd without AMC in Table II, it can be seen that our human motion prior model yields more accurate pose estimations. Since the AMC loss mainly focuses on dynamic occlusion recovery and local

TABLE IV
ABLATION EXPERIMENTS ON DIFFERENT VAE-BASED MOTION PRIORS IN THE VIRTUALCROWD DATASET.

Method	PA-PPDS \uparrow	PCOD \uparrow	MPJPE \downarrow	PA-MPJPE \downarrow	WA-MPJPE \downarrow	W-MPJPE \downarrow	ACCEL \downarrow
w/ NeMF [35]	88.72	91.86	84.63	59.43	80.64	99.44	13.87
w/ DMMR-VAE [57]	89.09	92.16	84.13	58.71	78.14	87.08	157.69
Ours	89.12	92.19	70.22	49.02	69.41	83.35	15.08

sequence consistency, it may cause slight fluctuations in global metrics such as PA-PPDS and W-MPJPE, and the introduction of more realistic but not completely smooth motion details leads to a slight increase in ACCEL. We further present the pose metrics for severely occluded instances in Table III, which demonstrate that our method achieves the best performance under severe occlusion, significantly improving robustness to occlusion.

Asynchronous Motion Consistency Loss. Our asynchronous motion consistency loss enables high-quality un-occluded motion segments to guide the motion recovery of occluded ones within groups that exhibit similar motion trajectories. This approach ensures robust and plausible motion recovery even under severe occlusion. Although the results in Table II indicate a slight compromise in performance compared to DyCrowd w/o AMC, Fig. 7 demonstrates that the motion under long-term occlusion is effectively restored by AMC loss, which we consider more meaningful. Additionally, Table III showcases that AMC loss further enhances the robustness against occlusion.

VAE-based Motion Prior. To validate the effectiveness of the VAE, we compared different VAE architectures. Specifically, we chose NeMF [35] and DMMR-VAE [57] as baselines, as these methods have recently been shown to be effective in human motion optimization tasks [12], [57]. As shown in Table IV, our VAE exhibits a significant advantage in pose reconstruction accuracy. These results stem from our key design for crowd motion optimization in large-scale scenes. Although the ACCEL index is slightly higher than that of some methods in certain cases, it mainly reflects the smoothness of motion. However, our method emphasizes retaining the real dynamic motion in complex large-scale scenes to obtain more motion details.

V. CONCLUSION AND DISCUSSION

Conclusion. We propose DyCrowd, the first framework to reconstruct spatio-temporally consistent 3D positions, poses, and shapes of crowd in a unified global space from a large-scene video. Our method incorporates prior knowledge of human motion and leverages the collective behavior of crowds in large scenes to facilitate the optimization of occluded human motion guided by visible human motion. This enables our method to not only reconstruct natural and continuous human motion but also achieve recovery of motion under long-term occlusion. We also contribute a virtual large-scene evaluation dataset named *VirtualCrowd* to facilitate future research on dynamic crowd reconstruction and analysis in large scenes.

Limitations and Future Work. Although our method achieves the reconstruction of crowd motion in large-scene




Fig. 7. Qualitative comparison of our key designs for occlusion.

videos, there are still some limitations. Firstly, our method may be severely affected in cases of catastrophic failures in 2D key-point detection and tracking. For instance, when a tracklet of a person is interrupted due to occlusion, we cannot recover the occluded individual's motion. Secondly, although our method can handle multi-ground scenes by region-level processing, it falls short in accurately capturing human interactions with complicated ground planes, such as climbing stairs. Lastly, our method employs a multi-stage optimization framework, which does not currently support real-time applications. In the future, we plan to integrate scene and environmental information into our framework. This will advance our ability to recover human motion in complex scene interactions.

ACKNOWLEDGMENTS

This work was supported in part by National Key R&D Program of China (2023YFC3082100), National Natural Science Foundation of China (62122058, 62171317), and Science Fund for Distinguished Young Scholars of Tianjin (No. 22JCJQC00040).

REFERENCES

- [1] H. Wen, J. Huang, H. Cui, H. Lin, Y.-K. Lai, L. Fang, and K. Li, "Crowd3D: Towards hundreds of people reconstruction from a single image," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2023, pp. 8937–8946.
- [2] B. Huang, J. Ju, Z. Li, and Y. Wang, "Reconstructing groups of people with hypergraph relational reasoning," in *Proc. IEEE International Conference on Computer Vision*, 2023, pp. 14 873–14 883.
- [3] M. Kocabas, N. Athanasiou, and M. J. Black, "VIBE: Video inference for human body pose and shape estimation," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 5253–5263.
- [4] H. Choi, G. Moon, J. Y. Chang, and K. M. Lee, "Beyond static features for temporally consistent 3D human pose and shape from a video," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1964–1973.
- [5] W.-L. Wei, J.-C. Lin, T.-L. Liu, and H.-Y. M. Liao, "Capturing humans in motion: Temporal-attentive 3D human pose and shape estimation from monocular video," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 13 211–13 220.

- [6] X. Shen, Z. Yang, X. Wang, J. Ma, C. Zhou, and Y. Yang, "Global-to-local modeling for video-based 3D human pose and shape estimation," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2023, pp. 8887–8896.
- [7] Y. Yuan, U. Iqbal, P. Molchanov, K. Kitani, and J. Kautz, "GLAMR: Global occlusion-aware human mesh recovery with dynamic cameras," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 11 038–11 049.
- [8] Y. Wang, Z. Wang, L. Liu, and K. Daniilidis, "TRAM: Global trajectory and motion of 3D humans from in-the-wild videos," *arXiv preprint arXiv:2403.17346*, 2024.
- [9] Y. Sun, Q. Bao, W. Liu, T. Mei, and M. J. Black, "TRACE: 5D temporal regression of avatars with dynamic cameras in 3D environments," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2023, pp. 8856–8866.
- [10] Z. Qiu, Q. Yang, J. Wang, H. Feng, J. Han, E. Ding, C. Xu, D. Fu, and J. Wang, "PSVT: End-to-end multi-person 3D pose and shape estimation with progressive video transformers," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 254–21 263.
- [11] V. Ye, G. Pavlakos, J. Malik, and A. Kanazawa, "Decoupling human and camera motion from videos in the wild," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 222–21 232.
- [12] M. Kocabas, Y. Yuan, P. Molchanov, Y. Guo, M. J. Black, O. Hilliges, J. Kautz, and U. Iqbal, "PACE: Human and camera motion estimation from in-the-wild videos," in *Proc. IEEE Int. Conf. 3D Vis.* IEEE, 2024, pp. 397–408.
- [13] B. Zhou, X. Tang, and X. Wang, "Measuring crowd collectiveness," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 3049–3056.
- [14] Y. Wu, Y. Ye, C. Zhao, and Z. Shi, "Collective density clustering for coherent motion detection," *IEEE Trans. Multimedia*, vol. 20, no. 6, pp. 1418–1431, 2017.
- [15] L. Mei, J. Lai, Z. Chen, and X. Xie, "Measuring crowd collectiveness via global motion correlation," in *Proc. IEEE International Conference on Computer Vision*, 2019, pp. 0–0.
- [16] Y. Zhang and S. Tang, "The wanderings of odysseus in 3D scenes," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 20 481–20 491.
- [17] K. Zhao, Y. Zhang, S. Wang, T. Beeler, and S. Tang, "Synthesizing diverse human motions in 3D indoor scenes," in *Proc. IEEE International Conference on Computer Vision*, 2023, pp. 14 738–14 749.
- [18] iCity3D, "iCity3D," 2024. [Online]. Available: <https://icity3d.com/>
- [19] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. Osman, D. Tzionas, and M. J. Black, "Expressive body capture: 3D hands, face, and body from a single image," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10 975–10 985.
- [20] D. J. Brady, M. E. Gehm, R. A. Stack, D. L. Marks, D. S. Kittle, D. R. Golish, E. Vera, and S. D. Feller, "Multiscale gigapixel photography," *Nature*, vol. 486, no. 7403, pp. 386–389, 2012.
- [21] X. Yuan, L. Fang, Q. Dai, D. J. Brady, and Y. Liu, "Multiscale gigapixel video: A cross resolution image matching and warping approach," in *2017 IEEE International Conference on Computational Photography (ICCP)*. IEEE, 2017, pp. 1–9.
- [22] X. Wang, X. Zhang, Y. Zhu, Y. Guo, X. Yuan, L. Xiang, Z. Wang, G. Ding, D. Brady, Q. Dai *et al.*, "PANDA: A gigapixel-level human-centric video dataset," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3268–3278.
- [23] C. Liu, H. Wei, J. Yang, J. Liu, W. Li, Y. Guo, and L. Fang, "GigaHumanDet: Exploring full-body detection on gigapixel-level images," in *Proc. AAAI Conference on Artificial Intelligence*, vol. 38, no. 9, 2024, pp. 10 092–10 100.
- [24] W. Mo, W. Zhang, H. Wei, R. Cao, Y. Ke, and Y. Luo, "PVDet: Towards pedestrian and vehicle detection on gigapixel-level images," *Engineering Applications of Artificial Intelligence*, vol. 118, p. 105705, 2023.
- [25] W. Li, R. Zhang, H. Lin, Y. Guo, C. Ma, and X. Yang, "SaccadeDet: A novel dual-stage architecture for rapid and accurate detection in gigapixel images," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2024, pp. 392–408.
- [26] J. Zhang, L. Gu, Y.-K. Lai, X. Wang, and K. Li, "Towards grouping in large scenes with occlusion-aware spatio-temporal transformers," *IEEE Transactions on Circuits and Systems for Video Technology*, 2023.
- [27] K. Li, Y. Liu, Y.-K. Lai, and J. Yang, "MILI: Multi-person inference from a low-resolution image," *Fundamental Research*, vol. 3, no. 3, pp. 434–441, 2023.
- [28] W.-L. Wei, J.-C. Lin, T.-L. Liu, and H.-Y. M. Liao, "Capturing humans in motion: Temporal-attentive 3D human pose and shape estimation from monocular video," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 13 211–13 220.
- [29] X. Xu, H. Chen, F. Moreno-Noguer, L. A. Jeni, and F. De la Torre, "3D human pose, shape and texture from low-resolution images and videos," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 9, pp. 4490–4504, 2021.
- [30] S. Shin, J. Kim, E. Halilaj, and M. J. Black, "WHAM: Reconstructing world-grounded humans with accurate 3D motion," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2024, pp. 2070–2080.
- [31] J. Zhang, D. Yu, J. H. Liew, X. Nie, and J. Feng, "Body meshes as points," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2021, pp. 546–556.
- [32] Y. Sun, Q. Bao, W. Liu, Y. Fu, M. J. Black, and T. Mei, "Monocular, one-stage, regression of multiple 3D people," in *Proc. IEEE International Conference on Computer Vision*, 2021, pp. 11 179–11 188.
- [33] Y. Sun, W. Liu, Q. Bao, Y. Fu, T. Mei, and M. J. Black, "Putting people in their place: Monocular regression of 3D people in depth," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 13 243–13 252.
- [34] D. Rempe, T. Birdal, A. Hertzmann, J. Yang, S. Sridhar, and L. J. Guibas, "HuMoR: 3D human motion model for robust pose estimation," in *Proc. IEEE International Conference on Computer Vision*, 2021, pp. 11 488–11 499.
- [35] C. He, J. Saito, J. Zachary, H. Rushmeier, and Y. Zhou, "NeMF: Neural motion fields for kinematic animation," *Proc. Adv. Neural Inform. Process. Syst.*, vol. 35, pp. 4244–4256, 2022.
- [36] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, "SMPL: A skinned multi-person linear model," *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, vol. 34, no. 6, pp. 248:1–248:16, Oct. 2015.
- [37] Y. Li, H. Mao, R. Girshick, and K. He, "Exploring plain vision transformer backbones for object detection," in *Proc. European Conference on Computer Vision*. Springer, 2022, pp. 280–296.
- [38] Z. Yang, A. Zeng, C. Yuan, and Y. Li, "Effective whole-body pose estimation with two-stages distillation," in *Proc. IEEE International Conference on Computer Vision*, 2023, pp. 4210–4220.
- [39] S. Goel, G. Pavlakos, J. Rajasegaran, A. Kanazawa, and J. Malik, "Humans in 4D: Reconstructing and tracking humans with transformers," in *Proc. IEEE International Conference on Computer Vision*, 2023, pp. 14 783–14 794.
- [40] J. Rajasegaran, G. Pavlakos, A. Kanazawa, and J. Malik, "Tracking people by predicting 3D appearance, location and pose," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 2740–2749.
- [41] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. A. Osman, D. Tzionas, and M. J. Black, "Expressive body capture: 3D hands, face, and body from a single image," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [42] X. Chen, B. Jiang, W. Liu, Z. Huang, B. Fu, T. Chen, and G. Yu, "Executing your commands via motion diffusion in latent space," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2023, pp. 18 000–18 010.
- [43] D. Rempe, L. J. Guibas, A. Hertzmann, B. Russell, R. Villegas, and J. Yang, "Contact and human dynamics from monocular video," in *Proc. European Conference on Computer Vision*. Springer, 2020, pp. 71–87.
- [44] S. Zhang, Y. Zhang, F. Bogo, M. Pollefeys, and S. Tang, "Learning motion priors for 4d human body capture in 3D scenes," in *Proc. IEEE International Conference on Computer Vision*, 2021, pp. 11 343–11 353.
- [45] C. Guo, S. Zou, X. Zuo, S. Wang, W. Ji, X. Li, and L. Cheng, "Generating diverse and natural 3D human motions from text," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5152–5161.
- [46] S. Zhang, B. L. Bhatnagar, Y. Xu, A. Winkler, P. Kadlecsek, S. Tang, and F. Bogo, "RoHM: Robust human motion reconstruction via diffusion," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2024, pp. 14 606–14 617.
- [47] N. Mahmood, N. Ghorbani, N. F. Troje, G. Pons-Moll, and M. J. Black, "AMASS: Archive of motion capture as surface shapes," in *Proc. IEEE International Conference on Computer Vision*, 2019, pp. 5442–5451.
- [48] P. Besse, B. Guillouet, J.-M. Loubes, and R. François, "Review and perspective for distance based trajectory clustering," *arXiv preprint arXiv:1508.04904*, 2015.
- [49] B. J. Frey and D. Dueck, "Clustering by passing messages between data points," *science*, vol. 315, no. 5814, pp. 972–976, 2007.
- [50] M. Cuturi and M. Blondel, "Soft-DTW: a differentiable loss function for time-series," in *International conference on machine learning*. PMLR, 2017, pp. 894–903.

- [51] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," *IEEE transactions on acoustics, speech, and signal processing*, vol. 26, no. 1, pp. 43–49, 1978.
- [52] Z. Yang, Z. Cai, H. Mei, S. Liu, Z. Chen, W. Xiao, Y. Wei, Z. Qing, C. Wei, B. Dai *et al.*, "SynBody: Synthetic dataset with layered human models for 3D human perception and modeling," in *Proc. IEEE International Conference on Computer Vision*, 2023, pp. 20 282–20 292.
- [53] X. Liu, T. Zhou, H. Kang, J. Ma, Z. Wang, J. Huang, W. Weng, Y.-K. Lai, and K. Li, "RESCUE: Crowd evacuation simulation via controlling sdm-united characters," in *2025 International Conference on Computer Vision (ICCV)*, 2025.
- [54] J. Zhen, Q. Fang, J. Sun, W. Liu, W. Jiang, H. Bao, and X. Zhou, "SMAP: Single-shot multi-person absolute 3D pose estimation," in *Proc. European Conference on Computer Vision*. Springer, 2020, pp. 550–566.
- [55] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimesheine, L. Antiga *et al.*, "PyTorch: An imperative style, high-performance deep learning library," *Proc. Adv. Neural Inform. Process. Syst.*, vol. 32, 2019.
- [56] G. Hinton, N. Srivastava, and K. Swersky, "Neural networks for machine learning lecture 6a overview of mini-batch gradient descent," *Cited on*, vol. 14, no. 8, p. 2, 2012.
- [57] B. Huang, J. Ju, Y. Shu, and Y. Wang, "Simultaneously recovering multi-person meshes and multi-view cameras with human semantics," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 34, no. 6, pp. 4229–4242, 2023.



Hao Wen received the B.E. degree from the Tianjin University of Mathematics, Tianjin, China in 2017. He is currently pursuing the Ph.D. degree with the computer technology in Tianjin University, Tianjin, China. His interests are mainly in 3D generative models, human motion generation, and human pose and shape estimation.




Hongbo Kang received the B.E. degree from Jishou University, Jishou, China, in 2021, and the Master's degree from Chongqing University of Technology, Chongqing, China, in 2024. He is currently pursuing a Ph.D. in Electronic Information at Tianjin University, Tianjin, China. His research interests include 3D human reconstruction and motion generation.



Jian Ma received the B.S. degree from Jiangsu University of Science and Technology, Zhenjiang, China, in 2016, the M.S. degree from Tianjin University, Tianjin, China, in 2019, and Ph.D degree from University of Bristol, Bristol, U.K. He is currently an assistant researcher in the College of Intelligence and Computing, Tianjin University, Tianjin, China. His research interests include video understanding.



Jing Huang received his BE degree from Tianjin University, Tianjin, China, in 2021. He is currently pursuing a PhD in computer science at Tianjin University. His research focuses on 3D reconstruction and motion generation.



Yuanwang Yang received the B.E. degree from Hebei University of Technology, Tianjin, China, in 2023. He is currently pursuing the M.E. degree in computer science in Tianjin University. His research interests include 3D human reconstruction and motion generation.



Haozhe Lin received the B.S degree from Central South University, and the Ph.D. degree from the Department of Automation at Tsinghua University, China. He is currently an assistant researcher with Beijing National Research Center for Information Science and Technology at Tsinghua University, China. He received the Best Student Paper Award from the 2018 IEEE International Conference on Web Services. His research interests include vision intelligence, gigapixel video processing, and complex scene understanding.



Yu-Kun Lai (Senior Member, IEEE) received his bachelor and Ph.D. degrees in computer science from Tsinghua University in 2003 and 2008, respectively. He is currently a Professor in the School of Computer Science & Informatics, Cardiff University. His research interests include computer graphics, geometry processing, image processing and computer vision. He is on the editorial boards of *IEEE Transactions on Visualization and Computer Graphics* and *The Visual Computer*.



Kun Li (Senior Member, IEEE) received the B.E. degree from Beijing University of Posts and Telecommunications, Beijing, China, in 2006, and the master and Ph.D. degrees from Tsinghua University, Beijing, in 2011. She is currently a Professor with the College of Intelligence and Computing, Tianjin University, Tianjin, China. She was the recipient of the CSIG Shi Qingyun Award for Women Scientists. Her research interests include 3D reconstruction and AIGC.