

# Motion Keyframe Interpolation for Any Human Skeleton via Temporally Consistent Point Cloud Sampling and Reconstruction

Clinton Mo<sup>1</sup>, Kun Hu<sup>1,\*</sup>, Chengjiang Long<sup>2</sup>, Dong Yuan<sup>1</sup>, and Zhiyong Wang<sup>1</sup>

<sup>1</sup> The University of Sydney, Darlington NSW 2008, Australia  
c1mo6615@uni.sydney.edu.au, {kun.hu, zhiyong.wang,  
dong.yuan}@sydney.edu.au

<sup>2</sup> Meta Reality Labs, Burlingame, CA, USA c1ong1@meta.com

**Abstract.** In the character animation field, modern supervised keyframe interpolation models have demonstrated exceptional performance in constructing natural human motions from sparse pose definitions. As supervised models, large motion datasets are necessary to facilitate the learning process; however, since motion is represented with fixed hierarchical skeletons, such datasets are incompatible for skeletons outside the datasets' native configurations. Consequently, the expected availability of a motion dataset for desired skeletons severely hinders the feasibility of learned interpolation in practice. To combat this limitation, we propose Point Cloud-based Motion Representation Learning (PC-MRL), an *unsupervised* approach to enabling cross-compatibility between skeletons for motion interpolation learning. PC-MRL consists of a skeleton obfuscation strategy using *temporal point cloud sampling*, and an *unsupervised skeleton reconstruction* method from point clouds. We devise a temporal point-wise K-nearest neighbors loss for unsupervised learning. Moreover, we propose First-frame Offset Quaternion (FOQ) and Rest Pose Augmentation (RPA) strategies to overcome necessary limitations of our unsupervised point cloud-to-skeletal motion process. Comprehensive experiments demonstrate the effectiveness of PC-MRL in motion interpolation for desired skeletons without supervision from native datasets.


**Keywords:** 3D point clouds · Human body motion · Dataset creation

## 1 Introduction


3D character animation workflows largely rely on the concepts of keyframing with interpolation, often referred to as the pose-to-pose principle [13]. By defining key poses at correct timings, algorithms can be employed to generate intermediate poses, and thereby eliminating the need for defining each frame individually [5, 40]. Although this keyframe-based workflow is less costly than frame-by-frame production, human motion is often complex, and natural depictions

---

\* Corresponding author.



**Fig. 1:** Motion interpolation of walking keyframes conducted by our approach.



**Fig. 2:** Our method leverages point cloud representations (c) to obfuscate the hierarchical dependencies of skeletal motion. Joint colours correspond to ordinal indices in original motion data representation. Point clouds are coloured by temporal index.

demand a significant number of keyframes, especially for realistic motions that adhere to physical properties and constraints. Although motion capture (MoCap) is frequently employed as an alternative for achieving realistic human motions, MoCap often nonetheless necessitates subsequent keyframing to correct unwanted motion elements and address other imperfections.

In recent years, the interpolation process has been widely studied, particularly with the advent of machine learning methods for sequential data. Compared to conventional methods such as linear interpolation (LERP), machine learning methods have demonstrated the capability to derive visually natural motions from sparser keyframe sets [12, 18, 33, 35, 38], which we can observe in Fig. 1. Data-driven interpolation approaches require large motion capture datasets to learn the necessary motion features for effective synthesis of transitions between keyframes. In the field of 3D motions, this aspect significantly restricts the feasibility of learned interpolation methods in animation practice, as motion/pose data are tied to specific skeletal configurations and are not cross-compatible. That is, the dataset’s native skeleton, i.e. *source skeleton*, is structurally different and incompatible with the desired skeleton, i.e. *target skeleton*.

Therefore, we propose a novel *point cloud-based motion representation learning* (PC-MRL) approach, in which skeletal hierarchies are obfuscated by sampling into the point cloud medium, and reconstructed into skeletal motion data using an *unsupervised* neural network. Unlike raw pose and motion data, the

point cloud format is a non-hierarchical representation, and effectively obscures any skeletal configuration as displayed in Fig. 2. PC-MRL samples a set of points around the rest position of skeletons using a combination of uniform and normal distributions, and geometrically parents each point to an associated bone in order to effectively represent the associated skeletons’ motion data in a temporally consistent manner. Subsequently, for the reconstruction of motion data into a given desired skeleton, we propose an *unsupervised* learning scheme, featuring a K-nearest neighbors (KNN) loss between cross-skeleton point clouds to optimize a transformer neural network. The network processes the point cloud-based motion representations and generates the corresponding target skeletal motion representations. By obscuring both source and predicted target motion data into the point cloud space, our strategy allows for a direct comparison, utilizing KNN to minimize the visual difference between the two point clouds.

The unsupervised nature of our target skeleton reconstruction scheme creates two main disparities between the expected and predicted motion features. Firstly, the exact roll rotation axis of bones cannot be represented in the point cloud format. To this end, we introduce *First-frame Offset Quaternion* (FOQ) representations to incorporate relative roll values, which are obtainable from *temporally consistent* point clouds. This approach is agnostic to absolute roll rotations, providing a standardisation mechanism for the resulting motion data within the context of motion interpolation learning. Second, KNN-based objectives learn skeletal representations in a geometrically optimal manner, which does not always reflect the desired skeletal behaviour. This disparity is commonly observed with smaller skeletal features, such as shoulder and hip bones. As such, we augment training motions using Rest Pose Augmentation (RPA) to increase the range of skeletal behaviours that our interpolation model learns. Extensive experiments demonstrate the effectiveness of our proposed method, achieving performance targets near the level of direct dataset supervision.

In summary, this paper presents the following key contributions:

1. We propose a novel method for achieving unsupervised human motion reconstruction from point cloud data, enabling skeleton-agnostic motion interpolation learning for the first time.
2. To train motion interpolation models, we formulate first-frame offset quaternions to represent bone rotations with relative roll data, as well as a rest position augmentation strategy to address skeletal configuration variability.
3. We perform comprehensive experiments to demonstrate our method’s effectiveness towards learning motion interpolation, despite the absence of directly compatible datasets in the training process.

## 2 Related works

We explore existing research on motion interpolation approaches, and *motion data modelling* methods in general. In addition, our proposed approach bears resemblance to *unsupervised approaches* for the motion re-targeting research problem, and as such, we will also analyse existing methods in this field.

## 2.1 Data-Driven Motion Modelling Methods


Contemporary research in motion data modelling predominantly explore learned data-driven methods, using neural networks to capture, model, and establish correlations between the motion features of various skeletal configurations [39, 49, 53]. Motion features alone have demonstrated considerable efficacy in cross-skeleton imitation using a variety of strategies, including task-driven objectives [19, 43], diffusion-based generation [25, 47, 50], and latent feature consistency [1–3, 7, 44]. Latent consistency in particular learns inter-skeleton correlations without the need for paired datasets, which reduces its barrier of entry in practice.

Likewise, motion interpolation strategies have evolved through the use of learning based approaches. The inherent numerical precision of motion data has traditionally posed a challenge for the adoption of the learning approaches, due to their approximate nature [21, 27]. This challenge is further amplified by the temporally sparse distribution of keyframes. Recently, a number of recurrent neural network based methods [17, 18, 41, 51] and transformer-based networks [12, 33, 35] have demonstrated encouraging interpolation performance for the transition between keyframes. Additionally, various other motion modelling methods are capable for keyframe based interpolation, albeit with limited intermediate frame representation capabilities [19, 25, 43].

A major constraint inherent in all data-driven motion modelling methods is their dependency on datasets featuring specific skeletal configurations. To address this, skeleton-free training strategies have been explored for elementary tasks involving 3D characters. These include vertex-based pose transfer [7, 8, 30, 45], and point cloud-based human shape reconstruction [23, 46]. Notably, in these methods, 3D mesh and point cloud coordinates serve as general 3D data representations. Concurrently, other research has significantly minimized the required volume of training data, bringing it down to single example sequences. This reduction has been achieved through the use of patch matching techniques [28] and imitation learning within physics-based simulations [29, 37]. Reinforcement learning has facilitated the generation of goal-driven motion without pre-existing datasets [31, 36, 48]. However, these simulation-based methods have typically been either exceedingly complex to implement in practice, or yield results that are too inflexible and/or prone to errors for animation workflows.

## 2.2 Motion Re-Targeting without Supervision

Conventional approaches in motion re-targeting have been a cornerstone of 3D animation for decades, primarily dedicated to converting motion capture data into 3D skeletal motions [4, 15]. Typically, raw motion capture data is optically recorded, yielding primarily spatial information. In contrast, motion data is largely rotational in nature [5], often represented in  $SO(3)$  space. Therefore, a principal objective of motion re-targeting solutions has been to effectively bridge these two distinct types of representations. In [14], extensive manually adjustable constraints were introduced for flexible re-targeting between skeletons of identical topology (i.e., same hierarchy, different bone lengths and rest poses). Key



**Fig. 3:** Illustration of point cloud-based motion representation learning (PC-MRL) pipeline for keyframe interpolation to any human skeleton. Our unsupervised point cloud-to-pose decoding approach **(I)** first adapts motion features from existing datasets to new skeletons. We then employ the adapted motion data to supervise motion interpolation models **(II)**, using FOQ and rest pose augmentations to address remaining discrepancies between the adapted and expected target skeleton motion patterns.

practices in this approach, including end-effector matching, joint contacts, and inter-skeleton point correlations, are still relevant to current animation pipelines. Point correlations using Inverse Kinematics (IK) have also been proposed as an intermediate representation to facilitate the adaptation between topologically distinct skeletons [9, 34], albeit with hierarchical pairing limitations.

In response, morphologically-independent motion control systems were proposed using IK joint chains and joint-wise constraints [20, 26]. Such systems often expect highly standardised motion skeleton features, and are generally impractical for re-targeting motions produced for different control schemes. Additionally, the reliance on Euler angles for extensive classical re-targeting constraints poses a variety of issues regarding rotational freedom [11], as well as compatibility with animation workflows [32], which predominantly use quaternion-based systems.

Despite these advancements, the challenge of arbitrary skeleton re-targeting from motion datasets still remains largely unresolved, limiting any widespread application of data-driven pipelines. Our paper aims to effectively address this pervasive limitation.

### 3 Methodology

As illustrated in Fig. 3, our proposed method consists of two main components: *point cloud-to-motion reconstruction* and *motion keyframe interpolation*.

#### 3.1 Quaternion-based Motion Representations

We first declare a set of notations for representing skeletons. A skeleton  $S$  is defined by a tree of bones  $\mathcal{B}$ . A bone  $\mathbf{v}_n \in \mathcal{B}$  is defined by its bone parent  $\mathbf{v}_m \in \mathcal{B}$ , and a *head position*  $\mathbf{H}_n \in \mathbb{R}^3$  which denotes the bone’s base offset from

its parent head  $\mathbf{H}_m$ . In addition, a *tail position*  $\mathbf{T}_n \in \mathbb{R}^3$  helps to define the end point of the bone’s line segment representation, starting from its head position, of  $\mathbf{v}_n$ . The root bone  $\mathbf{v}_0$  has no parent, and thus no head position  $\mathbf{H}_0$ . The set of  $\mathbf{H}_n$  and  $\mathbf{T}_n$  for each bone  $\mathbf{v}_n \in \mathcal{B}$  make up the *rest position* of the skeleton.

Next, we define the properties of motion data. At any given frame  $t$ , a skeleton’s *pose position* can be represented as global *3D positions*  $\mathbf{p}_{n,t} \in \mathbb{R}^3$  and global *quaternion rotations*  $\mathbf{q}_{n,t} \in \mathbb{R}^4$  for each bone  $\mathbf{v}_n \in \mathcal{B}$ . For a given motion sequence  $M$ , only the root bone is given a 3D position in pose position, i.e.  $\mathbf{p}_{0,t}$ , as all other pose positions are derived using Forwards Kinematics (FK [11,24]):

$$\mathbf{p}_{n,t} = \mathbf{p}_{m,t} + QR(\mathbf{H}_n, \mathbf{q}_{m,t}), \quad (1)$$

where  $QR$  is the quaternion rotation function, and  $\mathbf{v}_m$  is the parent joint. In summary, a motion sequence  $M$  of  $|M|$  frames is defined by:

$$M = \{\mathbf{p}_{0,t} \cup \{\mathbf{q}_{n,t} \mid \mathbf{v}_n \in \mathcal{B}\} \mid 0 \leq t < |M|\} \quad (2)$$

### 3.2 Point Cloud Obfuscation and Skeleton Reconstruction

Our approach is based on the hypothesis that 3D motion sequences can be visually represented by a more universal format: 3D point cloud sequences. By geometrically sampling the volume around a skeleton, point cloud data becomes a non-hierarchical representation of its pose(s), obfuscating the original skeletal configuration. Due to this decoupling between the skeleton configuration and motion representation, a successful motion data reconstruction from point cloud data would be the first step to enabling learnable cross-skeleton motion features. Formally, for a given source skeleton  $S_A$  with an associated motion sequence  $M_A$ , and any humanoid skeletal configuration  $S_B$ , the point cloud obfuscation and reconstruction module is to learn a function that can adequately perform  $F_{S_A \rightarrow S_B}(M_A) = \hat{M}_B$ , where  $\hat{M}_B$  is a visually similar motion adaptation of  $M_A$  on skeleton  $S_B$ , which is an estimation of  $M_B$ .

**Motion Data Obfuscation with Point Clouds** To sample the point cloud  $\mathcal{X}$ , we first generate a set of points  $u \in \mathcal{X}$  along the line segments defined by the global head and tail positions for each bone  $\mathbf{v}_n$ . In the bone-local space to  $\mathbf{v}_n$ , this would mean sampling  $u$  by a uniformly distributed factor  $\alpha \sim \mathcal{U}_{[0,1]}$ . Specifically,  $\alpha = 0$  aligns with the head position of the bone, and  $\alpha = 1$  indicates the tail position. By further sampling  $u$  via a normal distribution, our sampling strategy can capture the surrounding volume around the skeleton. Formally, to sample  $u$  from its associated bone  $\mathbf{v}_n$  within a standard deviation of  $\sigma$ , we have:

$$\begin{aligned} \mathbf{H}_x &\sim \mathcal{N}(\alpha \mathbf{T}_n, \sigma), \\ \mathbf{p}_{u,t} &= \mathbf{p}_{n,t} + QR(\mathbf{H}_x, \mathbf{q}_{n,t}) \end{aligned} \quad (3)$$

Notably, this sampling process as in Eq. 3 is associated with the skeletal configuration only, by which the sampled points are temporally consistent for further processing with the motion sequence. Specifically, each point maintains a

constant position from its associated bone in the bone-local space. The lateral distance of each point from their parent bone segment allows point clouds to react to the bone’s rotational roll axis.

In a pure point cloud representation, it is possible to remove the relationships between bones. However, to clearly distinguish symmetrical body features in different skeletons, we introduce and embed body group associations for the points during our reconstruction process. Specifically, every bone is categorised by one of these body groups: *[spine, left arm, right arm, left leg, right leg]*. Each sampled point is assigned a group attribute in line with the associated bone. The attribute is represented by a one-hot feature vector  $\mathbf{g}_u$ . To this end, we characterize  $u = \{\mathbf{p}_{u,t}, \mathbf{g}_u\} \in \mathcal{X}$ .

**Skeleton Reconstruction** Given a sampled point cloud representation  $\mathcal{X}_A$  from  $S_A$  and  $M_A$  as input, we aim to train a motion adaptation function  $F_{S_A \rightarrow S_B}$ , which is tasked with producing a visually similar skeletal motion  $\hat{M}_B$  for  $S_B$ , using a temporal and set-based neural network. As shown in Fig. 3, to construct  $F_{S_A \rightarrow S_B}$ , we employ a Point Transformer [52] to learn embeddings for the unordered point cloud sets frame-wisely, and a standard temporal transformer to process the sequence and decode  $\mathcal{X}_A$  into  $\hat{M}_B$ .

To optimize  $F_{S_A \rightarrow S_B}$ , our objective function maximises the similarity between the sampled point clouds of both  $M_A$  and  $\hat{M}_B$ .  $\mathcal{X}_A$  can be viewed as the ground truth to  $\mathcal{X}_B$ , which enables an unsupervised strategy for  $F_{S_A \rightarrow S_B}$ . In detail, we first sample a point cloud  $\mathcal{X}_B$  from  $S_B$  based on  $\hat{M}_B$  derived by  $F_{S_A \rightarrow S_B}$ . Next, a **temporally consistent KNN-Loss objective** is devised to minimise the  $\ell_2$  distance between any given point of  $\mathcal{X}_B$  from its  $K$  nearest neighbouring points in  $\mathcal{X}_A$ . The  $K$  nearest neighbours are determined based on inter-point distance throughout the entire motion sequence, instead of a frame-wise setting, as shown in Fig. 4. The points with differing body groups are excluded from the neighbour search. Mathematically, we have:

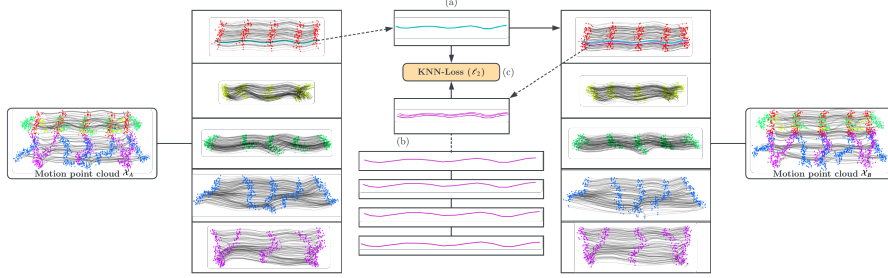
$$\delta(u_A, u_B) = \begin{cases} \sum_t \|\mathbf{p}_{u_A,t} - \mathbf{p}_{u_B,t}\|_2, & \text{if } \mathbf{g}_{u_A} = \mathbf{g}_{u_B}, \\ \inf, & \text{otherwise,} \end{cases} \quad (4)$$

$$\mathcal{L}_{\text{KNN}}(\mathcal{X}_A, \mathcal{X}_B) = \sum_{u_A \in \mathcal{X}_A, u_B \in N_k(u_A, \mathcal{X}_B)} \delta(u_A, u_B), \quad (5)$$

where  $N_k(u_A, \mathcal{X}_B)$  is the set containing points that are the  $k$ -nearest neighbors of  $u_A$  in  $\mathcal{X}_B$ , based on  $\delta(u_A, u_B)$  distance.

In addition to the KNN-Loss, we introduce an optional end-effector loss in the skeleton space for the hand, foot, and head bones, which provides direct positional guidance for matched end-effector pairs  $E_{AB}$ :

$$\mathcal{L}_{\text{end}} = \sum_{(\mathbf{v}_A, \mathbf{v}_B) \in E_{AB}} \frac{1}{|M_A|} \sum_t \|\mathbf{p}_{\mathbf{v}_A,t} - \mathbf{p}_{\mathbf{v}_B,t}\|_2. \quad (6)$$



**Fig. 4:** Visualisation of our temporally consistent KNN objective, processed independently per body group. For each point from  $\mathcal{X}_A$  (a), the total  $\ell_2$  distance throughout the sequence is measured between each point (b) from  $\mathcal{X}_B$ . The  $\ell_2$  distances from the K-nearest points are summed to produce KNN-Loss (c).

### 3.3 First-Frame Offset Quaternions

Point clouds lack absolute roll axis information for each bone. As such, we introduce a First-frame Offset Quaternion (FOQ) transform to incorporate the available *relative roll* data for motion modelling. We formulate all quaternions of a motion sequence relative to their values from the initial frame. In detail, for any quaternion sequence  $Q = \{\mathbf{q}_0, \mathbf{q}_1, \mathbf{q}_2, \dots\}$ , we can simply obtain its FOQ transform as  $\{\mathbf{q}_0 \times \mathbf{q}_0^{-1}, \mathbf{q}_1 \times \mathbf{q}_0^{-1}, \mathbf{q}_2 \times \mathbf{q}_0^{-1}, \dots\}$ , where  $\mathbf{q}_0^{-1}$  is the quaternion conjugate of  $\mathbf{q}_0$ . An FOQ sequence can easily be converted back to a quaternion sequence as long as any  $\mathbf{q}_t \in Q$  is known.

Crucially, FOQ is a **roll-invariant** rotation representation, that is, FOQ remains identical regardless of a bone’s absolute roll position in the original data. Trivially, to adjust the rotational roll of a bone  $\mathbf{v}_n$ , we perform the quaternion multiplication for each frame:  $\mathbf{q}_{n,t} \times \mathbf{r}$ , where  $\mathbf{r}$  is a roll quaternion. As a quaternion on the roll axis of  $\mathbf{v}_n$ ,  $\mathbf{r}$  is constrained within  $\mathbf{r} = \alpha + \beta(x\mathbf{i} + y\mathbf{j} + z\mathbf{k})$ , where  $x, y, z$  are the 3D XYZ values of  $\mathbf{T}_n$ , and  $\alpha, \beta$  are scalars that control the roll magnitude. Note that  $\mathbf{r}$  is constant throughout the sequence. Therefore, we can deduce the roll-invariant nature of any constant  $\mathbf{r}$ :

$$\begin{aligned} (\mathbf{q}_{n,t} \times \mathbf{r}) \times (\mathbf{q}_{0,t} \times \mathbf{r})^{-1} &= \mathbf{q}_{n,t} \times \mathbf{r} \times \mathbf{r}^{-1} \times \mathbf{q}_{0,t}^{-1} \\ &= \mathbf{q}_{n,t} \mathbf{q}_{0,t}^{-1} \end{aligned}$$

### 3.4 Motion Interpolation Transformer and Rest Pose Augmentation

To leverage our point cloud-based representation learning for cross-skeleton interpolation, a transformer model - CITL [33] is adapted, allowing for an efficient interpolation pipeline with existing datasets to alternative skeletons. We introduce two key modifications to CITL and its training strategy, focusing on enhancing motion quality. First, quaternion predictions are substituted with FOQ predictions. This is inspired by the RNN-based method  $\text{TG}_{complete}$  [18], which predicts



**Algorithm 1** IK target alignment with maximum real quaternion component

- 
- 1: **Input:**  $\mathbf{p}_{\text{IK}}$  = 3D IK target unit vector
  - 2:  $\mathbf{T}$  = resting tail position of associated bone
  - 3: **Output:**  $\mathbf{q}_{\text{max}}$  = output quaternion, to be applied on original bone quaternion
  - 4:  $\eta \leftarrow \frac{\mathbf{T}}{|\mathbf{T}|}$  ▷ 3D unit coordinates  $[\eta_x, \eta_y, \eta_z]$
  - 5:  $\omega \leftarrow \frac{\mathbf{p}_{\text{IK}}}{|\mathbf{p}_{\text{IK}}|}$  ▷ 3D unit coordinates  $[\omega_x, \omega_y, \omega_z]$
  - 6:  $\mathbf{q}_{\text{min}} \leftarrow \frac{0 + (\eta_x + \omega_x)\mathbf{i} + (\eta_y + \omega_y)\mathbf{j} + (\eta_z + \omega_z)\mathbf{k}}{2}$
  - 7:  $\mathbf{q}_{\text{max}} \leftarrow \mathbf{q}_{\text{min}} \times (0 + \eta_x\mathbf{i} + \eta_y\mathbf{j} + \eta_z\mathbf{k})$
- 

quaternion offsets from the preceding frame rather than raw quaternions, thereby achieving greater generalisability and superior quality in interpolation.

For our model’s CITL-based architecture, we substitute the original sinusoidal positional encoding scheme with learned relative positions [42]. While the original implementation relies on the continual nature of sinusoidal functions, we observe that relative positional embeddings with *zero initialisation* converges upon similarly continuous behaviour, and significantly lowers the complexity of positional attention. The lessened focus on positional relations allows the model to synthesise deeper behaviours within the pose space, improving overall generalisability. As a side effect, the model also accepts arbitrary length inputs beyond the training data length.

Notably, the intended rest pose of a skeletal configuration is often geometrically sub-optimal for point cloud matching, and can manifest in multiple forms. To address this, we devise a **rest position augmentation strategy**, leveraging a per-bone Inverse Kinematics (IK) system. RPA broadens the range of intended rest poses that the learned interpolation model can accommodate. In detail, during each phase of the Forward Kinematics (FK) process, we augment the bone tails with a random additional offset. We re-align the augmented global tail position as closely as possible to its original location using a quaternion rotation with the maximum possible real component. Algorithm 1 describes the method by which this quaternion can be derived. Since the real component describes, in cosine, the amount of rotation around an axis, this minimizes numerical disturbance from RPA while altering the visual representation to the desired extent.

## 4 Experimental Results

### 4.1 Datasets

To thoroughly evaluate the effectiveness of the proposed method, we conducted extensive experiments on the following widely used motion capture datasets:

- **LaFAN1** [18] contains long, high quality motions in controllable video game character styles. The dataset was recorded on a large motion capture stage, and manually cleaned to production standards. The native skeleton of LaFAN1 contains 6 spinal bones, and 4 bones for each limb.

- **Human3.6M** [6, 22] comprises of 3.6 million motion frames. The motions captured in Human3.6M tend to be less dynamic, as they were recorded on a relatively compact 4x3m stage. Its native skeleton includes 5 spinal bones, 4 bones for each limb, and an additional finger bone for each hand.
- **CMU Mocap** [10] is a diverse motion dataset encompassing 113 categories. It features a mix of static and dynamic motions, captured on a 3x8m stage. The skeleton is identical to that of Human3.6M, with the exception of the pelvis, which is divided into two hip bones and a lower back bone.

For PC-MRL, we designated the CMU MoCap skeleton as  $S_B$ . Its native dataset is exclusively utilized for testing purposes, while LaFAN1 and Human3.6M are used simultaneously for training.<sup>3</sup>

## 4.2 Implementation Details

In our experiments, point cloud obfuscation adopted a 256-point setting for the sampling process with  $\sigma = 0.05$  and each point was characterized by a 64-element vector. To maintain consistency, we ensure that all bone offsets and motion positions are defined in **metre** units. For skeletal reconstruction, 4 Point Transformer [52] layers was utilized to construct the neural network. Within each layer, the feature size was doubled, and the size of the point set was downsampled by a factor of 4 through a farthest-first traversal approach, eventually resulting in a singular feature vector. In the final step, two linear layers with ReLU activations were utilized to transform the resulted feature vector in a frame-wise manner, producing the output skeletal representation. During training, we set  $K = 8$  for the KNN loss to optimize the network. Additionally, the use of unit quaternion values in practice relies on modulus division for constraint, a process that can potentially lead to exploding gradients. To address this, we introduced an objective function  $\mathcal{L}_q$  that measures the  $\ell_2$  norms’ difference between the raw quaternion output and the expected unit quaternion. All experiments were trained and evaluated on a single NVIDIA GeForce RTX 3090 GPU.

## 4.3 Motion Interpolation Comparison against Supervised Methods

To demonstrate the efficacy of our method, we produce and compare our method against the conventional LERP method, as well as three state-of-the-art interpolation models trained on the original CMU MoCap dataset. Specifically, we measure the performance of the RNN-based  $TG_{\text{complete}}$  model [18], the BERT-based motion interpolation adaptation [12], and a transformer-based encoder-decoder approach [33]. We additionally provide results on a variant of our method trained without our RPA strategy. The experimental configuration for each state-of-the-art model is identical to their original implementations. To measure interpolation accuracy, we employ the standard  $\ell_2$  positional distance (L2P),  $\ell_2$  quaternion difference (L2Q), and NPSS [16] for visual similarity evaluation [18].


<sup>3</sup> The authors Clinton Mo and Zhiyong Wang are signatories of the dataset licenses and produced all experimental results of the paper. Meta Inc. was not granted access to the datasets.

Motion category	Method	L2P↓			L2Q↓			NPSS↓		
		5	15	30	5	15	30	5	15	30
Basketball	LERP	0.0787	0.4009	0.8033	0.1588	0.5834	0.9611	<b>0.1315</b>	0.6120	1.8202
	TG <sub>complete</sub> [18]	0.1050	0.4419	0.7104	0.1859	0.6439	0.9303	0.1600	0.6579	<b>1.2136</b>
	BERT [12]	0.0889	0.3737	0.6865	0.1585	0.5560	0.8992	0.1689	0.6177	1.4587
	CITL [33]	<b>0.0625</b>	<b>0.2973</b>	<b>0.5708</b>	<b>0.1410</b>	<b>0.4932</b>	<b>0.8011</b>	<b>0.1047</b>	<b>0.4098</b>	<b>1.0959</b>
	PC-MRL w/o RPA	0.0767	0.3534	0.6715	<b>0.1505</b>	<b>0.5431</b>	<b>0.9049</b>	0.1364	0.5596	1.3471
PC-MRL (Ours)	<b>0.0752</b>	<b>0.3443</b>	<b>0.6662</b>	0.1561	0.5495	0.9088	0.1397	<b>0.5572</b>	1.3015	
Golf	LERP	<b>0.0189</b>	0.1229	0.3187	<b>0.0455</b>	<b>0.1870</b>	0.3987	<b>0.0587</b>	0.3979	1.0312
	TG <sub>complete</sub> [18]	0.0473	0.2616	0.4546	0.0693	0.3191	0.5297	0.0954	0.5062	1.1534
	BERT [12]	0.0276	0.1159	0.2967	0.0612	0.2082	0.4258	0.1099	0.3906	0.9255
	CITL [33]	<b>0.0201</b>	<b>0.1043</b>	<b>0.2589</b>	<b>0.0476</b>	<b>0.1766</b>	<b>0.3552</b>	<b>0.0636</b>	<b>0.2581</b>	<b>0.7799</b>
	PC-MRL w/out RPA	0.0274	0.1304	0.3622	0.0555	0.2065	0.4923	0.0911	0.3759	1.2083
PC-MRL (Ours)	0.0273	<b>0.1130</b>	<b>0.2706</b>	0.0598	0.2012	<b>0.3844</b>	0.1049	<b>0.3716</b>	<b>0.8150</b>	
Swimming	LERP	<b>0.0731</b>	<b>0.3680</b>	<b>0.7374</b>	<b>0.1326</b>	<b>0.5327</b>	<b>0.9476</b>	0.2147	0.9222	1.6972
	TG <sub>complete</sub> [18]	0.1292	0.5796	0.9141	0.1845	0.7420	1.1199	0.2660	0.9716	1.7698
	BERT [12]	0.1088	0.3896	0.7516	0.1819	0.5518	0.9622	0.2877	0.9074	1.6682
	CITL [33]	0.0905	0.3895	0.7475	0.1484	0.5535	0.9533	0.2084	<b>0.8317</b>	<b>1.6122</b>
	PC-MRL w/o RPA	<b>0.0722</b>	0.3754	0.7459	<b>0.1345</b>	0.5442	0.9594	<b>0.1944</b>	0.8731	1.6608
PC-MRL (Ours)	0.0734	<b>0.3595</b>	<b>0.7139</b>	0.1355	<b>0.5331</b>	<b>0.9248</b>	<b>0.1956</b>	<b>0.8591</b>	<b>1.6370</b>	
Walking & Running Locomotion	LERP	0.0455	0.2505	0.5527	0.1120	0.3997	0.6233	0.0655	0.2682	0.9071
	TG <sub>complete</sub> [18]	0.0540	0.1786	0.3103	0.1147	0.2866	<b>0.3866</b>	0.0895	0.2616	<b>0.4265</b>
	BERT [12]	0.0552	0.1998	0.3544	0.1111	0.3101	0.4762	0.0901	0.2760	0.6458
	CITL [33]	<b>0.0328</b>	<b>0.1034</b>	<b>0.2427</b>	<b>0.0942</b>	<b>0.2102</b>	<b>0.3384</b>	<b>0.0526</b>	<b>0.1568</b>	<b>0.4080</b>
	PC-MRL w/o RPA	0.0397	0.1657	0.3417	<b>0.1058</b>	0.2879	0.4722	<b>0.0640</b>	0.2288	0.6021
PC-MRL (Ours)	<b>0.0377</b>	<b>0.1392</b>	<b>0.2848</b>	0.1064	<b>0.2778</b>	0.4281	0.0658	<b>0.2129</b>	0.5098	

**Table 1:** Performance comparisons for various motion categories, measured in L2P, L2Q, and NPSS. All methods *except PC-MRL* are trained on the original CMU MoCap dataset. Each motion contains 128 frames with keyframes placed every 5, 15, or 30 frames. The top 2 results for each test are highlighted in purple and blue respectively.

The results listed in Table 1 clearly demonstrate our method’s ability to approach the accuracy levels exhibited by directly supervised state-of-the-art methods. PC-MRL consistently outperforms both the LERP standard and RNN-based TG<sub>complete</sub> model, particularly in longer keyframe interval scenarios where the quality and depth of learned motion features is most critical and impactful. Likewise, the inclusion of RPA also tends to improve the long interval performance of PC-MRL, particularly when precise movements (i.e. golf) or deeper motion features (i.e. locomotion) are expected. Given this, we can observe that RPA definitively improves the overall consistency of the PC-MRL method in scenarios that may be unseen or less effectively represented by our point cloud-to-motion system.

Native supervision, i.e. with CITL, unsurprisingly produces the most accurate interpolation model, with a notable exception of swimming motions, where our PC-MRL supervision produces even stronger results. We believe the abundance of low crawl motions, a similar motion to swimming in LaFAN1, provides more meaningful supervision over the original CMU dataset, which conversely has few similar motions once swimming motions are filtered out. On the other end of the spectrum, i.e. for high data scenarios such as walking and running locomotion from Fig. 5, all models including our PC-MRL approach observe the strongest improvements for learned methods over the naive LERP method. Both cases strongly support our method’s capability to supervise highly intricate motion patterns, despite the incompleteness of the point cloud representation and reliance on relative rotations and geometric optimisation assumptions.



**Fig. 5:** Walking keyframe interpolation examples produced by methods from Table 1. Close-up layered comparisons against the original motion are provided on the right.


#### 4.4 Cross-Skeleton Motion Re-Targeting Experiments

We further conduct an experiment on motion re-targeting to directly benchmark our method against the sole existing state-of-the-art method for unpaired motion re-targeting with unrestricted skeletons, i.e. *primal skeletons* (PS) [1]. For direct motion re-targeting evaluation, we utilized the KNN-Loss and end-effector loss as metrics, owing to their effective measurement of visual similarity and non-sensitivity towards absolute roll axis correctness. Though absolute rolls are generally necessary for a complete motion re-targeting solution, this aspect is out of scope for our project as our main goal is motion interpolation.

Table 2 indicates the superior performance of our method over the state-of-the-art PS method in terms of visual similarity. At low  $\sigma$  values, the higher relative  $\mathcal{L}_{\text{KNN}}$  of the primal skeleton method underscores its geometric deviations

Model	$L_{\text{KNN}}(\mathcal{X}_A, \mathcal{X}_B) \downarrow$					$L_{\text{end}} \downarrow$
	$\sigma = 0$	$\sigma = 0.01$	$\sigma = 0.05$	$\sigma = 0.1$	$\sigma = 0.2$	
Original motion	1.0099	2.0234	5.7399	9.6636	16.9223	0
Primal skeleton [1]	8.1109	8.0353	9.4206	12.8291	20.6332	9.6990
PC-MRL (Ours)	<b>3.8714</b>	<b>4.1231</b>	<b>6.7277</b>	<b>10.3078</b>	<b>17.5261</b>	<b>5.6258</b>

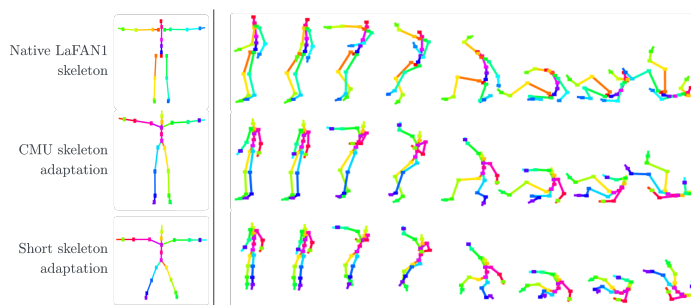
**Table 2:** Average unpaired LaFAN1 to CMU motion re-targeting performance of the primal skeleton method and our point cloud method.  $\mathcal{L}_{\text{KNN}}$  is measured using 1024-point clouds at  $K = 8$ . Independently sampled point clouds from the original LaFAN1 skeleton are provided as optimal  $\mathcal{L}_{\text{KNN}}$  reference. All values are scaled by  $\times 100$ .



**Fig. 6:** Jumping motion re-targeting examples produced on the CMU skeleton by the primal skeleton method (middle row) and our method (bottom row). Sampled point cloud representations of each motion sequence are provided on the right.

from the original motion. In contrast, our method demonstrates significantly better visual adherence to the original geometry. As  $\sigma$  increases, our method approaches near-perfect alignment with the original motion. Figure 6 visually indicates this, showing our method follows the original motion closely, while the latent consistency-based primal skeleton method struggles to generalise and decode the skeleton’s internal structures without directly supervised objectives, such as the end-effector positions provided during training. In addition, Fig 7 demonstrates that, like existing re-targeting approaches, our method is able to produce adequate results on disproportionate skeletons.

Since our approach is largely focused on enabling non-native motion interpolation supervision, we additionally compare our PC-MRL method against a CITL model supervised by PS in Table 3. Due to the suboptimal re-targeting performance of PS, such interpolation models exhibit expectedly poor accuracy.



**Fig. 7:** Re-targeted falling motion from LaFAN1 (top) on the CMU skeleton (middle) and a short custom skeleton (bottom), performed by our point cloud-to-motion model.

Method	L2P↓			L2Q↓			NPSS↓		
	5	15	30	5	15	30	5	15	30
CITL - PS	0.2043	0.6332	1.1524	0.2318	0.7428	1.1442	0.1122	0.4276	1.4622
CITL - Native	0.0495	0.2306	0.4898	0.0979	0.3604	0.6609	0.1149	0.4684	1.1591
PC-MRL	0.0494	0.2337	0.4960	0.1030	0.3771	0.6813	0.1370	0.5433	1.1760

**Table 3:** Average L2P, L2Q, and NPSS performance of CITL trained on native and PS-generated datasets, compared to our method.

## 5 Limitations

Due to the inherent limitations of the point cloud representation, certain constraints on the applicability of our point cloud-to-motion method are inevitable. Primarily, the reliance of our point cloud re-targeting method on the relative nature of First-frame Offset Quaternions (FOQ) for most motion learning tasks necessitates the presence of at least *one known pose*. In the absence of a known pose, relative quaternions are incapable of reconstructing absolute rotational values, which are essential for motion applications. Secondly, due to our method of sampling point clouds with a consistent deviation from each bone segment, the resultant representation inevitably obscures finer details. This includes elements like fingers and facial controls, to an extent that is beyond rectification.

Our proposed objective function relies entirely on geometric segments, necessitating that all bones possess a non-zero length. This stipulation is deemed a reasonable prerequisite for generic human skeletons, as each bone is typically responsible for manipulating a discernible segment of the human body. Due to technical considerations, the skeletal configurations in each of our datasets included at least one bone of zero length, which we ignored in our experiments in favour of global quaternion rotations of all bones with non-zero length.

As a result of these limitations, we refrained from claiming state-of-the-art performance for general motion re-targeting tasks.

## 6 Conclusion

This paper introduces a novel learning-based motion interpolation method designed to enable cross-skeleton compatibility with existing motion datasets. Our proposed method PC-MRL employs a process of point cloud obfuscation and skeleton reconstruction. The point cloud space represents human motions in a non-hierarchic and skeleton-agnostic manner. It enables a KNN-based objective to be optimised without dataset supervision, guiding a neural network to generate high-quality motion features with any target skeleton. We address the rotational information loss of our point cloud format by presenting an offset quaternion strategy compatible with concurrent transformer-based models. Through extensive experiments, we have demonstrated the efficacy of PC-MRL in performing motion interpolation without relying on native motion data. Moreover, PC-MRL has achieved superior visual similarity metrics in the domain of motion re-targeting. We concluded our work by discussing the limitations of PC-MRL.

## References

1. Aberman, K., Li, P., Lischinski, D., Sorkine-Hornung, O., Cohen-Or, D., Chen, B.: Skeleton-aware networks for deep motion retargeting **39**(4), 62–1 (2020) [4](#), [12](#)
2. Aberman, K., Wu, R., Lischinski, D., Chen, B., Cohen-Or, D.: Learning character-agnostic motion for motion retargeting in 2d **38**(4), 1–14 (jul 2019) [4](#)
3. Annabi, L., Ma, Z., et al.: Unsupervised motion retargeting for human-robot imitation. In: Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interactio. ACM (2024) [4](#)
4. Bodenheimer, B., Rose, C., Rosenthal, S., Pella, J.: The process of motion capture: Dealing with the data. In: Computer Animation and Simulation. pp. 3–18. Springer (1997) [4](#)
5. Burtnyk, N., Wein, M.: Computer-generated key-frame animation. Journal of the SMPTE **80**(3), 149–153 (1971) [1](#), [4](#)
6. Catalin Ionescu, Fuxin Li, C.S.: Latent structured models for human pose estimation. In: International Conference on Computer Vision (2011) [10](#)
7. Chen, H., Tang, H., Timofte, R., Gool, L.V., Zhao, G.: Lart: Neural correspondence learning with latent regularization transformer for 3d motion transfer. Advances in Neural Information Processing Systems **36** (2024) [4](#)
8. Chen, J., Li, C., Lee, G.H.: Weakly-supervised 3D pose transfer with keypoints. In: ICCV. pp. 15156–15165 (2023) [4](#)
9. Choi, K.J., Ko, H.S.: Online motion retargetting. The Journal of Visualization and Computer Animation **11**(5), 223–235 (2000) [5](#)
10. CMU: Carnegie-mellon university motion capture database (2003), <http://mocap.cs.cmu.edu/> [10](#)
11. Dam, E.B., Koch, M., Lillholm, M.: Quaternions, interpolation and animation, vol. 2. Citeseer (1998) [5](#), [6](#)
12. Duan, Y., Lin, Y., Zou, Z., Yuan, Y., Qian, Z., Zhang, B.: A unified framework for real time motion completion. In: AAAI. vol. 36, pp. 4459–4467 (2022) [2](#), [4](#), [10](#), [11](#)
13. Frank, T., Johnston, O.: Disney Animation: The Illusion of Life. Abbeville Publishing Group, New York (1981) [1](#)
14. Gleicher, M.: Retargetting motion to new characters. In: Computer Graphics and Interactive Techniques. pp. 33–42 (1998) [4](#)
15. Gleicher, M.: Animation from observation: Motion capture and motion editing **33**(4), 51–54 (1999) [4](#)
16. Gopalakrishnan, A., Mali, A., Kifer, D., Giles, L., Ororbia, A.G.: A neural temporal model for human motion prediction. In: CVPR. pp. 12116–12125 (2019) [10](#)
17. Harvey, F.G., Pal, C.: Recurrent transition networks for character locomotion. In: SIGGRAPH Asia 2018 Technical Briefs, pp. 1–4 (2018) [4](#)
18. Harvey, F.G., Yurick, M., Nowrouzezahrai, D., Pal, C.: Robust motion in-betweening **39**(4), 60–1 (2020) [2](#), [4](#), [8](#), [9](#), [10](#), [11](#)
19. He, C., Saito, J., Zachary, J., Rushmeier, H., Zhou, Y.: Nemf: Neural motion fields for kinematic animation **35**, 4244–4256 (2022) [4](#)
20. Hecker, C., Raabe, B., Enslow, R.W., DeWeese, J., Maynard, J., van Prooijen, K.: Real-time motion retargeting to highly varied user-created morphologies **27**(3), 1–11 (2008) [5](#)
21. Hernandez, A., Gall, J., Moreno-Noguer, F.: Human motion prediction via spatio-temporal inpainting. In: ICCV. pp. 7134–7143 (2019) [4](#)
22. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: 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 (jul 2014) [10](#)
23. Jiang, H., Cai, J., Zheng, J.: Skeleton-aware 3D human shape reconstruction from point clouds. In: ICCV. pp. 5431–5441 (2019) [4](#)
  24. Kantor, I.L., Solodovnikov, A.S., Shenitzer, A.: Hypercomplex numbers: an elementary introduction to algebras, vol. 302. Springer (1989) [6](#)
  25. Karunratanakul, K., Preechakul, K., Suwajanakorn, S., Tang, S.: Guided motion diffusion for controllable human motion synthesis. In: ICCV. pp. 2151–2162 (2023) [4](#)
  26. Kulpa, R., Multon, F., Arnaldi, B.: Morphology-independent representation of motions for interactive human-like animation. In: Eurographics (2005) [5](#)
  27. Lehrmann, A.M., Gehler, P.V., Nowozin, S.: Efficient nonlinear markov models for human motion. In: CVPR. pp. 1314–1321 (2014) [4](#)
  28. Li, P., Aberman, K., Zhang, Z., Hanocka, R., Sorkine-Hornung, O.: Ganimator: Neural motion synthesis from a single sequence **41**(4), 1–12 (2022) [4](#)
  29. Li, Z., Peng, X.B., Abbeel, P., Levine, S., Berseth, G., Sreenath, K.: Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control. arXiv preprint arXiv:2401.16889 (2024) [4](#)
  30. Liao, Z., Yang, J., Saito, J., Pons-Moll, G., Zhou, Y.: Skeleton-free pose transfer for stylized 3D characters. In: ECCV. pp. 640–656. Springer (2022) [4](#)
  31. Makoviychuk, V., Wawrzyniak, L., Guo, Y., Lu, M., Storey, K., Macklin, M., Hoeller, D., Rudin, N., Allshire, A., Handa, A., et al.: Isaac gym: High performance gpu-based physics simulation for robot learning. arXiv preprint arXiv:2108.10470 (2021) [4](#)
  32. Merry, B., Marais, P., Gain, J.: Animation space: A truly linear framework for character animation. ACM Transactions on Graphics (TOG) **25**(4), 1400–1423 (2006) [5](#)
  33. Mo, C.A., Hu, K., Long, C., Wang, Z.: Continuous intermediate token learning with implicit motion manifold for keyframe based motion interpolation. In: CVPR. pp. 13894–13903 (2023) [2](#), [4](#), [8](#), [10](#), [11](#)
  34. Monzani, J.S., Baerlocher, P., Boulic, R., Thalmann, D.: Using an intermediate skeleton and inverse kinematics for motion retargeting. In: Computer Graphics Forum. vol. 19, pp. 11–19. Wiley Online Library (2000) [5](#)
  35. Oreshkin, B.N., Valkanas, A., Harvey, F.G., Ménard, L.S., Bocquet, F., Coates, M.J.: Motion inbetweening via deep  $\delta$ -interpolator (2022) [2](#), [4](#)
  36. Peng, X.B., Abbeel, P., Levine, S., van de Panne, M.: Deepmimic: Example-guided deep reinforcement learning of physics-based character skills **37**(4), 143:1–143:14 (Jul 2018) [4](#)
  37. Peng, X.B., Ma, Z., Abbeel, P., Levine, S., Kanazawa, A.: Amp: Adversarial motion priors for stylized physics-based character control **40**(4) (Jul 2021) [4](#)
  38. Qin, J., Zheng, Y., Zhou, K.: Motion in-betweening via two-stage transformers **41**(6), 1–16 (2022) [2](#)
  39. Reda, D., Won, J., Ye, Y., van de Panne, M., Winkler, A.: Physics-based motion retargeting from sparse inputs. ACM Computer Graphics and Interactive Techniques **6**(3), 1–19 (2023) [4](#)
  40. Reeves, W.T.: Inbetweening for computer animation utilizing moving point constraints **15**(3), 263–269 (1981) [1](#)
  41. Ren, T., Yu, J., Guo, S., Ma, Y., Ouyang, Y., Zeng, Z., Zhang, Y., Qin, Y.: Diverse motion in-betweening from sparse keyframes with dual posture stitching. IEEE Transactions on Visualization and Computer Graphics (2024) [4](#)



42. Shaw, P., Uszkoreit, J., Vaswani, A.: Self-attention with relative position representations. In: NAACL. Association for Computational Linguistics (2018) [9](#)
43. Tiwari, G., Antić, D., Lenssen, J.E., Sarafianos, N., Tung, T., Pons-Moll, G.: Pose-ndf: Modeling human pose manifolds with neural distance fields. In: ECCV. pp. 572–589. Springer (2022) [4](#)
44. Villegas, R., Yang, J., Ceylan, D., Lee, H.: Neural kinematic networks for unsupervised motion retargeting. In: CVPR. pp. 8639–8648 (2018) [4](#)
45. Wang, J., Li, X., Liu, S., De Mello, S., Gallo, O., Wang, X., Kautz, J.: Zero-shot pose transfer for unrigged stylized 3D characters. In: CVPR. pp. 8704–8714 (2023) [4](#)
46. Wang, K., Xie, J., Zhang, G., Liu, L., Yang, J.: Sequential 3D human pose and shape estimation from point clouds. In: CVPR. pp. 7275–7284 (2020) [4](#)
47. Yuan, Y., Song, J., Iqbal, U., Vahdat, A., Kautz, J.: Physdiff: Physics-guided human motion diffusion model. In: ICCV. pp. 16010–16021 (2023) [4](#)
48. Yuan, Y., Wei, S.E., Simon, T., Kitani, K., Saragih, J.: Simpoe: Simulated character control for 3D human pose estimation. In: CVPR. pp. 7159–7169 (2021) [4](#)
49. Zhang, J., Weng, J., Kang, D., Zhao, F., Huang, S., Zhe, X., Bao, L., Shan, Y., Wang, J., Tu, Z.: Skinned motion retargeting with residual perception of motion semantics & geometry. In: CVPR. pp. 13864–13872 (2023) [4](#)
50. Zhang, M., Li, H., Cai, Z., Ren, J., Yang, L., Liu, Z.: Finemogen: Fine-grained spatio-temporal motion generation and editing. *Advances in Neural Information Processing Systems* **36** (2024) [4](#)
51. Zhang, X., van de Panne, M.: Data-driven autocompletion for keyframe animation. In: ACM SIGGRAPH Conference on Motion, Interaction and Games. pp. 1–11 (2018) [4](#)
52. Zhao, H., Jiang, L., Jia, J., Torr, P.H., Koltun, V.: Point transformer. In: ICCV. pp. 16259–16268 (2021) [7](#), [10](#)
53. Zhu, W., Yang, Z., Di, Z., Wu, W., Wang, Y., Loy, C.C.: Mocanet: Motion retargeting in-the-wild via canonicalization networks. In: AAAI. vol. 36, pp. 3617–3625 (2022) [4](#)