

# PP-Motion: Physical-Perceptual Fidelity Evaluation for Human Motion Generation

Sihan Zhao\*  
Tsinghua University  
Beijing, China  
zhaosh024@gmail.com

Zixuan Wang\*  
Tsinghua University  
Beijing, China  
wangzixu21@mails.tsinghua.edu.cn

Tianyu Luan†  
University at Buffalo  
Buffalo, NY, USA  
tianyulu@buffalo.edu


Jia Jia†  
BNRist, Tsinghua University  
Key Laboratory of Pervasive  
Computing, Ministry of Education  
Beijing, China  
jjia@tsinghua.edu.cn

Wentao Zhu  
Eastern Institute of Technology,  
Ningbo  
Ningbo, China  
wtzhu@eitech.edu.cn

Jiebo Luo  
University of Rochester  
Rochester, NY, USA  
jluo@cs.rochester.edu

Junsong Yuan  
University at Buffalo  
Buffalo, NY, USA  
jsyuan@buffalo.edu

Nan Xi  
University at Buffalo  
Buffalo, NY, USA  
nanxi@buffalo.edu



**Figure 1: A motion that looks realistic does not necessarily mean it is physically feasible. Top-left: Motion appears realistic and semantically meaningful to the human eye, yet fails in physics simulation, resulting in a fall (bottom-left). Top-right: Unnatural motion in human perception executes successfully in simulation (bottom-right). This reveals a discrepancy between human perception and physical laws.**

## Abstract

Human motion generation has found widespread applications in AR/VR, film, sports, and medical rehabilitation, offering a cost-effective alternative to traditional motion capture systems. However, evaluating the fidelity of such generated motions is a crucial,

\*Equal contribution.

†Co-corresponding authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

MM '25, Dublin, Ireland

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-2035-2/2025/10  
<https://doi.org/10.1145/3746027.3754940>

multifaceted task. Although previous approaches have attempted at motion fidelity evaluation using human perception or physical constraints, there remains an inherent gap between human-perceived fidelity and physical feasibility. Moreover, the subjective and coarse binary labeling of human perception further undermines the development of a robust data-driven metric. We address these issues by introducing a physical labeling method. This method evaluates motion fidelity by calculating the minimum modifications needed for a motion to align with physical laws. With this approach, we are able to produce fine-grained, continuous physical alignment annotations that serve as objective ground truth. With these annotations, we propose PP-Motion, a novel data-driven metric to evaluate both physical and perceptual fidelity of human motion. To effectively capture underlying physical priors, we employ Pearson's correlation loss for the training of our metric. Additionally, by incorporating a human-based perceptual fidelity loss, our metric

can capture fidelity that simultaneously considers both human perception and physical alignment. Experimental results demonstrate that our metric, PP-Motion, not only aligns with physical laws but also aligns better with human perception of motion fidelity than previous work.

## CCS Concepts

• **Computing methodologies** → **Motion processing; Physical simulation.**

## Keywords

Human Motion Evaluation, Fidelity Metrics

### ACM Reference Format:

Sihan Zhao, Zixuan Wang, Tianyu Luan, Jia Jia, Wentao Zhu, Jiebo Luo, Junsong Yuan, and Nan Xi. 2025. PP-Motion: Physical-Perceptual Fidelity Evaluation for Human Motion Generation. In *Proceedings of the 33rd ACM International Conference on Multimedia (MM '25)*, October 27–31, 2025, Dublin, Ireland. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3746027.3754940>

## 1 Introduction

Human motion generation has found widespread applications in modern industrial production. Whether in AR/VR games, film, content creation, sports, or medical rehabilitation, generated motions can replace complex motion capture systems and on-site filming with actors. Realistic human motion generation can produce large quantities of poses at very low cost, potentially saving significant labor and filming expenses. Therefore, evaluating the fidelity of generated motions is a problem closely tied to real-world applications. However, it remains a challenging task to design a comprehensive evaluation metric due to the multifaceted influence factors.

Previous methods, such as MotionCritic [72], have made good progress in assessing human motion fidelity. MotionCritic introduces a dataset, *MotionPercept*, where human subjects judge the fidelity of a motion. These ratings are later used as labels to train a motion fidelity metric that aligns with human perception. Although human perception of motion fidelity is reasonably accurate and practically useful, the fundamental standard for motion fidelity should not be based solely on human perception. It is even more important to consider whether the motion conforms to physical laws. A motion that looks realistic does not necessarily mean it is physically feasible. For example, as shown in Fig. 1, the motion on the top-left appears realistic and semantically meaningful to human eyes, yet when simulated in a physics engine, the motion cannot be completed and results in a fall on the ground (bottom-left). Conversely, the motion on the top-right may look unusual and meaningless, but it can be executed well in a physics simulation (bottom-right). These examples reveal a discrepancy between human perception and physical laws. Moreover, human annotations of fidelity are subjective; different annotators may have difficulty quantifying the fidelity of the same motion consistently. In MotionCritic [72], the dataset employs a binary “better/worse” classification to avoid the quantification issues of human labeling. However, such coarse labeling lacks fine-grained information and poses challenges for learning a data-driven metric.

To address these two issues, we propose a data-driven method to evaluate motion fidelity that aligns with physical laws. We achieve this by calculating the minimum distance between the test motion and a motion that complies with physical laws. A small minimal distance indicates a high fidelity of the input motion, whereas a large minimal distance implies low fidelity. With this physically grounded definition of fidelity, we can establish fine-grained continuous labels for physical law alignments. In addition, we design a framework that trains a metric, named PP-Motion, by simultaneously utilizing fine-grained physical labels and coarse, discrete human perceptual labels. By designing loss functions that better suit the fine-grained labels, we can more effectively learn the underlying physical law priors. With this approach, we are able to train our metric to better align with physical laws. Furthermore, since human perception is inherently correlated with physical feasibility, this approach also improves the potential for metric design to align with human judgments by learning physical principles.

Specifically, we adopt the motions generated in MotionCritic [72] and create new physically aligned annotations for them. Inspired by PHC [48], we refine every motion in our dataset using reinforcement learning with the help of a physics simulator. We use this approach to make only minimal adjustments while making each motion conform to physical laws. We then compare the difference between the adjusted motion and the original motion to serve as the annotation for physical alignment in the dataset. Such annotations not only closely adhere to our definition of physical fidelity and offer strong interpretability, but also provide continuous, fine-grained labels. These fine-grained annotations offer rich information for the supervision of subsequent metric training. To better learn from these fine-grained physical annotations, we design loss functions based on data correlation, such as Pearson’s correlation loss. Unlike previous classification losses, correlation loss can effectively capture the intrinsic correlations within the data rather than simply comparing categories or numeric values. Without the constraint on scales, the correlation loss can be more easily combined with existing human perception loss functions, enabling our metric to align with both human perception and physical laws, and providing the potential for mutual reinforcement between the two aspects.

In summary, our contributions are as follows:

- We propose a novel fidelity evaluation method, **PP-Motion**, for human motions, which takes into account both physical feasibility and human perception. Our method can evaluate whether a motion is realistically aligned with physical laws and human perception.
- We define and design a fine-grained physical alignment annotation and provide this annotation for existing datasets. This annotation serves as fine-grained physical ground truth for training our metric and has the potential to benefit subsequent metric design.
- We design an effective learning framework that leverages these fine-grained physical annotations. By incorporating correlation-based loss functions (i.e., Pearson’s correlation loss), our approach better learns the physical priors from the labels, while seamlessly combining with existing human perceptual loss functions. This design not only ensures that our metric adheres to physical laws but also has the potential to enhance human-perceived fidelity.

## 2 Related Work

### 2.1 Human Motion Generation

Human motion generation aims to automatically generate natural, fluent, and physically plausible human pose sequences, playing a central role in character animation, human-robot interaction, and embodied agents acting in complex environments. With wide applications and high practical value, motion generation is a foundational problem in both academic and industrial fields. Fueled by rapid advances in deep learning [35], especially generative models [16, 25, 33, 56, 59], extensive research has focused on generating human motions from multimodal signals, including text, action, speech, and music. Action-to-motion [5, 9, 21, 47] aims to synthesize motions from predefined action labels, evolving from retrieval-based to label-conditioned generative models. Text-to-motion [12, 18, 20, 28, 52–54, 63, 66, 67, 69, 74, 80, 83, 85] focuses on mapping natural language to motion, bridging linguistic semantics and physical embodiment. Beyond text, audio-driven motion generation has also seen progress. Music-to-dance methods [37, 39, 61, 65, 68] synthesize motions aligned with musical features such as beat and style. Speech-to-gesture approaches [2, 7, 15, 22, 34, 36, 55, 78] emphasize temporal alignment and semantic expressiveness to convey emotion. Further studies on controllable and editable motion generation [6, 11, 26, 31, 60, 71, 77, 84] focus on generating high quality long-term motions while maintaining maximum faithfulness to various multimodal control conditions. Considering the growing range of real-world applications and generative methods of motion generation, it's critical to establish comprehensive evaluations for generated motions.

### 2.2 Human Motion Evaluation

Designing evaluation metrics for human motion is a complicated and challenging problem. Existing evaluation metrics can mainly be divided into three categories: (1) distance-based metrics, (2) human-perception-based metrics, and (3) physical plausibility metrics. The most commonly used evaluation metrics in early works are distance-based. Distance-based metrics such as Position, Velocity, and Acceleration Errors [1, 4, 8, 10, 15, 27, 32, 34, 36, 43–46, 50, 52, 65, 73, 75, 81, 86] compare generated motions with ground truth, but struggle to capture the diversity of plausible motions. To address this, feature-based metrics [2, 7, 13, 14, 17, 22, 55, 78, 86] are proposed to provide more refined semantic abstraction to calculate distance for similarity assessment between generated and ground truth motions. Among all these distance-based metrics, Fréchet Inception Distance (FID) and average Euclidean distance (DIV) in feature space are widely used to evaluate motion quality and diversity [9, 19, 30, 37–39, 61, 67, 68, 82, 85, 86, 88]. However, FID and DIV measure the distance between or within distributions, making them unsuitable for evaluating the quality of a single action. Feature-space metrics also depend heavily on the effectiveness of the feature extractor, where certain features may lack interpretability. To better assess motion naturalness, smoothness, and plausibility, Wang et al. [72] propose MotionCritic, a data-driven model trained on human-annotated motion preferences. Although MotionCritic aligns well with human perception, there remains a huge gap between motions deemed feasible by humans and those

grounded in physical laws. For physical plausibility evaluation, researchers have proposed rule-based metrics such as foot-ground penetration [24, 57, 58, 64], foot contact [57, 58, 68], foot skating rate [3, 42, 76], and floating [23]. However, these metrics are often heuristic, threshold-sensitive, and too limited to capture overall physical fidelity. Physdiff [79] proposes an overall physics error metric, but remains a naive aggregation of penetration, foot skating, and floating metrics. Li et al. [40] further propose the imitation failure rate (IFR), using a physics engine to test whether a motion can be successfully simulated. However, IFR only gives a binary judgment, offering no graded assessment. Thus, a comprehensive evaluation method that jointly considers human perception and physical feasibility remains an open challenge.

## 3 Method

### 3.1 Problem Formulation

In our approach to designing a motion evaluation metric, the measurement system is defined as follows. For any human motion  $x$ , the aim is to establish a function  $F(\cdot)$  that evaluates the fidelity  $\hat{s}$  of the motion sequence:

$$\hat{s} = F(x; \theta), \quad (1)$$

where  $F(x; \theta)$  is a neural network architecture with parameters  $\theta$ .

The subsequent sections introduce our methodology for both architecting the measurement function  $F(x; \theta)$  and optimizing its parameters. The training procedure can be summarized by the optimization objective:

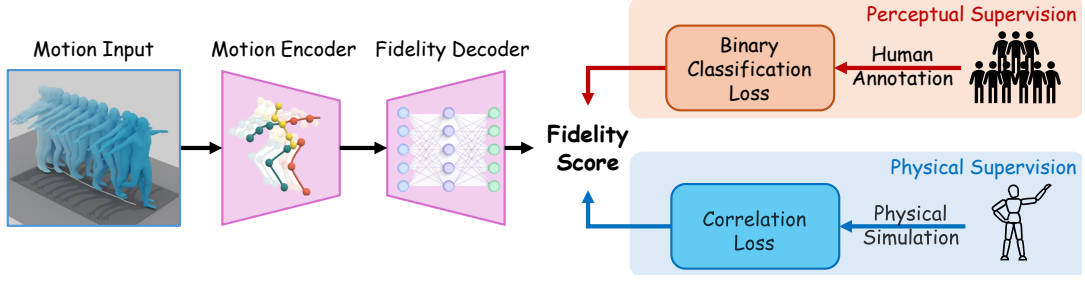
$$\min_{\theta} \mathcal{L}_{\text{prec}}(F(x; \theta), y_{\text{prec}}) + \lambda \mathcal{L}_{\text{phy}}(F(x; \theta), y_{\text{phy}}), \quad (2)$$

where  $\mathcal{L}_{\text{prec}}$  and  $\mathcal{L}_{\text{phy}}$  are the perceptual and physics losses,  $y_{\text{prec}}$  and  $y_{\text{phy}}$  are the perceptual and physics supervision, and  $\lambda$  is a balance weight for these two terms of loss functions. The objective is to jointly optimize the measurement accuracy for both physical and perceptual fidelity.

### 3.2 Physical-Perceptual Motion Metric

The design of our metric is illustrated in Fig. 2. Our network takes a human motion sequence as input. First, the motion is fed into a motion encoder to extract spatio-temporal features. A fidelity decoder is then used to decode these features into a fidelity score. We use annotations from two different sources, physical and perceptual, to supervise the metric training. On the physical side, we analyze motion fidelity using a physics simulator and generate fine-grained annotations for training supervision. A correlation loss between the metric output and the physical annotations encourages the metric to learn from those physical annotations effectively. Meanwhile, the network also learns fidelity from human annotations, ensuring that the fidelity score aligns closely with both physical annotation and human perception.

**Motion encoder.** The motion encoder plays a critical role in determining the quality of the motion features. To extract both spatial and temporal information necessary for fidelity evaluation, we adopt a state-of-the-art spatio-temporal motion encoder. In our experiments, we follow the design proposed in [87]. This encoder is built from  $N$  dual-stream fusion modules, each containing branches for spatial and temporal self-attention and MLP. The spatial layers



**Figure 2: Our metric design and training pipeline.** The network takes a human motion sequence as input, which is processed by a motion encoder to extract spatiotemporal features. These features are then decoded into a fidelity score by a fidelity decoder. The network is trained in a supervised fashion using fine-grained physical annotations alongside human perceptual labels.

capture correlations among different joints within the same time step, while the temporal layers focus on the dynamics of individual joints. This dual-stream design effectively captures comprehensive features required for assessing motion fidelity.

**Fidelity decoder.** We introduce a fidelity decoder module to interpret motion fidelity from the extracted features. Since our fidelity score is a fine-grained continuous value rather than a coarse classification, we design a fidelity decoder to extract the fidelity score from the features. Moreover, given the strong representation capability of our backbone, we can leverage the encoder to extract and fit the fidelity features during training, which allows us to simplify the fidelity decoder design. In practice, we adopt an MLP-based design for the fine-grained fidelity decoder. This simple design not only meets the requirement for producing detailed scores but also avoids imposing a significant extra burden on network training.

**Physical supervision.** Physical supervision is a core module of our method, and it consists of two parts: physical accuracy labels and the corresponding supervision strategy.

To evaluate physical fidelity, we develop a scoring system based on feedback from a physics simulator. For a given input motion, we generate a motion that is as close as possible to the input while satisfying the physical constraints of the simulator (i.e., the nearest regularized motion). The difference between the input pose and the nearest regularized motion represents the motion’s physical rationality. For more details on generating the nearest regularized motion, please refer to Sec. 4. Intuitively, if a pose that is initially physically implausible can be made physically reasonable with only minor adjustments, it is considered to have high physical fidelity (i.e., a small fidelity error). However, if major modifications are needed, the pose is considered to have low physical fidelity (i.e., a large fidelity error). This design enables fine-grained continuous annotation and measurement of physical fidelity.

Our approach to physical supervision primarily focuses on aligning the network’s output with the physical annotations. We have observed that, for designing a physically aligned fidelity score, the absolute value of the output score is less critical than its correlation with the physical labels. Therefore, our supervision targets the correlation between the fidelity score and the physical annotations rather than direct numerical differences as in traditional regression tasks. To this end, we use Pearson’s correlation loss as the core training loss, as detailed in Sec. 3.3.

**Perceptual supervision.** For human-aligned fidelity evaluation, our network also leverages the annotations and training strategy provided by [72]. Specifically, the fidelity annotations used in this branch are derived from human subjects. During training, given a pair of motions with “better” and “worse” labels provided by human subjects, we train a network using a perceptual loss, which encourages the model to assign higher scores to the “better” motion than to the “worse” motion. Notably, even though we did not modify the original training strategy or annotations for this part, our joint training under both physical and perceptual supervision results in a metric that aligns with human perception even better than a model optimized solely for human alignment. This outcome demonstrates that physical and perceptual annotations are well aligned and that fine-grained physical alignment can further boost human perceptual alignment.

### 3.3 Training Loss

Our loss design is mainly divided into two parts: a perceptual loss based on binary human fidelity scoring labels (i.e., better/worse in each pair), and a physical loss based on continuous physical labels.

**Perceptual Loss.** Our perceptual loss follows the design in [72]. For a better-worse motion pair  $x^{(h)}$  and  $x^{(l)}$ , the perceptual loss is defined as:

$$\mathcal{L}_{\text{percept}} = -\mathbb{E}_{(x^{(h)}, x^{(l)})} \left[ \log \sigma \left( F(x^{(h)}) - F(x^{(l)}) \right) \right], \quad (3)$$

where  $F(\cdot)$  is our designed metric, and  $\sigma(\cdot)$  is the sigmoid function.

**Physical Loss.** Our physical loss is designed to learn from the fine-grained physical annotation. We use a Pearson’s correlation loss [51] to learn for physical annotation. The correlation loss is defined as:

$$\mathcal{L}_{\text{corr}} = -\frac{\sum_{i=1}^n (\hat{x}_i - \bar{\hat{x}})(x_i - \bar{x})}{\sqrt{\sum_{i=1}^n (\hat{x}_i - \bar{\hat{x}})^2} \sqrt{\sum_{i=1}^n (x_i - \bar{x})^2}}, \quad (4)$$

where  $\hat{x}_i$  and  $x_i$  indicate the predicted and ground truth motion fidelity scores of the sample  $i$  in the dataset.  $n$  is the total number of samples in the dataset. The average predicted motion fidelity and ground truth motion fidelity are then defined as  $\bar{\hat{s}} = \frac{1}{n} \sum_{i=1}^n \hat{s}_i$  and  $\bar{s} = \frac{1}{n} \sum_{i=1}^n s_i$ .

**Total Loss.** Our total loss function can then be represented as:

$$\mathcal{L} = \mathcal{L}_{\text{percept}} + \lambda \mathcal{L}_{\text{corr}}, \quad (5)$$

**Table 1: Annotation statistics on MotionPercept dataset.**

Annoation	MotionCritic	Ours
Grain	Binary	Continuous
Annoation type	Perceptual	Physical
Categorization	Quadruples	Per-prompt
Score Distribution	$p(0) = 0.75, p(1) = 0.25$	$\sim N(0, 1)$

**Table 2: Quantitative results on imitating motion sequences of MotionPercept, which has three subsets: MDM-Train, MDM-Val, and FLAME. We use pose-based metrics to compare the imitation performance between using only the pretrained model and applying per-sequence fine-tuning. Recon. Err., MPJPE, and PA-MPJPE are measured in millimeters (mm).**

MotionPercept-MDM-Train					
	Recon. Err.↓	MPJPE↓	PA-MPJPE↓	$e_{acc}$ ↓	$e_{vel}$ ↓
Whole dataset pretrain	55.72	36.76	30.60	4.37	6.23
Per data fine-tune	<b>49.65</b>	<b>32.95</b>	<b>27.37</b>	<b>4.03</b>	<b>5.76</b>
MotionPercept-MDM-Val					
	Recon. Err.↓	MPJPE↓	PA-MPJPE↓	$e_{acc}$ ↓	$e_{vel}$ ↓
Whole dataset pretrain	55.49	36.88	30.68	4.27	6.07
Per data fine-tune	<b>50.90</b>	<b>34.10</b>	<b>28.20</b>	<b>4.16</b>	<b>5.89</b>
MotionPercept-FLAME					
	Recon. Err.↓	MPJPE↓	PA-MPJPE↓	$e_{acc}$ ↓	$e_{vel}$ ↓
Whole dataset pretrain	69.20	49.80	37.91	5.63	7.94
Per data fine-tune	<b>53.76</b>	<b>38.33</b>	<b>31.35</b>	<b>5.32</b>	<b>7.26</b>

where  $\lambda$  is the loss weight. The implementation details of model training are provided in the supplementary material.

## 4 Dataset

### 4.1 MotionPercept Dataset

We use the MotionPercept [72] dataset to provide perceptual annotations and also provide generated motions for our physical annotation. MotionPercept is a large-scale dataset of motion perceptual evaluation, in which real humans are invited to select the best or the worst from given motion sets. Each motion set has 4 motions generated by the state-of-the-art motion generation models MDM [67] and FLAME [32] with the same action label or text prompt. Specifically, the MDM model is trained on HumanAct12 [21] and UESTC [29], resulting in a total of 17521 groups of motions, each containing 4 motions. The FLAME model is trained on HumanML3D [19], resulting in 201 groups of motions, each group also comprising 4 motions.

### 4.2 Physical Annotation Generation

We propose a novel annotation method that leverages a physics simulator to provide a fine-grained and interpretable measurement of a motion’s physical accuracy. Specifically, we use the simulator to identify a physically plausible motion  $x'$  that is as close as possible to the input motion  $x$ . The physical error  $e_p$  is then defined as the  $l_2$  norm of the difference between the two motions:

$$e_p = \|x - x'\|_2, \quad (6)$$

where  $\|\cdot\|_2$  is  $l_2$  norm. To obtain an  $x'$  that closely approximates  $x$ , we propose using a physical correction network,  $F_p(x)$ , is defined as:

$$x' = F_p(x), \quad (7)$$

where  $x'$  is the output motion that is aligned with physical laws in physical simulators.

The training of  $F_p(x)$  would require feedback from a physical simulator. Since most state-of-the-art simulators are not designed to backpropagate gradients, we used reinforcement learning to get the physical fidelity feedback from the physical simulator. Moreover, we also constrain the output motion  $x'$  to be close to the input motion  $x$  in terms of translation, rotation, linear velocity, and angular velocity. Specifically, we use the network in PHC [48] as our physical correction network to find  $x'$  for each pose in the MotionPercept [72] dataset. To achieve a lower difference while still aligning with physical constraints, we use the physical reward in PHC [48]. For timestamp  $t$ , the reward can be represented as:

$$r_t = w_{jp}e^{-100\|p'_t - p_t\|} + w_{jr}e^{-10\|q'_t \ominus q_t\|} + w_{jv}e^{-0.1\|v'_t - v_t\|} + w_{j\omega}e^{-0.1\|\omega'_t - \omega_t\|}, \quad (8)$$

where  $p_t$ ,  $q_t$ ,  $v_t$ , and  $\omega_t$  are the translation, rotation, linear velocity, and angular velocity of input motion  $x$  at timestamp  $t$ .  $p'_t$ ,  $q'_t$ ,  $v'_t$ , and  $\omega'_t$  are the translation, rotation, linear velocity, and angular velocity of the output motion  $x'$  from the physical simulator at timestamp  $w_{jp}$ .  $w_{jr}$ ,  $w_{jv}$ , and  $w_{j\omega}$  are the loss weights.  $\ominus$  is the difference between 2 rotations. This reward function makes sure the physical simulator output is close to the original pose, and still aligns with physical laws. In training, we summarize the reward function for all timestamps  $t$  and get the optimal motion  $x'$ .

Our annotation generation contains 2 steps: First, we pretrain the PHC network on the whole MotionCritic dataset using the reward function Eq. (8). Second, we use the same reward function to optimize every single motion in the MotionCritic dataset, so that it will achieve a closer corrected motion for each input motion in the dataset. In Tab. 2 we report the difference between the input motion  $x$  and the corrected motion  $x'$  after the whole dataset pretrain (step 1) and the per-data fine-tune (step 2), respectively. In these tables, we use IsaacGym [49] simulator for physical simulation. Reconstruction error (Recon. Err.) calculates the absolute mean per-joint position error in world coordinates. MPJPE (Mean Per Joint Position Error) measures the mean per-joint position error relative to the root joint, while PA-MPJPE (Procrustes-Aligned MPJPE) evaluates position error after optimal rigid alignment. We also compute acceleration error  $e_{acc}$  and velocity error  $e_{vel}$ . From the results, we observe that our per-data fine-tune result provides closer corrected motion than the whole data pretrain step.

In Tab. 1, we also provide simple statistics and a comparison with the previous annotation in MotionPercept. Our annotation is fine-grained, physically aligned, and is normalized to a  $N(0, 1)$  normal distribution. In Fig. 3, we visualize the input motion along with the final fine-tuned corrected motion. Our corrected motion well aligns with physical laws.

## 5 Experiment Results

**Comparison with previous metrics.** We verify the human perceptual and physical alignment of our metric PP-Motion and previous metrics on 2 subsets of MotionPercept: MDM and FLAME. The results are shown in Tab. 3. For human perceptual alignment, we evaluated the “better/worse” classification accuracy. For physical alignment, we evaluated 3 correlation methods (i.e., PLCC [51], SROCC [62], KROCC [41] between PP-Motion results and physical

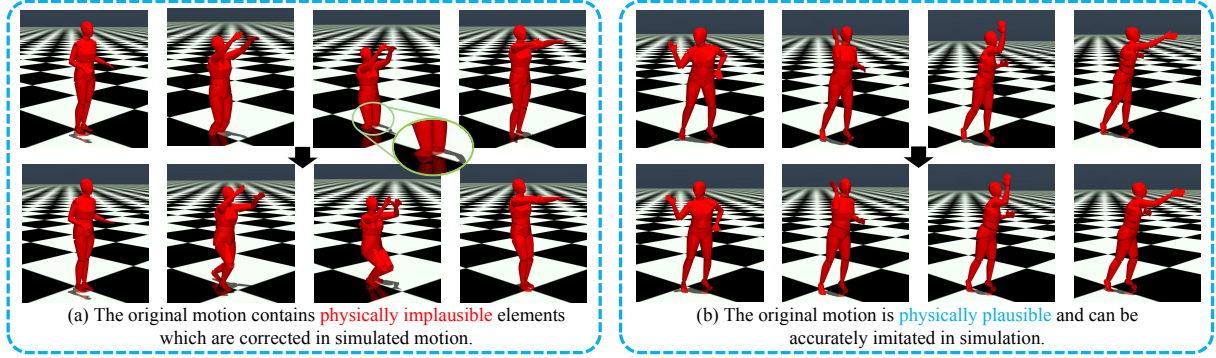


Figure 3: Visualized examples for our dataset labeling. Above: Original motion sequence from the MotionPercept dataset. Below: The motion sequence imitated in the simulator using a per-sequence fine-tuned model.

Table 3: Quantitative comparison of our metric with previous metrics. We report human perceptual accuracy [72] and physical correlation PLCC [51], SROCC [62], and KROCC [41] on 2 datasets, MotionPercept-MDM and MotionPercept-FLAME. Bold numbers indicate the best results.

Metrics	MotionPercept-MDM				MotionPercept-FLAME			
	Accuracy(%) $\uparrow$	PLCC $\uparrow$	SROCC $\uparrow$	KROCC $\uparrow$	Accuracy(%) $\uparrow$	PLCC $\uparrow$	SROCC $\uparrow$	KROCC $\uparrow$
Root AVE [72]	59.47	0.323	0.223	0.150	48.43	0.048	0.135	0.089
Root AE [72]	61.79	0.436	0.412	0.295	59.54	0.135	0.304	0.208
Joint AVE [72]	56.77	0.322	0.239	0.164	44.61	0.072	0.112	0.079
Joint AE [72]	62.73	0.467	0.456	0.327	58.37	0.236	0.377	0.262
PFC [72]	64.79	0.441	0.504	0.364	66.00	0.298	0.451	0.325
Penetration [70]	50.88	0.169	0.082	0.058	56.72	0.229	0.215	0.152
Skating [70]	52.46	0.219	0.132	0.096	56.72	0.092	0.190	0.137
Floating [70]	55.13	0.382	0.318	0.230	55.06	0.478	0.426	0.305
MotionCritic [72]	85.07	0.329	0.316	0.220	67.66	0.152	0.280	0.188
PP-Motion (Ours)	<b>85.18</b>	<b>0.727</b>	<b>0.622</b>	<b>0.461</b>	<b>68.82</b>	<b>0.657</b>	<b>0.660</b>	<b>0.487</b>

Table 4: Pearson’s Correlation Coefficients (PLCC) results on 12 different prompts on HumanAct12 and on the total 40 prompts of UESTC. HumanAct12 and UESTC are 2 subsets of MotionPercept-MDM. Bold numbers indicate the best results, and underline numbers indicate the second best results.

Metrics	HumanAct12												UESTC	Total
	P00	P01	P02	P03	P04	P05	P06	P07	P08	P09	P10	P11		
Root AVE [72]	0.610	0.158	0.129	0.187	0.559	0.559	-0.140	0.050	-0.007	-0.154	0.458	0.181	0.355	0.323
Joint AVE [72]	0.509	0.115	0.177	0.344	0.456	0.576	-0.127	<b>0.278</b>	-0.004	-0.141	0.571	0.018	0.350	0.322
Joint AE [72]	<u>0.738</u>	0.455	<u>0.555</u>	0.220	0.535	0.552	-0.182	-0.010	-0.047	-0.229	<u>0.647</u>	0.147	0.522	<u>0.467</u>
Root AE [72]	0.714	<u>0.564</u>	<u>0.475</u>	0.171	<u>0.568</u>	<u>0.642</u>	-0.242	0.181	-0.019	-0.240	0.629	0.184	0.476	0.436
PFC [72]	0.521	0.099	0.330	0.286	0.478	0.587	0.458	-0.044	-0.012	0.100	0.437	0.255	<u>0.486</u>	0.441
Penetration [70]	0.220	-0.190	-0.137	0.005	0.155	0.275	-0.394	0.173	-0.116	0.059	-0.120	0.219	0.216	0.169
Skating [70]	0.320	-0.257	-0.045	-0.039	0.371	0.201	-0.226	-0.090	-0.077	-0.171	0.078	0.271	0.277	0.219
Floating [70]	0.601	0.341	0.534	<b>0.594</b>	0.344	0.570	<u>0.754</u>	0.021	<u>-0.003</u>	-0.004	<b>0.664</b>	<u>0.425</u>	0.375	0.382
MotionCritic [72]	0.385	0.525	0.438	0.096	0.328	0.334	0.688	-0.284	-0.274	0.163	0.223	0.217	0.302	0.287
PP-Motion (Ours)	<b>0.760</b>	<b>0.983</b>	<b>0.808</b>	<u>0.541</u>	<b>0.699</b>	<b>0.664</b>	<b>0.782</b>	<u>0.272</u>	<b>0.123</b>	<b>0.663</b>	0.568	<b>0.515</b>	<b>0.760</b>	<b>0.727</b>

annotation. The correlation methods are defined in supplementary materials. Our metric is trained only on the MDM training set and directly tested on the MotionPercept-MDM validation and the MotionPercept-FLAME dataset without further finetuning. The results show that our metric outperforms previous works on physical alignment. Notably, on human perceptual alignment, our metric slightly outperforms the baseline method, MotionCritic, which proves that the physical annotation has the potential to improve the human perceptual alignment of our metric. We also compare our metric with existing pose-based metrics (i.e., Root AVE, Root

AE, Joint AVE, Joint AE) and physics-based metrics (i.e., Penetration, Float, Skate). Detailed definitions of these metrics are provided in supplementary materials.

**Per-category physical alignment.** As shown in Tab. 4, Tab. 5, and Tab. 6, we further report the per-category correlation results on 12 different prompts of HumanAct12 [21] and 40 prompts of UESTC [29]. We observe that on most categories, our metric has the best or the second best correlation with the physical annotation among all metrics. This further proves the generalizability of our metric.



**Table 5: Spearman’s Ranking Order Correlation Coefficients (SROCC) on 12 different prompts on HumanAct12 and on UESTC.**

Metrics	HumanAct12												UESTC	Total
	P00	P01	P02	P03	P04	P05	P06	P07	P08	P09	P10	P11		
Root AVE [72]	0.450	0.052	-0.185	0.037	0.372	0.418	-0.280	0.081	0.104	-0.292	0.166	0.262	0.260	0.223
Joint AVE [72]	0.085	0.011	-0.161	0.235	0.040	0.304	-0.073	<b>0.357</b>	-0.117	-0.316	0.338	0.100	0.290	0.239
Joint AE [72]	0.612	0.591	0.578	0.292	0.301	0.485	0.059	-0.028	-0.207	-0.310	0.339	0.203	0.520	0.456
Root AE [72]	<u>0.632</u>	0.537	0.561	0.304	0.509	0.469	-0.055	<u>0.183</u>	-0.184	-0.376	0.269	0.294	0.458	0.412
PFC [72]	<b>0.659</b>	0.404	<u>0.632</u>	0.350	<u>0.533</u>	<b>0.610</b>	0.601	0.004	0.015	-0.197	<b>0.640</b>	<u>0.330</u>	<u>0.541</u>	<u>0.504</u>
Penetration [70]	0.090	-0.394	-0.069	-0.152	0.199	0.032	-0.285	0.127	-0.172	-0.137	-0.048	0.121	0.124	0.082
Skating [70]	0.156	-0.390	0.018	-0.145	0.317	-0.007	-0.057	-0.087	-0.022	-0.276	0.194	0.227	0.173	0.132
Floating [70]	0.330	<b>0.760</b>	0.467	<b>0.569</b>	0.176	<u>0.528</u>	<u>0.716</u>	-0.020	<b>0.109</b>	0.006	0.274	0.236	0.310	0.318
MotionCritic [72]	0.450	0.460	0.400	0.071	0.402	0.373	0.696	-0.357	-0.323	<u>0.052</u>	0.288	0.249	0.299	0.283
PP-Motion (Ours)	0.551	<u>0.593</u>	<b>0.716</b>	<u>0.435</u>	<b>0.596</b>	0.455	<b>0.778</b>	-0.156	-0.013	<b>0.203</b>	<u>0.426</u>	<b>0.545</b>	<b>0.681</b>	<b>0.622</b>

**Table 6: Kendall’s Ranking Order Correlation Coefficients (KROCC) on 12 different prompts on HumanAct12 and on UESTC.**

Metrics	HumanAct12												UESTC	Total
	P00	P01	P02	P03	P04	P05	P06	P07	P08	P09	P10	P11		
Root AVE [72]	0.316	0.043	-0.139	0.032	0.249	0.286	-0.189	0.065	<b>0.075</b>	-0.202	0.102	0.166	0.175	0.150
Joint AVE [72]	0.063	0.014	-0.112	0.157	0.028	0.202	-0.068	<b>0.247</b>	-0.058	-0.222	0.246	0.057	0.199	0.164
Joint AE [72]	0.433	0.425	0.419	0.213	0.206	<u>0.339</u>	0.048	-0.021	-0.146	-0.186	0.236	0.147	0.373	0.327
Root AE [72]	0.457	0.387	0.405	0.217	0.355	0.326	-0.033	<u>0.116</u>	-0.126	-0.244	0.199	0.202	0.326	0.295
PFC [72]	<b>0.482</b>	0.275	0.454	0.230	0.386	0.420	0.405	0.016	0.007	-0.128	<b>0.438</b>	0.230	0.392	0.364
Penetration [70]	0.057	-0.271	-0.049	-0.111	0.136	0.021	-0.203	0.084	-0.128	-0.095	-0.027	0.086	0.088	0.058
Skate [70]	0.107	-0.272	0.015	-0.101	0.225	-0.007	-0.025	-0.077	-0.008	-0.200	0.133	0.165	0.125	0.096
Float [70]	0.236	<b>0.572</b>	0.308	<b>0.409</b>	0.123	<b>0.373</b>	<u>0.521</u>	-0.017	0.072	0.009	0.196	0.161	0.225	0.230
MotionCritic [72]	0.318	0.315	0.278	0.038	0.279	0.245	0.507	-0.232	-0.221	<u>0.038</u>	0.191	0.167	0.208	0.197
PP-Motion (Ours)	0.398	<u>0.431</u>	<b>0.515</b>	<u>0.300</u>	<b>0.435</b>	0.320	<b>0.586</b>	-0.101	-0.015	<b>0.140</b>	<u>0.285</u>	<b>0.349</b>	<b>0.509</b>	<b>0.461</b>

**Table 7: Ablation studies on different loss functions and training strategies.**

Metrics	MotionPercept-MDM				MotionPercept-FLAME			
	Accuracy(%) $\uparrow$	SROCC $\uparrow$	KROCC $\uparrow$	PLCC $\uparrow$	Accuracy(%) $\uparrow$	SROCC $\uparrow$	KROCC $\uparrow$	PLCC $\uparrow$
MotionCritic	85.07	0.3290	0.3160	0.2200	67.66	0.1520	0.2797	0.1875
w/o prompt categorization	<b>85.61</b>	0.5191	0.3791	0.6146	<b>70.98</b>	0.6422	0.4649	0.6347
MSE loss	84.29	0.6000	0.4446	0.6357	69.48	0.6059	0.4446	0.5797
PP-Motion (Ours)	85.18	<b>0.6223</b>	<b>0.4612</b>	<b>0.7268</b>	68.82	<b>0.6598</b>	<b>0.4873</b>	<b>0.6567</b>

**Table 8: Comparison of MDM model performance before and after finetuning with our metric. PP-Motion is the average predicted score of our metric. Mean MPJPE is the mean per-joint position error between simulated motion and ground truth motion.**

	PP-Motion $\uparrow$	Mean MPJPE $\downarrow$
Before Fine-tuning	-0.09	76.06
Fine-tune 100 steps	0.61	63.33

**Ablation studies.** We perform ablation studies to validate our metric designs, with results presented in Tab. 7. First, in “w/o prompt categorization”, we examine the impact of a prompt-categorized training strategy by training directly on the MotionPercept-MDM dataset without categorizing by prompt labels. The prompt-categorized training approach is detailed in supplementary materials. The results demonstrate that computing PLCC loss within the same-label motion groups achieves better results in physical correlation metrics. Second, in “MSE loss”, we train the metric with MSE loss, which calculates the  $l_2$  distance between the predicted score and GT annotation. The results show that replacing PLCC loss with conventional MSE optimization leads to noticeable

degradation in both physical plausibility assessment and human evaluation metrics.

**Improving motion generation with PP-Motion.** We try to improve the motion generation method MDM [67] to further verify our metric’s physical fidelity alignment. We fine-tune the MDM network using our metric, along with critic loss and KL loss from [72]. The critic loss for the input pose  $x'_0$  is defined as:

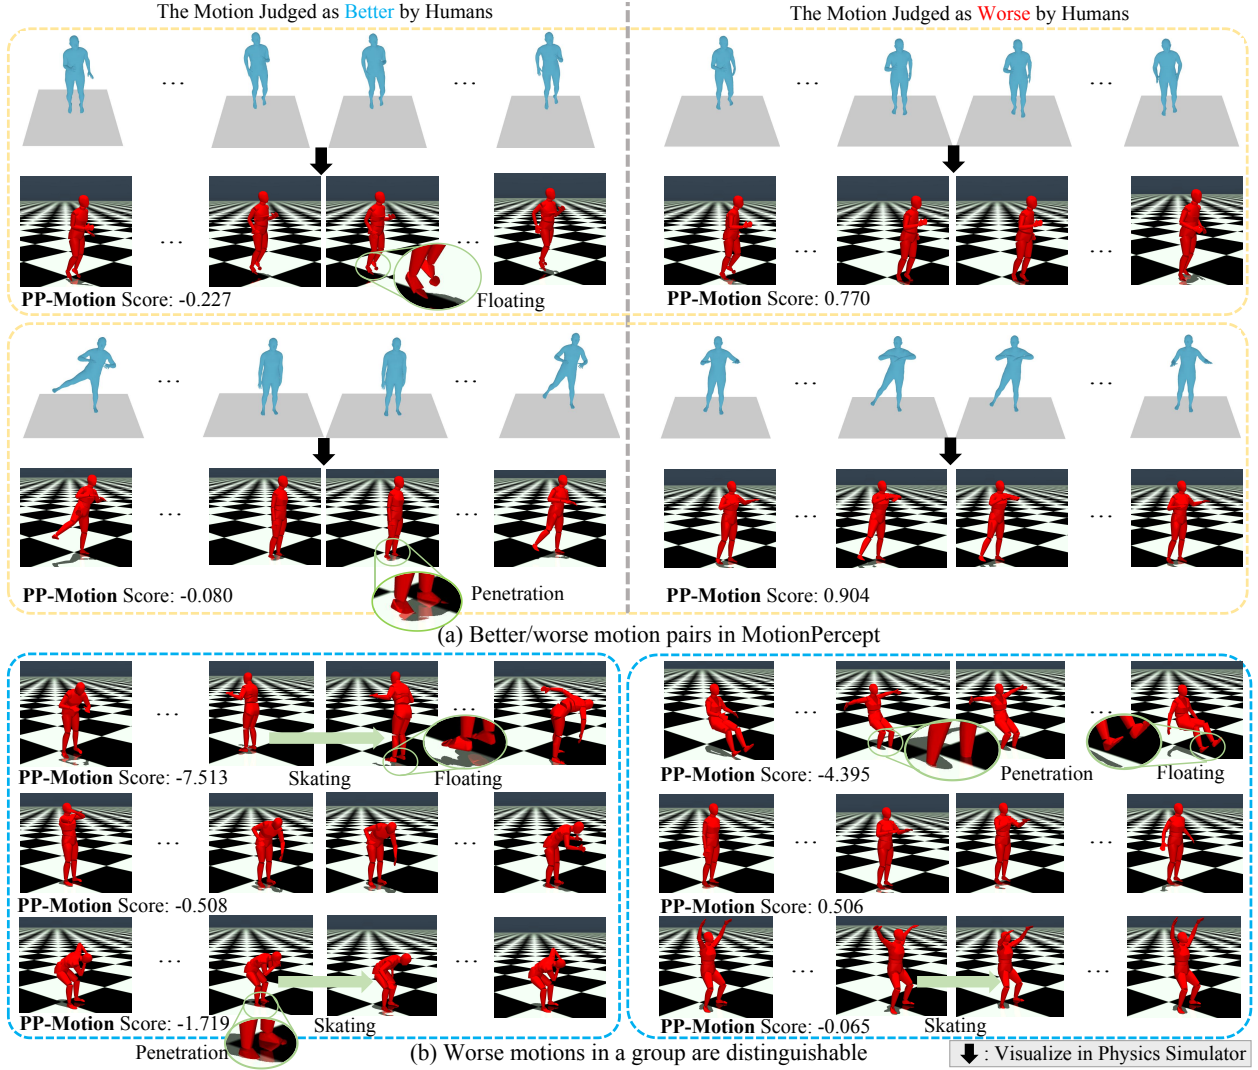
$$\mathcal{L}_{\text{Critic}} = \mathbb{E} \left[ -\sigma \left( \tau - F \left( x'_0 \right) \right) \right], \quad (9)$$

where  $\tau$  is the threshold of sigmoid function  $\sigma(\cdot)$ , and  $\mathbb{E}$  is the expectation on the whole dataset. The KL loss is defined as

$$\mathcal{L}_{\text{KL}} = \mathbb{E} \left[ D_{\text{KL}} \left( p(x'_0) \mid p(\tilde{x}'_0) \right) \right], \quad (10)$$

where  $D_{\text{KL}}$  is KL divergence, and  $\tilde{x}'_0$  is the pervious iteration of  $x'_0$ .

We fine-tune the MDM model for 100 steps and generate 120 motion sequences each using the MDM baseline model and the fine-tuned model. Then we fine-tune each motion on PHC, following the procedure described in Sec. 4, and calculate the mean MPJPE between the motion simulated in Isaac and the ground truth motion generated by MDM. The results reported in Tab. 8 show that our metric can improve physical alignment in motion generation.



**Figure 4:** (a) A MotionPercept better/worse data pair: the motion on the left, annotated as ‘better’ and visually superior, exhibits physics issues (e.g. floating, skating). The motion on the right, annotated as ‘worse’ and visually inferior, shows greater physical plausibility. (b) Three MotionPercept samples in a group, all annotated as ‘worse’ by humans, reveal different physical characteristics in the simulator. Our PP-Motion scores (higher means better) successfully capture these physical distinctions.

**Visualization.** Fig. 4 (a) shows a better/worse data pair sampled from the MotionPercept dataset. The motion on the left (annotated as ‘better’ and visually superior) exhibits physics issues (e.g. floating, skating) in the simulator, while the motion on the right (annotated as ‘worse’ and visually inferior) demonstrates greater physical plausibility when simulated. Fig. 4 (b) shows three MotionPercept samples in a group that are all annotated as ‘worse’, but reveal different physical characteristics in the simulator. Our PP-Motion scores successfully capture these physical distinctions. Note that for our PP-Motion score, the higher reveals the better.

## 6 Conclusion

In this work, we address the challenges in evaluating the fidelity of generated human motions by bridging the gap between human

perception and physical feasibility. We introduce a novel physical labeling method that computes the minimum adjustments needed for a motion to adhere to physical laws, thereby producing fine-grained, continuous physical alignment annotations as objective ground truth. Our framework leverages Pearson’s correlation loss to capture the underlying physical priors, while integrating a human-based perceptual fidelity loss to ensure that the evaluation metric reflects both physical fidelity and human perception. Experimental results validate that our metric not only complies with physical laws but also demonstrates superior alignment with human perception compared to previous approaches.



## Acknowledgments

This work is supported by the National Key R&D Program of China under Grant No.2024QY1400, and the National Natural Science Foundation of China No. 62425604. This work is also supported by Tsinghua University Initiative Scientific Research Program and the Institute for Guo Qiang at Tsinghua University.

## References

- [1] Chaitanya Ahuja and Louis-Philippe Morency. 2019. Language2pose: Natural language grounded pose forecasting. In *2019 International conference on 3D vision (3DV)*. IEEE, 719–728.
- [2] Tenglong Ao, Qingzhe Gao, Yuke Lou, Baoquan Chen, and Libin Liu. 2022. Rhythmic gesticulator: Rhythm-aware co-speech gesture synthesis with hierarchical neural embeddings. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–19.
- [3] Joao Pedro Araújo, Jiaman Li, Karthik Vetrivel, Rishi Agarwal, Jiajun Wu, Deepak Gopinath, Alexander William Clegg, and Karen Liu. 2023. Circle: Capture in rich contextual environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 21211–21221.
- [4] Nikos Athanasiou, Mathis Petrovich, Michael J Black, and Gül Varol. 2022. Teach: Temporal action composition for 3d humans. In *2022 International Conference on 3D Vision (3DV)*. IEEE, 414–423.
- [5] Nikos Athanasiou, Mathis Petrovich, Michael J. Black, and Gül Varol. 2022. TEACH: Temporal Action Compositions for 3D Humans. In *International Conference on 3D Vision (3DV)*.
- [6] German Barquero, Sergio Escalera, and Cristina Palmero. 2024. Seamless Human Motion Composition with Blended Positional Encodings. (2024).
- [7] Uttaran Bhattacharya, Elizabeth Childs, Nicholas Rewkowski, and Dinesh Manocha. 2021. Speech2affectivegestures: Synthesizing co-speech gestures with generative adversarial affective expression learning. In *Proceedings of the 29th ACM International Conference on Multimedia*. 2027–2036.
- [8] Zhe Cao, Hang Gao, Karttikeya Mangalam, Qi-Zhi Cai, Minh Vo, and Jitendra Malik. 2020. Long-term human motion prediction with scene context. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*. Springer, 387–404.
- [9] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. 2023. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 18000–18010.
- [10] Enric Corona, Albert Pumarola, Guillem Alenya, and Francesc Moreno-Noguer. 2020. Context-aware human motion prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 6992–7001.
- [11] Wenxun Dai, Ling-Hao Chen, Jingbo Wang, Jinpeng Liu, Bo Dai, and Yansong Tang. 2025. Motionlcm: Real-time controllable motion generation via latent consistency model. In *ECCV*. 390–408.
- [12] Xuehao Gao, Yang Yang, Zhenyu Xie, Shaoyi Du, Zhongqian Sun, and Yang Wu. 2024. Guess: Gradually enriching synthesis for text-driven human motion generation. *IEEE Transactions on Visualization and Computer Graphics* 30, 12 (2024), 7518–7530.
- [13] Anindita Ghosh, Noshaba Cheema, Cennet Oguz, Christian Theobalt, and Philipp Slusallek. 2021. Synthesis of compositional animations from textual descriptions. In *Proceedings of the IEEE/CVF international conference on computer vision*. 1396–1406.
- [14] Anindita Ghosh, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Philipp Slusallek. 2023. IMoS: Intent-Driven Full-Body Motion Synthesis for Human-Object Interactions. In *Computer Graphics Forum*, Vol. 42. Wiley Online Library, 1–12.
- [15] Shiry Ginosar, Amir Bar, Gefen Kohavi, Caroline Chan, Andrew Owens, and Jitendra Malik. 2019. Learning individual styles of conversational gesture. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 3497–3506.
- [16] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014).
- [17] Anand Gopalakrishnan, Ankur Mali, Dan Kifer, Lee Giles, and Alexander G Ororbia. 2019. A neural temporal model for human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12116–12125.
- [18] Chuan Guo, Yuxuan Mu, Muhammad Gohar Javed, Sen Wang, and Li Cheng. 2024. Momask: Generative masked modeling of 3d human motions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1900–1910.
- [19] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. 2022. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5152–5161.
- [20] Chuan Guo, Xinxin Zuo, Sen Wang, and Li Cheng. 2022. Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *European Conference on Computer Vision*. Springer, 580–597.
- [21] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. 2020. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*. 2021–2029.
- [22] Ikhsanul Habibie, Weipeng Xu, Dushyant Mehta, Lingjie Liu, Hans-Peter Seidel, Gerard Pons-Moll, Mohamed Elgharib, and Christian Theobalt. 2021. Learning speech-driven 3d conversational gestures from video. In *Proceedings of the 21st ACM International Conference on Intelligent Virtual Agents*. 101–108.
- [23] Gao Han, Mingjiang Liang, Jinglei Tang, Yongkang Cheng, Wei Liu, and Shaoli Huang. 2024. Reindiffuse: Crafting physically plausible motions with reinforced diffusion model. *arXiv preprint arXiv:2410.07296* (2024).
- [24] Mohamed Hassan, Duygu Ceylan, Ruben Villegas, Jun Saito, Jimei Yang, Yi Zhou, and Michael J Black. 2021. Stochastic scene-aware motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 11374–11384.
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems* 33 (2020), 6840–6851.
- [26] Nhat M Hoang, Kehong Gong, Chuan Guo, and Michael Bi Mi. 2024. Motionmix: Weakly-supervised diffusion for controllable motion generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 2157–2165.
- [27] Siyuan Huang, Zan Wang, Puhao Li, Baoxiong Jia, Tengyu Liu, Yixin Zhu, Wei Liang, and Song-Chun Zhu. 2023. Diffusion-based generation, optimization, and planning in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16750–16761.
- [28] Yiheng Huang, Hui Yang, Chuanchen Luo, Yuxi Wang, Shibiao Xu, Zhaoxiang Zhang, Man Zhang, and Junran Peng. 2024. Stablemofusion: Towards robust and efficient diffusion-based motion generation framework. In *Proceedings of the 32nd ACM International Conference on Multimedia*. 224–232.
- [29] Yanli Ji, Feixiang Xu, Yang Yang, Fumin Shen, Heng Tao Shen, and Wei-Shi Zheng. 2018. A large-scale RGB-D database for arbitrary-view human action recognition. In *Proceedings of the 26th ACM international Conference on Multimedia*. 1510–1518.
- [30] Biao Jiang, Xin Chen, Wen Liu, Jingyi Yu, Gang Yu, and Tao Chen. 2023. Motiongpt: Human motion as a foreign language. *Advances in Neural Information Processing Systems* 36 (2023), 20067–20079.
- [31] Korrae Karunratanakul, Konpat Preechakul, Emre Aksan, Thabo Beeler, Supasorn Suwajanakorn, and Siyu Tang. 2024. Optimizing diffusion noise can serve as universal motion priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1334–1345.
- [32] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. 2023. Flame: Free-form language-based motion synthesis & editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 8255–8263.
- [33] Diederik P Kingma, Max Welling, et al. 2013. Auto-encoding variational bayes.
- [34] Taras Kucherenko, Dai Hasegawa, Gustav Eje Henter, Naoshi Kaneko, and Hedvig Kjellström. 2019. Analyzing input and output representations for speech-driven gesture generation. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*. 97–104.
- [35] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.
- [36] Jing Li, Di Kang, Wenjie Pei, Xuefei Zhe, Ying Zhang, Zhenyu He, and Linchao Bao. 2021. Audio2gestures: Generating diverse gestures from speech audio with conditional variational autoencoders. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 11293–11302.
- [37] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. 2021. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF international conference on computer vision*. 13401–13412.
- [38] Ronghui Li, YuXiang Zhang, Yachao Zhang, Hongwen Zhang, Jie Guo, Yan Zhang, Yebin Liu, and Xiu Li. 2024. Lodge: A coarse to fine diffusion network for long dance generation guided by the characteristic dance primitives. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1524–1534.
- [39] Ronghui Li, Junfan Zhao, Yachao Zhang, Mingyang Su, Zeping Ren, Han Zhang, Yansong Tang, and Xiu Li. 2023. Finedance: A fine-grained choreography dataset for 3d full body dance generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10234–10243.
- [40] Zhuo Li, Mingshuang Luo, Ruibing Hou, Xin Zhao, Hao Liu, Hong Chang, Zimo Liu, and Chen Li. 2024. Morph: A Motion-free Physics Optimization Framework for Human Motion Generation. *arXiv preprint arXiv:2411.14951* (2024).
- [41] Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017).
- [42] Nathan Louis, Mahzad Khoshhessan, and Jason J Corso. 2025. Measuring Physical Plausibility of 3D Human Poses Using Physics Simulation. *arXiv preprint arXiv:2502.04483* (2025).
- [43] Tianyu Luan, Zhongpai Gao, Luyuan Xie, Abhishek Sharma, Hao Ding, Benjamin Planche, Meng Zheng, Ange Lou, Terrence Chen, Junsong Yuan, et al. 2024. Divide and Fuse: Body Part Mesh Recovery from Partially Visible Human Images.

- In *European Conference on Computer Vision*. Springer, 350–367.
- [44] Tianyu Luan, Yali Wang, Junhao Zhang, Zhe Wang, Zhipeng Zhou, and Yu Qiao. 2021. Pc-hmr: Pose calibration for 3d human mesh recovery from 2d images/videos. In *AAAI*. 2269–2276.
  - [45] Tianyu Luan, Yuanhao Zhai, Jingjing Meng, Zhong Li, Zhang Chen, Yi Xu, and Junsong Yuan. 2023. High Fidelity 3D Hand Shape Reconstruction via Scalable Graph Frequency Decomposition. In *CVPR*. 16795–16804.
  - [46] Tianyu Luan, Yuanhao Zhai, Jingjing Meng, Zhong Li, Zhang Chen, Yi Xu, and Junsong Yuan. 2025. Scalable High-Fidelity 3D Hand Shape Reconstruction Via Graph-Image Frequency Mapping and Graph Frequency Decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025).
  - [47] Thomas Lucas, Fabien Baradel, Philippe Weinzaepfel, and Grégory Rogez. 2022. Posegpt: Quantization-based 3d human motion generation and forecasting. In *European Conference on Computer Vision*. Springer, 417–435.
  - [48] Zhengyi Luo, Jinkun Cao, Alexander W. Winkler, Kris Kitani, and Weipeng Xu. 2023. Perpetual Humanoid Control for Real-time Simulated Avatars. In *International Conference on Computer Vision (ICCV)*.
  - [49] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, et al. 2021. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470* (2021).
  - [50] Wei Mao, Richard I Hartley, Mathieu Salzmann, et al. 2022. Contact-aware human motion forecasting. *Advances in Neural Information Processing Systems* 35 (2022), 7356–7367.
  - [51] Karl Pearson. 1920. Notes on the history of correlation. *Biometrika* (1920), 25–45.
  - [52] Mathis Petrovich, Michael J Black, and Gül Varol. 2022. Temos: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision*. Springer, 480–497.
  - [53] Mathis Petrovich, Michael J. Black, and Gül Varol. 2022. TEMOS: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision (ECCV)*.
  - [54] Ekkasit Pinyoanuntapong, Pu Wang, Minwoo Lee, and Chen Chen. 2024. Mmm: Generative masked motion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1546–1555.
  - [55] Shenhan Qian, Zhi Tu, Yihao Zhi, Wen Liu, and Shenghua Gao. 2021. Speech drives templates: Co-speech gesture synthesis with learned templates. In *Proceedings of the IEEE/CVF international conference on computer vision*. 11077–11086.
  - [56] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. (2018).
  - [57] Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J Guibas. 2021. Humor: 3d human motion model for robust pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*. 11488–11499.
  - [58] Davis Rempe, Leonidas J Guibas, Aaron Hertzmann, Bryan Russell, Ruben Villegas, and Jimei Yang. 2020. Contact and human dynamics from monocular video. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V* 16. Springer, 71–87.
  - [59] Danilo Rezende and Shakir Mohamed. 2015. Variational inference with normalizing flows. In *International conference on machine learning*. PMLR, 1530–1538.
  - [60] Yoni Shafir, Guy Tevet, Roy Kapon, and Amit Haim Bermano. 2024. Human Motion Diffusion as a Generative Prior. In *The Twelfth International Conference on Learning Representations*.
  - [61] Li Siyao, Weijiang Yu, Tianpei Gu, Chunze Lin, Quan Wang, Chen Qian, Chen Change Loy, and Ziwei Liu. 2022. Bailando: 3d dance generation by actor-critic gpt with choreographic memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11050–11059.
  - [62] Charles Spearman. 1910. Correlation calculated from faulty data. *British journal of psychology* (1910), 271.
  - [63] Haowen Sun, Ruikun Zheng, Haibin Huang, Chongyang Ma, Hui Huang, and Ruizhen Hu. 2024. LGTM: Local-to-Global Text-Driven Human Motion Diffusion Model. In *ACM SIGGRAPH 2024 Conference Papers*. 1–9.
  - [64] Omid Taheri, Vasileios Choutas, Michael J Black, and Dimitrios Tzionas. 2022. Goal: Generating 4d whole-body motion for hand-object grasping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13263–13273.
  - [65] Taoran Tang, Jia Jia, and Hanyang Mao. 2018. Dance with melody: An lstm-autoencoder approach to music-oriented dance synthesis. In *Proceedings of the 26th ACM international conference on Multimedia*. 1598–1606.
  - [66] Guy Tevet, Brian Gordon, Amir Hertz, Amit H Bermano, and Daniel Cohen-Or. 2022. Motionclip: Exposing human motion generation to clip space. In *European Conference on Computer Vision*. Springer, 358–374.
  - [67] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. 2022. Human motion diffusion model. *arXiv preprint arXiv:2209.14916* (2022).
  - [68] Jonathan Tseng, Rodrigo Castellon, and Karen Liu. 2023. Edge: Editable dance generation from music. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 448–458.
  - [69] Kengo Uchida, Takashi Shibuya, Yuhta Takida, Naoki Murata, Julian Tanke, Shusuke Takahashi, and Yuki Mitsufuji. 2024. Mola: Motion generation and editing with latent diffusion enhanced by adversarial training. *arXiv preprint arXiv:2406.01867* (2024).
  - [70] Nicolas Ugrinovic, Boxiao Pan, Georgios Pavlakos, Despoina Paschalidou, Bokui Shen, Jordi Sanchez-Riera, Francesc Moreno-Noguer, and Leonidas Guibas. 2024. MultiPhys: multi-person physics-aware 3D motion estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2331–2340.
  - [71] Weilin Wan, Zhiyang Dou, Taku Komura, Wenping Wang, Dinesh Jayaraman, and Lingjie Liu. 2024. Tlcontrol: Trajectory and language control for human motion synthesis. In *European Conference on Computer Vision*. Springer, 37–54.
  - [72] Haoru Wang, Wentao Zhu, Luyi Miao, Yishu Xu, Feng Gao, Qi Tian, and Yizhou Wang. 2025. Aligning Motion Generation with Human Perceptions. In *International Conference on Learning Representations (ICLR)*.
  - [73] Jiajun Wang, Huazhe Xu, Jingwei Xu, Sifei Liu, and Xiaolong Wang. 2021. Synthesizing long-term 3d human motion and interaction in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9401–9411.
  - [74] Yuan Wang, Zhao Wang, Junhao Gong, Di Huang, Tong He, Wanli Ouyang, Jile Jiao, Xuetao Feng, Qi Dou, Shixiang Tang, et al. 2024. Holistic-motion2d: Scalable whole-body human motion generation in 2d space. *arXiv preprint arXiv:2406.11253* (2024).
  - [75] Zan Wang, Yixin Chen, Tengyu Liu, Yixin Zhu, Wei Liang, and Siyuan Huang. 2022. Humanise: Language-conditioned human motion generation in 3d scenes. *Advances in Neural Information Processing Systems* 35 (2022), 14959–14971.
  - [76] Yan Wu, Jiahao Wang, Yan Zhang, Siwei Zhang, Otmar Hilliges, Fisher Yu, and Siyu Tang. 2022. Saga: Stochastic whole-body grasping with contact. In *European Conference on Computer Vision*. Springer, 257–274.
  - [77] Yiming Xie, Varun Jampani, Lei Zhong, Deqing Sun, and Huaizu Jiang. 2024. OmniControl: Control Any Joint at Any Time for Human Motion Generation. In *The Twelfth International Conference on Learning Representations*.
  - [78] Youngwoo Yoon, Bok Cha, Joo-Haeng Lee, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. 2020. Speech gesture generation from the trimodal context of text, audio, and speaker identity. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–16.
  - [79] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. 2023. Physdiff: Physics-guided human motion diffusion model. In *Proceedings of the IEEE/CVF international conference on computer vision*. 16010–16021.
  - [80] Yuanhao Zhai, Mingzhen Huang, Tianyu Luan, Lu Dong, Ifeoma Nwogu, Siwei Lyu, David Doermann, and Junsong Yuan. 2023. Language-guided human motion synthesis with atomic actions. In *Proceedings of the 31st ACM International Conference on Multimedia*. 5262–5271.
  - [81] Junhao Zhang, Yali Wang, Zhipeng Zhou, Tianyu Luan, Zhe Wang, and Yu Qiao. 2021. Learning dynamical human-joint affinity for 3d pose estimation in videos. *IEEE TIP* (2021).
  - [82] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Yong Zhang, Hongwei Zhao, Hongtao Lu, Xi Shen, and Ying Shan. 2023. Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 14730–14740.
  - [83] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. 2024. Motiondiffuse: Text-driven human motion generation with diffusion model. *IEEE transactions on pattern analysis and machine intelligence* 46, 6 (2024), 4115–4128.
  - [84] Mingyuan Zhang, Huirong Li, Zhongang Cai, Jiawei Ren, Lei Yang, and Ziwei Liu. 2023. FineMoGen: Fine-Grained Spatio-Temporal Motion Generation and Editing. *NeurIPS* (2023).
  - [85] Zeyu Zhang, Yiran Wang, Wei Mao, Danning Li, Rui Zhao, Biao Wu, Zirui Song, Bohan Zhuang, Ian Reid, and Richard Hartley. 2025. Motion Anything: Any to Motion Generation. *arXiv preprint arXiv:2503.06955* (2025).
  - [86] Zixiang Zhou and Baoyuan Wang. 2023. Ude: A unified driving engine for human motion generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5632–5641.
  - [87] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang. 2023. Learning human motion representations: A unified perspective. In *Proc. Int. Conf. Comput. Vis.* 15085–15099.
  - [88] Wentao Zhu, Xiaoxuan Ma, Dongwoo Ro, Hai Ci, Jinlu Zhang, Jiaxin Shi, Feng Gao, Qi Tian, and Yizhou Wang. 2023. Human motion generation: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46, 4 (2023), 2430–2449.