


# MacDiff: Unified Skeleton Modeling with Masked Conditional Diffusion

Lehong Wu<sup>1,2</sup>, Lilang Lin<sup>1</sup>, Jiahang Zhang<sup>1</sup>, Yiyang Ma<sup>1</sup>, and Jiaying Liu<sup>1\*</sup> 

<sup>1</sup> Wangxuan Institute of Computer Technology, Peking University

<sup>2</sup> School of Electronics Engineering and Computer Science, Peking University  
aladonwlh@stu.pku.edu.cn

{linlilang, zjh2020, myy12769, liujiaying}@pku.edu.cn

**Abstract.** Self-supervised learning has proved effective for skeleton-based human action understanding. However, previous works either rely on contrastive learning that suffers false negative problems or are based on reconstruction that learns too much unessential low-level clues, leading to limited representations for downstream tasks. Recently, great advances have been made in generative learning, which is naturally a challenging yet meaningful pretext task to model the general underlying data distributions. However, the representation learning capacity of generative models is under-explored, especially for the skeletons with spacial sparsity and temporal redundancy. To this end, we propose Masked Conditional Diffusion (MacDiff) as a unified framework for human skeleton modeling. For the first time, we leverage diffusion models as effective skeleton representation learners. Specifically, we train a diffusion decoder conditioned on the representations extracted by a semantic encoder. Random masking is applied to encoder inputs to introduce an information bottleneck and remove redundancy of skeletons. Furthermore, we theoretically demonstrate that our generative objective involves the contrastive learning objective which aligns the masked and noisy views. Meanwhile, it also enforces the representation to complement for the noisy view, leading to better generalization performance. MacDiff achieves state-of-the-art performance on representation learning benchmarks while maintaining the competence for generative tasks. Moreover, we leverage the diffusion model for data augmentation, significantly enhancing the fine-tuning performance in scenarios with scarce labeled data. Our project is available at <https://github.com/LehongWu/MacDiff>.

**Keywords:** Self-supervised learning · Unified skeleton modeling · Diffusion model

## 1 Introduction

Human action understanding has been a crucial problem in computer vision. Skeletons use 3D coordinates to represent human joints, providing a lightweight,

---

\* Corresponding author.

compact, and privacy-preserving data modality of human representation. Owing to these advantages, skeletons are widely used for human action analysis in real-world applications such as human-robotics interaction [24], autonomous driving [2] and video surveillance [13]. Since well-annotated data are expensive to obtain, self-supervised methods have been proposed to extract meaningful representations from unlabeled skeletons.

Prevalent high-performing self-supervised methods for skeletons mainly contain two paradigms, *i.e.*, contrastive learning-based methods and reconstruction-based methods. Contrastive learning (CL) trains the model to capture shared information between two augmented views of the same sample. Numerous works [8, 27, 40, 48, 49] focus on designing skeleton-specific augmentations to benefit CL performance. Reconstruction-based methods [17, 32, 50, 55] carefully design reconstruction targets that prompt the model to capture the spatial-temporal correlations of skeletons. However, existing self-supervised methods all focus on specific semantics required by the pre-defined tasks. For example, contrastive learning is proved to only learn the discriminative information for positive/negative pairs [53], which limits its generalization ability and makes it sensitive to the augmentation design. Reconstruction-based methods overly focus on full signal reconstruction, causing the representations to contain too much low-level information irrelevant for high-level action understanding.

Generative models, aiming to approximate the real-world data distribution, naturally form a more general self-supervised learning target than carefully designed tasks. Thus, introducing generative models in self-supervised learning would force the model extract better representations with richer semantics. Autoencoder-based generative models like VAEs [22, 37] rely on the information bottleneck to obtain a compact meaningful latent representation, resulting in a trade-off between discriminability and generation authenticity. More recently, diffusion models [11, 18, 41, 42] have shown remarkable generative capacities. However, diffusion models which directly predict the noise contained in the noisy states do not explicitly learn a meaningful latent representation tailored for discriminative tasks. Moreover, due to the spatial sparsity and temporal redundancy of skeletons, diffusion models on skeletons maintain too much irrelevant low-level appearance information, which fails to construct a tight bottleneck and obtain meaningful representation. Therefore, efforts should be made to mine the potential of such powerful diffusion models to obtain more powerful representations.

To tackle these challenges, we propose Masked Conditional Diffusion (MacDiff) which, for the first time, tames diffusion for both skeleton representation learning and skeleton generation. Specifically, we train a semantic encoder to guide a diffusion decoder. The encoder-decoder design serves to disentangle the high-level representation learning with low-level generative training. Meanwhile, this architecture mitigates the conflict between discrimination and generation authenticity, allowing for a more flexible design of the encoder. Considering the spatial-temporal correlation of skeletons, a tighter information bottleneck should be imposed in addition to restricting the dimensions of representations. To this

end, we propose to apply random masking on patchified skeletons with a high masking ratio to constrain the representation’s dimension.

We conduct theoretical analyses on MacDiff from a mutual information perspective to prove the effectiveness of our framework. We demonstrate that the training of MacDiff is equivalent to a combination of the contrastive learning and reconstruction objective. The contrastive learning objective, conducted on the masked and noisy views of skeletons, enforces the representation to capture shared information between different views, while the very small portion of reconstruction objective enriches the representation by complementing the missing information in the noisy view, including more downstream task-relevant information in the representation.

We provide thorough experiments and results on NTU RGB+D [29, 43] and PKUMMD [28] datasets to demonstrate the effectiveness and versatility of our method. MacDiff achieves remarkable results on self-supervised learning benchmarks. In addition, we utilize the pre-trained diffusion for generating label-preserving training data in scenarios with limited training data, which brings significant performance gain in semi-supervised protocols. Our contributions can be summarized as follows:

- We propose Masked Conditional Diffusion (MacDiff), a unified framework for human skeleton modeling. A semantic encoder is employed, learning high-level compact representations, to assist the conditioned generative learning of the diffusion decoder. By virtue of this, our model learns powerful representations for both discriminative and generative downstream tasks.
- We theoretically demonstrate that the generative objective of MacDiff involves both diffusion learning and contrastive learning that aligns the representation of the masked view with the noisy view. Moreover, MacDiff is capable of preserving more semantics in the learned representation, leading to better downstream performance than contrastive-only paradigms.
- MacDiff achieves state-of-the-art performance on three large-scale benchmarks. Remarkably, we leverage diffusion-based data augmentation for encoder fine-tuning, significantly improving the action recognition performance with scarce labeled data.

## 2 Related Work

**Self-Supervised Learning for Skeletons.** Self-supervised learning aims to extract meaningful representations from unlabeled data to facilitate downstream tasks. Prevalent methods for skeleton representation learning can be divided into contrastive methods and reconstruction methods. Contrastive learning (CL) extracts meaningful representations by discriminating positive/negative sample pairs from different augmented views [4, 6]. To leverage CL for skeletons, numerous works [27, 40, 48, 49] focus on developing skeleton-specific data augmentations. Other works extract shared information between different skeleton modalities [25, 33]. Most of these methods use RNNs or GCNs as backbone. Among reconstruction methods, LongT GAN [62] and P&C [23] design

reconstruction tasks on autoencoders to learn compressed representations. GL-Transformer [21] designs prediction tasks with Transformer backbone to capture spatial-temporal dynamics of skeletons. Some works combine reconstruction or prediction tasks with contrastive learning, including [7, 26, 58]. More recently, Masked Autoencoders (MAEs) are introduced to skeletons by SkeletonMAE [55] and MAMP [32], which achieve remarkable performance by modeling spatial-temporal correlations of skeletons with Transformers.

**Diffusion Models.** Diffusion models are a type of generative model that gradually maps a noise prior to a target distribution. They have demonstrated superior capacity in image generation [11, 18, 41, 42] and a variety of applications [31, 34]. Therefore, diffusion models have been widely used for text-guided human motion generation [5, 16, 47, 59], motion prediction [3] and anomaly detection [13]. These efforts adapt image diffusion models to skeletons with modifications including specialized loss functions for skeletons and leveraging Transformer architecture.


Among a few works that pioneer in leveraging diffusion representations, most works directly use intermediate representations of pre-trained diffusion models or fine-tune the model. While proved effective for dense prediction tasks like segmentation [36] and dense matching [35], these representations are not compact enough compared with other self-supervised methods. Another paradigm [61] is to distill information from fixed diffusion models. More recently, efforts that jointly train a separate encoder to guide diffusion yield promising results [20, 39]. In this paper, we propose the first framework that tailors diffusion models for skeleton representation learning, to the best of our knowledge. Besides, we provide theoretical analysis for the effectiveness of the proposed framework.

**Unified Skeleton Modeling.** Recent advancements in Natural Language Processing have demonstrated the potential for building models that unify multiple tasks. For a general-purpose human representation learning, several works [7, 26, 58] leverage multiple tasks as self-supervised training. MotionBERT [63] learns unified 2D skeleton representation by training on 2D-to-3D lifting task. UniHCP [10] provides a unified model for several human-centric tasks, *e.g.*, pose estimation, ReID and pedestrian detection. UPS [14] forms skeletons and action labels as language tokens. Skeleton-in-Context [54] leverages in-context learning to unify multiple estimation and prediction tasks. However, most unified methods are restrained to the specific tasks they are trained on. GFPose [9] utilizes score-based generative model to learn a unified human pose prior but only focus on single-frame skeletons, limiting its applications. In this paper, we explore the capability of a versatile generative model, diffusion model, to unify skeleton representation learning with generation.

### 3 Method

#### 3.1 Diffusion Models Preliminary

Diffusion models are a family of generative models that learn the target distribution by performing data denoising. Denoising Diffusion Probabilistic Model



**Fig. 1:** The overview of the proposed method. We train a diffusion decoder conditioned on the representations extracted by a semantic encoder. In the above stream, we embed the input skeletons into tokens and employ random masking. The global representation is obtained by pooling the local representations extracted by the semantic encoder. In the below stream, we train a conditional diffusion model. We sample the noisy skeleton  $\mathbf{x}_t$  following the diffusion process  $q(\mathbf{x}_t|\mathbf{x}_0)$ . The diffusion decoder predicts the noise  $\epsilon$  from  $\mathbf{x}_t$  guided by the learned representation  $\mathbf{z}$ . The pre-trained encoder can be utilized independently in downstream discriminative tasks.

(DDPM) [18] employs a forward (diffusion) process that sequentially corrupts the data distribution  $q(x_0)$  to the standard Gaussian distribution  $\mathcal{N}(\mathbf{0}, \mathbf{I})$  with the conditional distribution  $q(x_t|x_{t-1})$ . The noise level of the timesteps  $t$  is defined by a fixed increasing variance schedule  $\{\beta_t\}_{t=0}^T$ . This process can be denoted as:

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}). \quad (1)$$

With a given  $t$ , the forward process allows sampling  $\mathbf{x}_t$  directly from  $\mathbf{x}_0$  with  $q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I})$ , where  $\alpha_t = 1 - \beta_t$  and  $\bar{\alpha}_t = \prod_{i=0}^t \alpha_i$ .

The reverse process is defined as another Markov Chain parameterized by  $\theta$ , which maps the standard Gaussian distribution to the distribution of clean data by gradually denoising. Each step is a Gaussian distribution with a predicted mean  $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)$  and covariance matrix  $\sigma_t^2\mathbf{I}$ :

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \sigma_t^2\mathbf{I}). \quad (2)$$

The final training objective is derived from optimizing the variational bound on  $\mathbb{E}[-\log p_\theta(\mathbf{x}_0)]$  [18], where  $\gamma_{1:T}$  are positive coefficients depending on  $\alpha_{1:T}$ :

$$\mathcal{L}(\theta) = \mathbb{E}_{\mathbf{x}_0, t, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\gamma_t \|\epsilon - \boldsymbol{\epsilon}_\theta(\sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t)\|^2]. \quad (3)$$

Other implementations beyond  $\epsilon$ -prediction include  $\mathbf{x}_{t-1}$ - and  $\mathbf{x}_0$ -prediction, which are all equivalent to predicting  $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)$  in mathematical formulation.

For sampling, we adopt Denoising Diffusion Implicit Model (DDIM) [45], which shares the same training objective with DDPM but defines a non-Markov diffusion process. DDIM allows for deterministic sampling with better quality.

### 3.2 Masked Conditional Diffusion

In this section, we describe the proposed method, Masked Conditional Diffusion (MacDiff), as a unified framework for human skeleton modeling. Fig. 1 illustrates the overall pipeline of our method.

Concretely, we jointly optimize a semantic encoder  $\mathcal{E}$  and a denoising decoder  $\mathcal{D}$  (also denoted as  $\epsilon_\theta$ ) by training a conditional diffusion model  $\epsilon_\theta(\mathbf{x}_t, t, \mathbf{z})$ . First, the encoder extracts a latent representation from the masked *view* of skeleton data  $\mathbf{x}$ , formulated as  $\mathbf{z} = \mathcal{E}(\mathcal{M}(\mathbf{x}))$  where  $\mathcal{M}(\cdot)$  denotes the masking operation. Then, the decoder predicts the noise contained in  $\mathbf{x}_t$  conditioned on  $\mathbf{z}$ . With the aforementioned pre-training, the encoder learns to extract compact representation applicable for discriminative downstream tasks.

**Patchify and Embedding.** Given the input skeleton  $\mathbf{x} \in \mathbb{R}^{T_0 \times V \times 3}$ , we first divide it into non-overlapping patches with equal length along the temporal dimension  $\mathbf{x}' \in \mathbb{R}^{T_p \times V \times (l \times 3)}$ .  $T_0$  is the number of frames,  $V$  is the number of joints,  $l$  is the patch length, and  $T_p = T_0/l$ . The patches are then flattened and embedded to  $C$  dimensions with a trainable linear projection:  $\mathbf{E} = \text{Embedding}(\mathbf{x}') \in \mathbb{R}^{T_p \times V \times C}$ . Now we have  $T \cdot V$  tokens of dimension  $C$  as the encoder’s input.

**Random Masking.** Considering the redundancy in the skeleton, especially in the temporal dimension, naive reconstruction training can result in the undesirable shortcuts and degrade the model performance. Meanwhile, based on the information bottleneck principle [15,30], the model learns a compact representation if information compression is introduced to remove redundancy. Therefore, we employ random masking for the encoder input with a masking ratio  $r$ , retaining a total of  $K = \lceil (1-r) \cdot T_p \cdot V \rceil$  tokens. In practice, we adopt an extremely high masking ratio  $r = 90\%$ , enforcing a tight bottleneck on the representation. In addition, the masking operation significantly speeds up training by reducing the encoder’s computational graph.

**Model Architecture.** Our encoder and decoder both follow a vanilla Transformer [51] architecture. First, trainable spatial and temporal positional embeddings  $\mathbf{E}_{pos}^s \in \mathbb{R}^{1 \times V \times C}$ ,  $\mathbf{E}_{pos}^t \in \mathbb{R}^{T_p \times 1 \times C}$  are added to the tokens  $\mathbf{E}$  with broadcasting. Note that this is implemented before the masking operation for the encoder. The Transformer network consists of alternating layers of multi-head self-attention (MSA) and multi-layer perceptron (MLP) with residual connection. Layer Norm (LN) is applied before each layer and after the last layer.

The encoder output  $\mathbf{z}_{local} \in \mathbb{R}^{K \times C}$  is the local representations corresponding to unmasked patches. The global representation  $\mathbf{z}_{global}$  is obtained by pooling all tokens. For the decoder output, we linearly project it to the final prediction of the same shape as  $\mathbf{x}_0$ .

**Conditioning.** To incorporate the condition  $\mathbf{z}$  into the denoising decoder, we replace Layer Norm with Adaptive Layer Norm (AdaLN), following [38]:

$$\text{AdaLN}(\mathbf{h}, \mathbf{z}, t) = \mathbf{z}_s \cdot (t_s \cdot \text{LN}(\mathbf{h}) + t_b) + \mathbf{z}_b, \quad (4)$$

where  $\mathbf{h}$  is the hidden representation,  $(t_s, t_b)$  and  $(\mathbf{z}_s, \mathbf{z}_b)$  are obtained from linear projection of the timestep embedding  $t$  and condition  $\mathbf{z}$ , respectively. Through AdaLN layers, the condition  $\mathbf{z}$  guides the denoising process by scaling and shifting normalized hidden representation. To further disentangle the contributions of the encoder and decoder to the final prediction, we dropout  $\mathbf{z}$  with a probability of 0.1, meanwhile enabling unconditional generation.

In practice, we observe an over-smoothing problem of the encoder if we simply utilize  $\mathbf{z}_{global}$  (with broadcast) as  $\mathbf{z}$ . Over-smoothing is a common problem

for GCNs and Transformers [12, 19] that degrades performance and can be reflected by high similarity between tokens. We aim to increase token uniformity by preserving more local information in tokens in addition to global information. To this end, we unshuffle the unmasked tokens  $\mathbf{z}_{local}$  and fill the masked positions with  $\mathbf{z}_{global}$  to form  $\mathbf{z}$ . Since a high masking ratio is implemented, the model prioritizes optimizing the global representation while also attempts to benefit from local information by directly optimizing local representations.

**Diffusion Training.** Considering the discrepancy between the statistics of original skeleton data and the standard Gaussian distribution, we normalize the data using the mean  $\mu$  and standard variation  $\sigma$  calculated from the training set:

$$\mathbf{x}_0 = \frac{\mathbf{x}_{orig} - \mu}{\sigma}, \quad \mu, \sigma \in \mathbb{R}^{1 \times 1 \times 3}. \quad (5)$$

For detailed settings of diffusion, we adopt some common practices including  $\epsilon$ -prediction, total timesteps  $T = 1000$  and loss weights  $\gamma_t = 1$  [18]. Our training objective is a simplified and conditional version of Eq. (3):

$$\mathcal{L} = \mathbb{E}_{\mathbf{x}_0, t, \epsilon} \left[ \|\epsilon - \mathcal{D}(\sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{1 - \alpha_t} \epsilon, t, \mathcal{E}(\mathcal{M}(\mathbf{x}_0)))\|^2 \right]. \quad (6)$$

A minority choice in our work regarding diffusion is the inverse-cosine schedule [20]. The inverse-cosine schedule pulls the noise level of all timesteps towards a medium level compared with commonly-used cosine [1] or linear [18] schedules. We verify this choice with experiments (see Sec. 4.5).

### 3.3 Information Analysis on MacDiff

In this section, we conduct information-theoretic analyses on our proposed framework MacDiff. We formulate the generative objective of MacDiff as an improvement of contrastive learning (CL) objectives, leading to a better guarantee of downstream performance. For mathematical formulation, we use random variables  $X, X_t, X_m$  and  $Z$  to denote the original view, noisy view, masked view, and latent representation of the skeleton data.

**The Training Objective of MacDiff.** A generative model (*e.g.*, VAEs and MAEs) that predicts some target  $X$  from latent code  $V$  maximizes their mutual information (MI), *i.e.*  $I(X; V)$ . In the case of MacDiff,  $V$  takes the form of  $(Z, X_t)$ . MacDiff can thereby be described as:

$$\max_{\mathcal{E}, \mathcal{D}} I(X; (Z, X_t)), \quad Z = \mathcal{E}(X_m). \quad (7)$$

This training objective is further formulated as follows:

$$I(X; (Z, X_t)) = I(X; Z) + I(X; X_t | Z). \quad (8)$$

The first term  $I(X; Z)$  trains the encoder to contain more information about  $X$  in representation  $Z$ , while the second term trains the decoder to predict  $X$  from  $X_t$  conditioned on  $Z$ . By employing the encoder-decoder design and dropout on

the representation, we further disentangle the encoder’s high-level representation learning from the decoder’s low-level prediction objective.

**MacDiff as an Improvement of Contrastive Learning.** We can further decompose  $I(X; Z)$  into two terms with the intermediate variable  $X_t$ :

$$I(X; Z) = I(X_t; Z) + I(X; Z|X_t). \quad (9)$$

We point out that the first term is consistent with the contrastive learning objective. CL assumes that the information needed for downstream task  $Y$  is shared between two views  $X_1, X_2$  [44, 52, 53]. Therefore, CL aims to optimize the MI between their representations  $Z_1, Z_2$ .

$$I(Z_1; Z_2) \leq I(Z_1; X_2) \leq I(X_1; X_2) = I(X; Y). \quad (10)$$

In our case,  $X_1, X_2$  correspond to the masked and noisy views  $X_m, X_t$ . MacDiff does not directly extract representation from  $X_t$ . It instead optimizes  $I(Z_1, X_2) = I(Z, X_t)$  as a tighter lower bound of  $I(X_1; X_2)$ , which avoids feature collapse.

Moreover, the second term aims to contain more information about  $X$  in  $Z$  that is complementary to  $X_t$ . CL is proved to suffer from the *discriminative information overfitting* problem [30] that the model is biased to extract only the discriminative information of two views. However, optimizing  $I(X; Z|X_t)$  requires the representation to contain more task-relevant information that is not shared between views. Therefore, the generative objective of MacDiff can be viewed as an improvement of CL methods and provides a better theoretical guarantee of downstream performance as discussed next.

**Relation to Downstream Performance.** The Bayes error rate  $P_e$  is the lowest error that can be achieved by any classifier trained on the given data representations, defined as  $P_e = 1 - \mathbb{E}_{z \sim p(Z)}[\max_{y \in Y} p(y|z)]$ , where  $Z$  denotes the representation and  $Y$  denotes the labels. Then, we can prove that (refer to the supplementary material):

**Theorem 1.** (*Bayes Error Rate of Representations*) *For arbitrary data representation distribution  $Z$ , and  $V$  denotes a certain view of the data, its Bayes error rate can be estimated as:*

$$P_e \leq 1 - e^{-(H(Y) - I(Z; Y))} \quad (11)$$

$$\leq 1 - e^{-(H(Y) - I(Z; Y; V) - I(Z; Y|V))}. \quad (12)$$

We set  $V = X_t$  in this theorem. Thus, our goal is to increase the terms  $I(Z; Y; X_t)$  and  $I(Z; Y|X_t)$ . These two terms are bounded by  $I(Z; X_t)$  and  $I(Z; X|X_t)$  respectively. As shown in Eq. (9), the MacDiff objective directly increases both terms. Note that the contrastive learning objective merely optimizes the first term  $I(Z; X_t)$ , which explains the improved downstream performance of our method compared to contrastive-only methods.

### 3.4 Diffusion-Based Data Augmentation

For generative self-supervised methods (*e.g.*, MAEs), only the encoders are utilized for downstream tasks, while the rest of the models are completely discarded.



In MacDiff, we propose that the denoising decoder can be used for data augmentation when fine-tuning the encoder. We focus on scenarios where labeled data is scarce, which is quite common in reality, especially for skeletons.

To synthesize training data that follows the real-data distribution  $q(\mathbf{x}, y)$ , generative methods have to be label-preserving. Existing methods achieve this by leveraging pre-trained text-guided diffusion [56,60]. However, our semi-supervised setting does not provide text guidance, which forms a more general challenge of synthesizing label-preserving samples given labeled ones.

To this end, we propose diffusion-based data augmentation based on the assumption that samples generated with the same representation guidance are label-consistent. Specifically, we take two steps: (1) pre-calculate the representations of labeled samples, and (2) generate new samples with the decoder conditioned on these representations. Nevertheless, sampling from Gaussian noise is time-consuming and thus can only be performed before training. We find that one-step denoising from some medium timestep  $t_s$  yields similar effects with smaller computational cost, which allows for generating diverse augmented data at different epochs.

## 4 Experiments

### 4.1 Experimental Setup

**Datasets.** For evaluation, our experiments are conducted on the following three datasets: NTU RGB+D 60 dataset (NTU 60) [43], NTU RGB+D 120 dataset (NTU 120) [29] and PKU Multi-Modality Dataset (PKUMMD) [28]. All three datasets use 25 joints to represent the human body.

NTU 60 is a large-scale dataset for human action recognition with 60 categories and 56,578 videos. We follow the widely-used evaluation protocols, cross-subject (xsub) and cross-view (xview). The former uses action sequences from half of the 40 subjects for training, and the rest for testing. The latter uses sequences from camera 2,3 for training and sequences from camera 1 for testing.

NTU 120 is an extension of NTU 60, with 120 categories and 114,480 videos from 106 subjects. Evaluation protocols on NTU 120 are cross-subject (xsub) and cross-setup (xset). Specifically, xset divides sequences into 32 setups based on the camera distance and background, half of which are used for training and the rest for testing.

PKUMMD covers a multi-modality 3D understanding of human actions, with 52 categories and almost 20,000 instances. PKUMMD is divided into part I and II, and part II is more challenging due to the noise caused by view variation. We split training and testing sets according to the cross-subject protocol.

**Implementation Details.** The input sequence of 300 frames is cropped and interpolated to 120 frames, and the patch length  $l = 4$ . Apart from random crop, we use random rotation and small Gaussian noise ( $\sigma = 0.005$ ) as data augmentation. Note that the small noise is only added to encoder inputs. For Transformer architecture, the embedding dimension is 256, the MLP hidden

**Table 1:** Comparison of linear evaluation results on NTU 60, NTU 120, and PKUMMD datasets. 3s- represents the ensemble results of joint(J), bone(B) and motion(M) streams. **Bold** and underlined indicate the best and second best results, respectively. The same notation applies throughout.

Method	Stream	NTU 60		NTU 120		PKU I
		xsub	xview	xsub	xset	
3s-CrosSCLR [25]	J+M+B	77.8	83.4	67.9	66.7	84.9
3s-AimCLR [49]	J+M+B	78.9	83.8	68.2	68.8	87.4
3s-SkeleMixCLR [49]	J+M+B	82.7	87.1	70.5	70.7	91.1
3s-ActCLR [27]	J+M+B	84.3	88.8	74.3	75.7	-
3s-CPM [57]	J+M+B	83.2	87.0	73.0	74.0	90.7
LongT GAN [62]	J	39.1	48.1	-	-	67.7
MS <sup>2</sup> L [26]	J	52.6	-	-	-	64.9
AS-CAL [40]	J	58.5	64.8	48.6	49.2	-
ISC [48]	J	76.3	85.2	67.1	67.9	80.9
GL-Transformer [21]	J	76.3	83.8	66.0	68.7	-
CMD [33]	J	79.4	86.9	70.3	71.5	-
PCM <sup>3</sup> [58]	J	83.9	<u>90.4</u>	76.5	77.5	-
SkeletonMAE [55]	J	74.8	77.7	72.5	73.5	82.8
MAMP [32]	J	<u>84.9</u>	89.1	<u>78.6</u>	<u>79.1</u>	<u>92.2</u>
MacDiff (Ours)	J	<b>86.4</b>	<b>91.0</b>	<b>79.4</b>	<b>80.2</b>	<b>92.8</b>

dimension is 1024, and the number of heads in MSA is 8. By default, the encoder and decoder have 8 layers and 5 layers, respectively. We train our model on four NVIDIA TITAN Xp GPUs with a total batch size of 128 for 500 epochs. The AdamW optimizer is adopted with the learning rate decreasing from 1e-3 to 1e-5.

## 4.2 Self-Supervised Learning Evaluation

**Linear Evaluation.** In the linear evaluation protocol, a linear classifier is post-attached to the encoder to classify the learned representations. We fix the encoder and train the classifier for 100 epochs with the SGD optimizer and a learning rate of 0.1. We compare MacDiff with latest methods, with action recognition accuracy reported as a measurement.

As shown in Tab. 1, our method surpasses high-performing reconstruction-based methods, *e.g.*, SkeletonMAE [55] and MAMP [55]. With only the joint stream, our method also outperforms multi-stream contrastive learning methods, *e.g.* 3s-AimCLR [49], 3s-CMD [33] and 3s-ActCLR [27]. The result demonstrates that MacDiff captures the spatial-temporal correlation of skeletons better than existing methods, and also confirms our theoretical analysis that MacDiff provides a better framework than contrastive-only paradigms.

**Supervised Fine-tuning Evaluation.** In the fine-tuning evaluation protocol, we attach an MLP head to the pre-trained encoder and train the whole model for another 100 epochs with the AdamW optimizer and the learning rate decreasing from 3e-4 to 1e-5.

As shown in Tab. 2, our method yields comparable results to MAMP and outperforms other existing methods. We point out that our performance gap in the fine-tuning protocol with MAMP is trivial since we share the same encoder architecture and the learned representation is disrupted during fine-tuning.

**Table 2:** Comparison of supervised fine-tuning evaluation results on NTU 60 xsub and xview datasets. The bottom four rows are unified models, among which UPS is a supervised method.

Method	Backbone	NTU 60	
		xsub	xview
CrosSCLR [25]	3s-ST-GCN	86.2	92.5
AimCLR [49]	3s-ST-GCN	86.9	92.8
ActCLR [27]	3s-ST-GCN	88.2	93.9
MCC [46]	2s-AGCN	89.7	96.3
SkeletonMAE [55]	Transformer	88.5	94.7
MAMP [55]	Transformer	<b>93.1</b>	<b>97.5</b>
Hi-TRS [7]	3s-Transformer	90.0	95.7
MotionBERT [63]	DSTformer	<u>93.0</u>	97.2
UPS [14]	Transformer	92.6	97.0
<b>MacDiff (Ours)</b>	Transformer	92.7	<u>97.3</u>

**Table 3:** Comparison of transfer learning results on PKUMMD II dataset. The source datasets are NTU 60, NTU 120 and PKUMMD I. All datasets use the xsub split.

Method	To PKU II		
	NTU 60	NTU 120	PKU I
LongT GAN [62]	44.8	-	43.6
MS <sup>2</sup> L [26]	45.8	-	44.1
ISC [48]	51.1	52.3	45.1
CMD [33]	56.0	57.0	-
SkeletonMAE [55]	58.4	61.0	62.5
MAMP [32]	<u>70.6</u>	<u>73.2</u>	<u>70.1</u>
<b>MacDiff (Ours)</b>	<b>72.2</b>	<b>73.4</b>	<b>71.4</b>

**Table 4:** Comparison of semi-supervised fine-tuning results on NTU 60 datasets. By default we set the starting timestep  $t_s = 500$  and augment-to-real ratio  $\lambda = 2.0, 0.5, 0.25$  for  $p = 1\%, 2\%, 10\%$ . The results are averaged over 5 runs. For MacDiff, results with and without diffusion-based augmentation are both reported.

Method	NTU 60 xsub			NTU 60 xview		
	1%	2%	10%	1%	2%	10%
ISC [48]	35.7	-	65.1	38.1	-	72.5
3s-AimCLR [49]	54.8	-	78.2	54.3	-	81.6
3s-CMD [33]	55.6	-	79.0	55.5	-	82.4
CPM [57]	56.7	-	73.0	57.5	-	77.1
PCM <sup>3</sup> [58]	53.8	-	77.1	53.1	-	82.8
SkeletonMAE [55]	54.4	-	80.6	54.6	-	83.5
MAMP [32]	<u>66.0</u>	80.3	88.0	68.7	83.5	91.5
<b>MacDiff w/o aug</b>	65.6	<u>80.7</u>	<u>88.2</u>	<u>77.3</u>	<u>84.1</u>	<u>92.5</u>
<b>MacDiff</b>	<b>72.0</b>	<b>82.1</b>	<b>89.2</b>	<b>79.2</b>	<b>85.6</b>	<b>93.1</b>

Meanwhile, MacDiff outperforms other unified models such as Hi-TRS [7], MotionBERT [63] and UPS [14].

**Transfer Learning Evaluation.** In the transfer learning evaluation protocol, the backbones are pre-trained on a source dataset and evaluated on a target dataset following linear evaluation protocol to keep the backbone intact.

As shown in Tab. 2, our method achieves significant transfer learning performance on the challenging PKUMMD II, demonstrating the generalization ability and robustness of our method.

### 4.3 Semi-Supervised Fine-tuning with Diffusion-based Data Augmentation

We next evaluate the effectiveness of the diffusion-based data augmentation in scenarios with limited labeled data. We report results following the semi-supervised protocol, which is consistent with the fine-tuning protocol except

that only a proportion  $p$  of the training set is used. We evaluate our methods when  $p = 1\%, 2\%, 10\%$ .

We explore different augment-to-real ratios  $\lambda$ , defined as the ratio of augmented data to real data, for different proportions of training data. We empirically find that the optimal augment-to-real ratio  $\lambda$  declines as  $p$  increases. Intuitively, when labeled data is severely scarce, augmented data that falls in the neighborhood of labeled data representations helps the classifier to learn more robust boundaries. However, as the cardinality of the training set increases, since the dataset inherently contains ambiguous class boundaries (*e.g.*, "phone call" and "play with phone/tablet" in NTU 60), generating new data around them further confuse the classifier. See more experiments in Sec. 4.5.

As shown in Tab. 4, with the aid of augmented training samples, MacDiff outperforms state-of-the-art MAMP by 6.0%, 1.8%, 1.2% in NTU 60 xsub, and 10.5%, 2.1%, 1.6% in NTU 60 xview within the three settings, respectively. Compared with encoder-only MacDiff, the augmentation brings significant performance gain of 6.4%, 1.4%, 1.0% in NTU 60 xsub, and 1.9%, 1.5%, 0.6% in NTU 60 xview.

#### 4.4 Generative Evaluation

In this section, we implement MacDiff for motion reconstruction and motion generation tasks. We compare our method with reconstruction-based method SkeletonMAE [55] and diffusion-based methods DDIM [45] and MDM [47]. Note that MAMP [32] cannot be applied for either task because it predicts the normalized motion. For fair comparison, all methods are implemented with our Transformer decoder architecture. In addition, we also implement the original MDM (denoted as MDM-orig) with temporal-only attention, 8 layers, and 512 hidden dimensions. Please find implementation details in the supplementary material. All experiments are conducted on the testing set of NTU 60 xsub.

**Motion Reconstruction.** Real-world skeleton data suffer from occlusions, resulting in incomplete sequences. We evaluate motion reconstruction in two types of occlusions: (1) random consecutive frames, and (2) a random body part from {trunk, left arm, right arm, left leg, right leg}, following the division of previous works. We follow a diffusion-based inpainting paradigm [31] for DDIM and MacDiff. The MacDiff decoder is fine-tuned for another 100 epochs with the encoder fixed and only global representations.

We report Mean Per Joint Position Error (MPJPE) as our metric. As shown in Tab. 5, MacDiff is capable of recovering incomplete skeletons as a unified framework and surpasses reconstruction-based SkeletonMAE and DDIM.

**Motion Generation.** We utilize the MacDiff decoder for unconditional motion generation. We report four metrics FID, KID, diversity, and precision/recall. Please refer to the supplementary material for detailed implementation of these metrics. Note that reconstruction-based methods are not capable of unconditional generation. As shown in Tab. 6, MacDiff achieves comparable results with DDIM and MDM.

**Table 5:** Comparison of motion reconstruction results on NTU 60 xsub. MPJPE is reported on four occlusion settings: 10 frames, 20 frames, 40 frames and 1 body part.

	10 frames	20 frames	40 frames	1 body part
SkeletonMAE [55]	0.191	0.221	0.255	0.319
DDIM [45]	<u>0.041</u>	<b>0.087</b>	<u>0.205</u>	<u>0.251</u>
<b>MacDiff (Ours)</b>	<b>0.033</b>	<u>0.089</u>	<b>0.147</b>	<b>0.241</b>

**Table 6:** Comparison of motion generation results on NTU 60 xsub. MacDiff is evaluated both with and without fine-tuning.

Method	Prediction	FID↓	KID↓	Diversity↑	Precision↑ Recall↑
Real data	-	0.05	0.0001	2.09	0.886, 0.908
MDM-orig [47]	$x_0$	3.49	0.0328	1.51	0.047, 0.030
MDM [47]	$x_0$	<b>1.06</b>	<b>0.0059</b>	1.43	<b>0.828</b> , 0.070
DDIM [45]	$\epsilon$	1.32	0.0082	1.89	<u>0.576</u> , <b>0.537</b>
<b>MacDiff (Ours)</b>	$\epsilon$	1.52	0.0085	<u>1.96</u>	0.414, 0.289
<b>+ fine-tune</b>	$\epsilon$	<u>1.30</u>	<u>0.0070</u>	<b>1.97</b>	0.460, <u>0.505</u>

**Qualitative Results.** We provide visualization of motion reconstruction, unconditional motion generation and one-step denoising results (for data augmentation) in the supplementary material.

#### 4.5 Ablation Study

**Masking Strategy and Ratio.** In Tab. 7, we compare the results of different masking strategies, including random masking, temporal-only masking, tube masking [50], spatial-temporal masking and motion-aware random masking [32]. For tube masking, the tube length is set to 5. For spatial-temporal masking, we keep 8 out of 25 joints and 10 out of 30 temporal patches. For motion-aware masking, we follow the implementation of MAMP. The results show that the simple random masking works best as a spacetime-agnostic masking. We also compare different masking ratios and find a high masking ratio of 90% works best, which coincides with the findings in the video field.

**Noise Schedule.** We construct a series of noise schedules as linear combinations of the inverse-cosine [20] and cosine schedule controlled by  $\tau$  (see definition in the supplementary material).  $\tau = 1$  and  $\tau = -1$  represents the inverse-cosine and cosine schedule, respectively. We compare these schedules with the widely used linear schedule. As shown in Tab. 8, cosine-based schedule performs better than the linear schedule, and the performance peaks at  $\tau = 1$ , indicating that medium noise levels are preferred for representation learning.

**Diffusion-based Data Augmentation.** For the ablation study of our diffusion-based data augmentation, We compare the effects of different starting timestep  $t_s$  and augment-to-real ratio  $\lambda$ . As shown in Tab. 9, the performance gain is highest when  $t_s = 500$ . Intuitively, an overly large  $t_s$  may introduce too much noise since we implement one-step denoising, while an overly small  $t_s$  fails to

**Table 7:** Ablation study on the masking strategy and the masking ratio. We report results on NTU 60 xsub under the linear evaluation protocol.

Strategy	NTU 60 xsub	Ratio	NTU 60 xsub
Temporal	84.1	0	79.3
Tube	83.1	50%	82.7
Spatial-temporal	85.3	80%	83.8
Random	<b>86.4</b>	90%	<b>86.4</b>
Motion-aware	85.5	95%	83.6

**Table 8:** Ablation study on the noise schedule. We report results on NTU 60 xsub under the linear evaluation protocol.

Noise Schedule	NTU 60 xsub
$\tau = 1.5$	85.7
$\tau = 1.0$ (Inverse)	<b>86.4</b>
$\tau = 0.5$	85.8
$\tau = -0.5$	85.2
$\tau = -1.0$ (Cosine)	83.8
Linear	83.4

**Table 9:** Ablation study on the starting timestep  $t_s$  and augment-to-real ratio  $\lambda$  of diffusion-based augmentation. We report results on NTU 60 xsub under the semi-supervised 1% and 10% protocol.  $t_s = 0$  or  $\lambda = 0$  means results without augmentation.

$t_s$	semi 1%	$\lambda$	semi 1%	semi 10%
0	65.6	0	65.6	88.2
100	66.4	0.25	67.8	89.1
300	71.4	0.5	68.2	<b>89.2</b>
500	<b>72.0</b>	1.0	70.0	88.5
900	68.7	2.0	<b>72.0</b>	-

**Table 10:** Ablation study on the depth of the decoder. We report results on NTU 60 xsub under the linear evaluation protocol.

Depth	NTU 60 xsub
2	84.8
3	<b>86.4</b>
4	86.0
5	85.9

provide sufficient strength of augmentation. For larger proportion of training data, the optimal  $\lambda$  is smaller, which is consistent with our analysis in Sec. 4.3.

**Decoder Design.** Tab. 10 reports the effects of different decoder depths on representation learning. The best result is achieved with a depth of 3, but generally our method is robust to the decoder depth. Therefore, we adopt a depth of 5 by default considering the generative capability.

## 5 Conclusion

We present MacDiff, a novel generative framework to enhance skeleton representation learning for human action understanding. By training a diffusion decoder guided by the representation from the encoder, the encoder is enforced to contain rich semantics in the representation. We formulate the objective of MacDiff as an improvement of the contrastive learning objective, theoretically demonstrating the effectiveness of the proposed framework.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant No.62172020, and in part by the Key Laboratory of Science, Technology and Standard in Press Industry (Key Laboratory of Intelligent Press Media Technology).

## References

1. Alex Nichol, P.D.: Improved denoising diffusion probabilistic models. arXiv preprint arXiv:2102.09672 (2021)
2. Camara, F., Bellotto, N., Cosar, S., Nathanael, D., Althoff, M., Wu, J., Ruenz, J., Dietrich, A., Fox, C.W.: Pedestrian models for autonomous driving part I: low-level models, from sensing to tracking. *IEEE Transactions on Intelligent Transportation Systems* **22**(10), 6131–6151 (2020)
3. Chen, L.H., Zhang, J., Li, Y., Pang, Y., Xia, X., Liu, T.: HumanMAC: Masked motion completion for human motion prediction. In: *Int. Conf. Comput. Vis.* (2023)
4. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *Int. Conf. Mach. Learn.* (2020)
5. Chen, X., Jiang, B., Liu, W., Huang, Z., Fu, B., Chen, T., Yu, J., Yu, G.: Executing your commands via motion diffusion in latent space. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2023)
6. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297 (2020)
7. Chen, Y., Zhao, L., Yuan, J., Tian, Y., Xia, Z., Geng, S., Han, L., Metaxas, D.N.: Hierarchically self-supervised transformer for human skeleton representation learning. In: *Eur. Conf. Comput. Vis.* (2022)
8. Chen, Z., Liu, H., Guo, T., Chen, Z., Song, P., Tang, H.: Contrastive learning from spatio-temporal mixed skeleton sequences for self-supervised skeleton-based action recognition. arXiv preprint arXiv:2207.03065 (2022)
9. Ci, H., Wu, M., Zhu, W., Ma, X., Dong, H., Zhong, F., Wang, Y.: GFPose: Learning 3D human pose prior with gradient fields. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2023)
10. Ci, Y., Wang, Y., Chen, M., Tang, S., Bai, L., Zhu, F., Zhao, R., Yu, F., Qi, D., Ouyang, W.: UniHCP: A unified model for human-centric perceptions. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2023)
11. Dhariwal, P., Nichol, A.: Diffusion models beat GANs on image synthesis. In: *Adv. Neural Inform. Process. Syst.* (2021)
12. Dong, Y., Cordonnier, J.B., Loukas, A.: Attention is not all you need: Pure attention loses rank doubly exponentially with depth. arXiv preprint arXiv:2103.03404 (2021)
13. Flaborea, A., Collorone, L., di Melendugno Stefano D’Arrigo, G.M.D., Prenkaj, B., Galasso, F.: Multimodal motion conditioned diffusion model for skeleton-based video anomaly detection. In: *Int. Conf. Comput. Vis.* (2023)
14. Foo, L.G., Li, T., Rahmani, H., Ke, Q., Liu, J.: Unified pose sequence modeling. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2023)
15. Goldfeld, Z., Polyanskiy, Y.: The information bottleneck problem and its applications in machine learning. *IEEE JSAIT* **1**(1), 19–38 (2020)
16. Gong, J., Foo, L.G., Fan, Z., Ke, Q., Rahmani, H., Liu, J.: DiffPose: Toward more reliable 3D pose estimation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2023)
17. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2022)
18. Ho, J., Jain, A., bbeel, P.: Denoising diffusion probabilistic models. arXiv preprint arXiv:2006.11239 (2020)
19. Huang, W., Rong, Y., Xu, T., Sun, F., Huang, J.: Tackling over-smoothing for general graph convolutional networks. arXiv preprint arXiv:2008.09864 (2020)

20. Hudson, D.A., Zoran, D., Malinowski, M., Lampinen, A.K., Jaegle, A., McClelland, J.L., Matthey, L., Hill, F., Lerchner, A.: SODA: Bottleneck diffusion models for representation learning. arXiv preprint arXiv:2311.17901 (2023)
21. Kim, B., Chang, H.J., Kim, J., , Choi, J.Y.: Global-local motion transformer for unsupervised skeleton-based action learning. In: Eur. Conf. Comput. Vis. (2022)
22. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
23. Kun Su, Xiulong Liu, E.S.: PREDICT & CLUSTER: Unsupervised skeleton based action recognition. arXiv preprint arXiv:1911.12409 (2019)
24. Lee, J., Ahn, B.: Real-time human action recognition with a low-cost RGB camera and mobile robot platform. *Sensors* **20**(10), 2886 (2020)
25. Li, L., Wang, M., Ni, B., Wang, H., Yang, J., Zhang, W.: 3D human action representation learning via cross-view consistency pursuit. In: IEEE Conf. Comput. Vis. Pattern Recog. (2021)
26. Lin, L., Song, S., Yang, W., Liu, J.: MS2L: Multi-task self-supervised learning for skeleton based action recognition. In: ACM Int. Conf. Multimedia (2020)
27. Lin, L., Zhang, J., Liu, J.: Actionlet-dependent contrastive learning for unsupervised skeleton-based action recognition. In: IEEE Conf. Comput. Vis. Pattern Recog. (2023)
28. Liu, J., Song, S., Liu, C., Li, Y., , Hu, Y.: A benchmark dataset and comparison study for multi-modal human action analytics. In: ACM TOMM (2020)
29. Liu, J., Shahroudy, A., Perez, M., Wang, G., Duan, L.Y., Kot, A.C.: NTU RGB+D 120: A large-scale benchmark for 3D human activity understanding. *IEEE Trans. Pattern Anal. Mach. Intell.* (2019)
30. Liu, Z., Li, B., Han, C., Guo, T., Nie, X.: Masked reconstruction contrastive learning with information bottleneck principle. arXiv preprint arXiv:2211.09013 (2022)
31. Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Gool, R.T.L.V.: RePaint: Inpainting using denoising diffusion probabilistic models. In: IEEE Conf. Comput. Vis. Pattern Recog. (2022)
32. Mao, Y., Deng, J., Zhou, W., Fang, Y., Ouyang, W., Li, H.: Masked motion predictors are strong 3D action representation learners. In: Int. Conf. Comput. Vis. (2023)
33. Mao, Y., Zhou, W., Lu, Z., Deng, J., Li, H.: CMD: Self-supervised 3D action representation learning with cross-modal mutual distillation. In: Eur. Conf. Comput. Vis. (2022)
34. Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: SDEdit: Guided image synthesis and editing with stochastic differential equations. arXiv preprint arXiv:2108.01073 (2021)
35. Nam, J., Lee, G., Kim, S., Kim, H., Cho, H., Kim, S., Kim, S.: Diffusion model for dense matching. In: Int. Conf. Learn. Represent. (2024)
36. Namekata, K., Sabour, A., Fidler, S., Kim, S.W.: EmerDiff: Emerging pixel-level semantic knowledge in diffusion models. In: Int. Conf. Learn. Represent. (2024)
37. van den Oord, A., Vinyals, O., koray kavukcuoglu: Neural discrete representation learning. In: Adv. Neural Inform. Process. Syst. (2017)
38. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: Int. Conf. Comput. Vis. (2023)
39. Preechakul, K., Chatthee, N., Wizadwongsa, S., Suwajanakorn, S.: Diffusion autoencoders: Toward a meaningful and decodable representation. In: IEEE Conf. Comput. Vis. Pattern Recog. (2022)



40. Rao, H., Xu, S., Hu, X., Cheng, J., Hu, B.: Augmented skeleton based contrastive action learning with momentum lstm for unsupervised action recognition. In: Information Sciences (2021)
41. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. arXiv preprint arXiv:2112.10752 (2021)
42. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S.K.S., Ayan, B.K., Mahdavi, S.S., Lopes, R.G., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M.: Photorealistic text-to-image diffusion models with deep language understanding. arXiv preprint arXiv:2205.11487 (2022)
43. Shahroudy, A., Liu, J., Ng, T.T., Wang, G.: NTU RGB+D: A large scale dataset for 3D human activity analysis. In: IEEE Conf. Comput. Vis. Pattern Recog. (2016)
44. Shwartz-Ziv, R., LeCun, Y.: To compress or not to compress - self-supervised learning and information theory: A review. arXiv preprint arXiv:2304.09355 (2023)
45. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
46. Su, Y., Lin, G., Wu, Q.: Self-supervised 3D skeleton action representation learning with motion consistency and continuity. In: Int. Conf. Comput. Vis. (2021)
47. Tevet, G., Raab, S., Gordon, B., Shafir, Y., Cohen-or, D., Bermano, A.H.: Human motion diffusion model. In: Int. Conf. Learn. Represent. (2023)
48. Thoker, F.M., Doughty, H., Snoek, C.G.: Skeleton-contrastive 3D action representation learning. In: ACM Int. Conf. Multimedia (2021)
49. Tianyu, G., Hong, L., Zhan, C., Mengyuan, L., Tao, W., Runwei, D.: Contrastive learning from extremely augmented skeleton sequences for self-supervised action recognition. In: AAAI (2022)
50. Tong, Z., Song, Y., Wang, J., Wang, L.: VideoMAE: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In: Adv. Neural Inform. Process. Syst. (2022)
51. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Łukasz Kaiser, Polosukhin, I.: Attention is all you need. In: Adv. Neural Inform. Process. Syst. (2017)
52. Wang, H., Guo, X., Deng, Z.H., Lu, Y.: What makes for good representations for contrastive learning. In: Adv. Neural Inform. Process. Syst. (2020)
53. Wang, H., Guo, X., Deng, Z., Lu, Y.: Rethinking minimal sufficient representation in contrastive learning. In: IEEE Conf. Comput. Vis. Pattern Recog. (2022)
54. Wang, X., Fang, Z., Li, X., Li, X., Chen, C., Liu, M.: Skeleton-in-context: Unified skeleton sequence modeling with in-context learning. arXiv preprint arXiv:2312.03703 (2023)
55. Wu, W., Hua, Y., Zheng, C., Wu, S., Chen, C., Lu, A.: SkeletonMAE: Spatial-temporal masked autoencoders for self-supervised skeleton action recognition. arXiv preprint arXiv:2209.02399 (2022)
56. Yuan, J., Zhang, J., Sun, S., Torr, P., Zhao, B.: Real-Fake: Effective training data synthesis through distribution matching. In: Int. Conf. Learn. Represent. (2024)
57. Zhang, H., Hou, Y., Zhang, W., Li, W.: Contrastive positive mining for unsupervised 3D action representation learning. In: Eur. Conf. Comput. Vis. (2022)
58. Zhang, J., Lin, L., Liu, J.: Prompted contrast with masked motion modeling: Towards versatile 3D action representation learning. In: ACMMM (2023)
59. Zhang, M., Cai, Z., Pan, L., Hong, F., Guo, X., Yang, L., Liu, Z.: MotionDiffuse: Text-driven human motion generation with diffusion model. IEEE Trans. Pattern Anal. Mach. Intell. pp. 1–15 (2024)
60. Zhang, Y., Zhou, D., Hooi, B., Wang, K., Feng, J.: Expanding small-scale datasets with guided imagination. In: Adv. Neural Inform. Process. Syst. (2023)

61. Zhang, Z., Zhao, Z., Lin, Z.: Unsupervised representation learning from pre-trained diffusion probabilistic models. In: Adv. Neural Inform. Process. Syst. (2022)
62. Zheng, N., Wen, J., Liu, R., Long, L., Dai, J., , Gong, Z.: Unsupervised representation learning with long-term dynamics for skeleton based action recognition. In: AAAI (2018)
63. Zhu, W., Ma, X., Liu, Z., Liu, L., Wu, W., Wang, Y.: MotionBERT: A unified perspective on learning human motion representations. In: Int. Conf. Comput. Vis. (2023)