

# Bridging the Skeleton-Text Modality Gap: Diffusion-Powered Modality Alignment for Zero-shot Skeleton-based Action Recognition

Jeonghyeok Do

Munchurl Kim\*

Korea Advanced Institute of Science and Technology

{ehwjdgur0913, mkimee}@kaist.ac.kr

[https://kaist-viclab.github.io/TDSM\\_site](https://kaist-viclab.github.io/TDSM_site)

## Abstract

In zero-shot skeleton-based action recognition (ZSAR), aligning skeleton features with the text features of action labels is essential for accurately predicting unseen actions. ZSAR faces a fundamental challenge in bridging the modality gap between the two-kind features, which severely limits generalization to unseen actions. Previous methods focus on direct alignment between skeleton and text latent spaces, but the modality gaps between these spaces hinder robust generalization learning. Motivated by the success of diffusion models in multi-modal alignment (e.g., text-to-image, text-to-video), we firstly present a diffusion-based skeleton-text alignment framework for ZSAR. Our approach, Triplet Diffusion for Skeleton-Text Matching (TDSM), focuses on cross-alignment power of diffusion models rather than their generative capability. Specifically, TDSM aligns skeleton features with text prompts by incorporating text features into the reverse diffusion process, where skeleton features are denoised under text guidance, forming a unified skeleton-text latent space for robust matching. To enhance discriminative power, we introduce a triplet diffusion (TD) loss that encourages our TDSM to correct skeleton-text matches while pushing them apart for different action classes. Our TDSM significantly outperforms very recent state-of-the-art methods with significantly large margins of 2.36%-point to 13.05%-point, demonstrating superior accuracy and scalability in zero-shot settings through effective skeleton-text matching.

## 1. Introduction

Human action recognition [31, 50, 58, 63] focuses on classifying actions from movements, with RGB videos commonly used due to their accessibility. However, recent advancements in depth sensors [28] and pose estimation algorithms [4, 57] have driven the adoption of skeleton-based

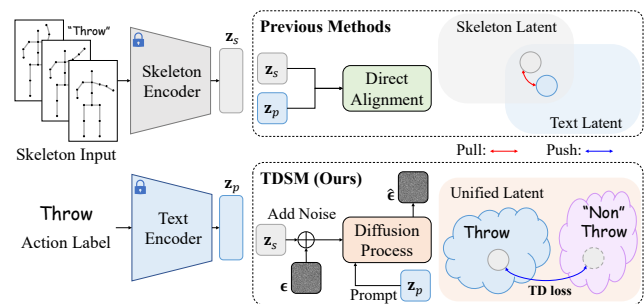


Figure 1. Overview of our Triplet Diffusion for Skeleton-Text Matching (TDSM) pipeline versus previous methods. While the previous methods rely on direct alignment between skeleton and text latent spaces, thus suffering from modality gaps that limit generalization, our TDSM pipeline overcomes this challenge by utilizing the cross-modality alignment power of diffusion models, establishing a more unified and robust skeleton-text representation for effective cross-modal matching.

action recognition. Skeleton data offers several advantages: it captures only human poses without background noise, ensuring a compact representation. Furthermore, 3D skeletons remain invariant to environmental factors such as lighting, background, and camera angles, providing consistent 3D coordinates across conditions [13].

Despite these benefits, the fully supervised skeleton-based action recognition methods [7, 9, 10, 12, 13, 71, 75, 78] tend to perform well, but annotating every possible action is impractical for a large number of possible action classes. In addition, retraining models for new classes incurs a significant cost. So, zero-shot skeleton-based action recognition (ZSAR) [8, 20, 32, 36, 38, 54, 69, 77, 79] addresses this issue by enabling predictions for unseen actions without requiring explicit training data, making it valuable for applications such as surveillance, robotics, and human-computer interaction, where continuous learning is infeasible [16, 64]. Importantly, ZSAR is possible because human actions often share common skeletal movement patterns across related actions. By leveraging these shared

\*Corresponding author.

patterns, ZSAR methods align pre-learned skeleton features with text-based action descriptions, allowing the models to extrapolate from seen actions to unseen ones. This alignment-based approach reinforces the model’s discriminative power, ensuring scalability and reliable zero-shot recognition in real-world scenarios. However, achieving the effective alignment between skeleton data and text features entails significant challenges. While skeleton data captures temporal and spatial motion patterns, the text descriptions for action labels carry high-level semantic information. This modality gap makes it difficult to align their corresponding latent spaces effectively, thus hindering the generalization learning for unseen actions.

Diffusion models [15, 52] have demonstrated strong cross-modal alignment capabilities by incorporating conditioning signals such as text, images, audio, or video to guide the generative process. This conditioning mechanism enables precise *cross-modality alignment*, ensuring that generated outputs adhere closely to the given condition. Inspired by this property, we propose a novel framework: Triplet Diffusion for Skeleton-Prompt Matching (TDSM), which *firstly* adopts diffusion models to ZSAR by conditioning the denoising process on text prompts. Our approach utilizes the reverse diffusion process to implicitly align skeleton and text features within a shared latent space, overcoming the challenges of direct feature space alignment.

Fig. 1 illustrates the key differences between previous methods and our proposed method. The previous methods [8, 20, 32, 36, 38, 54, 69, 77, 79] attempt to directly align skeleton and text features within separate latent spaces. However, this approach struggles with generalization due to the inherent modality gap between skeletal motion and textual semantics. On the other hand, our TDSM leverages a reverse diffusion training scheme to *implicitly* align skeleton features with their corresponding text prompts, producing discriminatively fused representations within a unified latent space. More specifically, our TDSM learns to denoise noisy skeleton features conditioned on the corresponding text prompts, embedding the prompts into a unified skeleton-text latent space to better capture the semantic meaning of action labels. This implicit alignment mitigates the limitations of direct latent space mapping while enhancing robustness. Additionally, we introduce a triplet diffusion (TD) loss, which encourages tighter alignment for correct skeleton-text pairs and pushes apart incorrect ones, further improving the model’s discriminative power. As an additional benefit, the stochastic nature of diffusion process, driven by the random noise added during training, acts as a natural regularization mechanism. This prevents overfitting and enhances the model’s ability to generalize effectively to unseen actions. Our contributions are threefold:

- We *firstly* present a diffusion-based action recognition with zero-shot learning for skeleton inputs, called a

Triplet Diffusion for Skeleton-Text Matching (TDSM) which is the first framework to apply diffusion models and to implicitly align the skeleton features with text prompts (action labels) by fully taking the advantage of excellent text-image correspondence learning in generative diffusion process, thus being able to learn fused discriminative features in a unified latent space.

- We introduce a reformulated triplet diffusion (TD) loss to enhance the model’s discriminative power by ensuring accurate denoising for correct skeleton-text pairs while suppressing it for incorrect pairs.
- Our TDSM *significantly* outperforms the very recent state-of-the-art (SOTA) methods with large margins of 2.36%-point to 13.05%-point across multiple benchmarks, demonstrating scalability and robustness under various seen-unseen split settings.

## 2. Related Work

### 2.1. Zero-shot Skeleton-based Action Recognition

Zero-shot Skeleton-based Action Recognition (ZSAR) aims to recognize human actions from skeleton sequences without requiring labeled training data for unseen action categories. Most of the existing works focus on aligning the skeleton latent space with the text latent space. These approaches can be categorized broadly into VAE-based methods [20, 36, 38, 54] and contrastive learning-based methods [8, 32, 69, 77, 79].

**VAE-based.** The previous work, CADA-VAE [54], leverages VAEs [30] to align skeleton and text latent spaces, ensuring that each modality’s decoder can generate useful outputs from the other’s latent representation. SynSE [20] refines this by introducing separate VAEs for verbs and nouns, improving the structure of the text latent space. MSF [36] extends this approach by incorporating action and motion-level descriptions to enhance alignment. SA-DVAE [38] disentangles skeleton features into semantic-relevant and irrelevant components, aligning text features exclusively with relevant skeleton features for improved performance.

**Contrastive learning-based.** Contrastive learning-based methods align skeleton and text features through positive and negative pairs [5]. SMIE [77] concatenates skeleton and text features, and applies contrastive learning by treating masked skeleton features as positive samples and other actions as negatives. PURLS [79] incorporates GPT-3 [1] to generate text descriptions based on body parts and motion evolution, using cross-attention to align text descriptions with skeleton features. STAR [8] extends this idea with GPT-3.5 [1], generating text descriptions for six distinct skeleton groups, and introduces learnable prompts to enhance alignment. DVTA [32] introduces a dual alignment strategy, performing direct alignment between skeleton and text features, while also generating augmented text features

via cross-attention for improved alignment. InfoCPL [69] strengthens contrastive learning by generating 100 unique sentences per action label, enriching the alignment space.

While most existing methods rely on direct alignment between skeleton and text latent spaces, they often struggle with generalization due to inherent differences between the two modalities. In contrast, our TDSM leverages diffusion models for *alignment* rather than generation. By conditioning the reverse diffusion process on action labels, we guide the denoising of skeleton features to implicitly align them with their corresponding semantic contexts of action labels. This enables more robust skeleton-text matching and improves generalization to unseen actions.

## 2.2. Diffusion Models

Diffusion models have become a fundamental to generative tasks by learning to reverse a noise-adding process for original data recovery. Denoising Diffusion Probabilistic Models (DDPMs) [22] introduced a step-by-step denoising framework, enabling the modeling of complex data distributions and establishing the foundation for diffusion-based generative models. Building on this, Latent Diffusion Models (LDMs) [15, 52] improve computational efficiency by operating in lower-dimensional latent spaces while maintaining high-quality outputs. LDMs have been successful in various generation tasks (e.g., text-to-image, text-to-video, image-to-video), showing the potential of diffusion models for cross-modal alignment tasks. In the reverse diffusion process, LDMs employ a denoising U-Net [53] where text prompts are integrated with image features through cross-attention blocks, effectively guiding the model to align the two modalities. Further extending this line of research, Diffusion Transformers (DiTs) [48] integrate transformer architectures into diffusion processes. In this work, we leverage the aligned fusion capabilities of diffusion models, focusing on the learning process during reverse diffusion rather than their generative power. Specifically, we utilize a DiT-based network as a denoising model, where text prompts guide the denoising of noisy skeleton features. This approach embeds text prompts into the unified latent space in the reverse diffusion process, ensuring robust fusion of the two modalities and enabling effective generalization to unseen actions in zero-shot recognition settings.

**Zero-shot tasks with diffusion models.** Recently, diffusion models have also been extended to zero-shot tasks in RGB-based vision applications, such as semantic correspondence [73], segmentation [2, 59], image captioning [72], and image classification [11, 35]. These approaches often rely on large-scale pretrained diffusion models, such as LDMs [15, 52], trained on datasets like LAION-5B [55] with billions of text-image pairs. In contrast, our approach demonstrates that diffusion models can be effectively applied to smaller, domain-specific tasks, such as skeleton-

based action recognition, without the need for large-scale finetuning. This highlights the versatility of diffusion models beyond large-scale vision-language tasks, providing a practical solution to zero-shot generalization with limited data resources.

## 3. Preliminaries: Diffusion Process

Diffusion models [22, 48, 52] are a class of generative models that progressively denoise a noisy sample to generate a target distribution. The diffusion process [22] consists of two main stages: a forward process and a reverse process.

**Forward diffusion process.** In the forward process, Gaussian noise is incrementally added to the target data  $\mathbf{x}_0$  over discrete timesteps  $t$ , gradually transforming it into a Gaussian noise distribution  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . This can be formulated as:  $q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1-\beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I})$ , where  $q(\mathbf{x}_t|\mathbf{x}_{t-1})$  follows a Markov chain that progressively corrupts  $\mathbf{x}_0$  into noise, and  $\beta_t$  controls the noise schedule. By reparameterizing the forward process,  $\mathbf{x}_t$  can be directly expressed in terms of  $\mathbf{x}_0$  as:  $\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t}\epsilon$ , where  $\bar{\alpha}_t = \prod_{s=1}^t(1-\beta_s)$  controls the noise level at step  $t$ .

**Reverse diffusion process.** The reverse process learns to recover the original data  $\mathbf{x}_0$  from a noisy sample by estimating the denoising step conditioned on previous timesteps as:  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t))$ , where  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$  follows a Gaussian conditional distribution that gradually reconstructs  $\mathbf{x}_0$  by estimating the denoised mean  $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)$  and covariance  $\boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)$ . Here,  $\boldsymbol{\mu}_\theta$  and  $\boldsymbol{\Sigma}_\theta$  are predicted by a neural network trained to approximate the inverse of the forward noise addition, enabling the model to progressively refine  $\mathbf{x}_t$  into a clean representation.

**Objective function.** To train the reverse process, the objective function is derived by minimizing the variational bound on the negative log-likelihood  $\mathbb{E}[-\log p_\theta(\mathbf{x}_0)]$ . It can be simplified by reparameterizing  $\boldsymbol{\mu}_\theta$  as a noise prediction network  $\epsilon_\theta$ , leading to the following loss as:  $\mathcal{L}_{\text{diff}} = \|\epsilon_\theta(\mathbf{x}_t, t) - \epsilon\|_2$ .

**Cross-modality conditioning in diffusion models.** An important property of diffusion models is their ability to incorporate conditioning signals  $\mathbf{c}$  such as text, images, audio, or video to guide the generative process. This conditioning mechanism enables strong *cross-modality alignment*, allowing the models to generate outputs that closely align with the given conditions  $\mathbf{c}$ . To integrate conditioning into the diffusion process, the noise prediction model  $\epsilon_\theta$  is conditioned on  $\mathbf{c}$ , modifying the objective function as follows:  $\mathcal{L}_{\text{diff}} = \|\epsilon_\theta(\mathbf{x}_t, t; \mathbf{c}) - \epsilon\|_2$ . This formulation ensures that the denoising process is guided by the conditions  $\mathbf{c}$ , enforcing the alignment between the generated output and the conditioning signals  $\mathbf{c}$ .

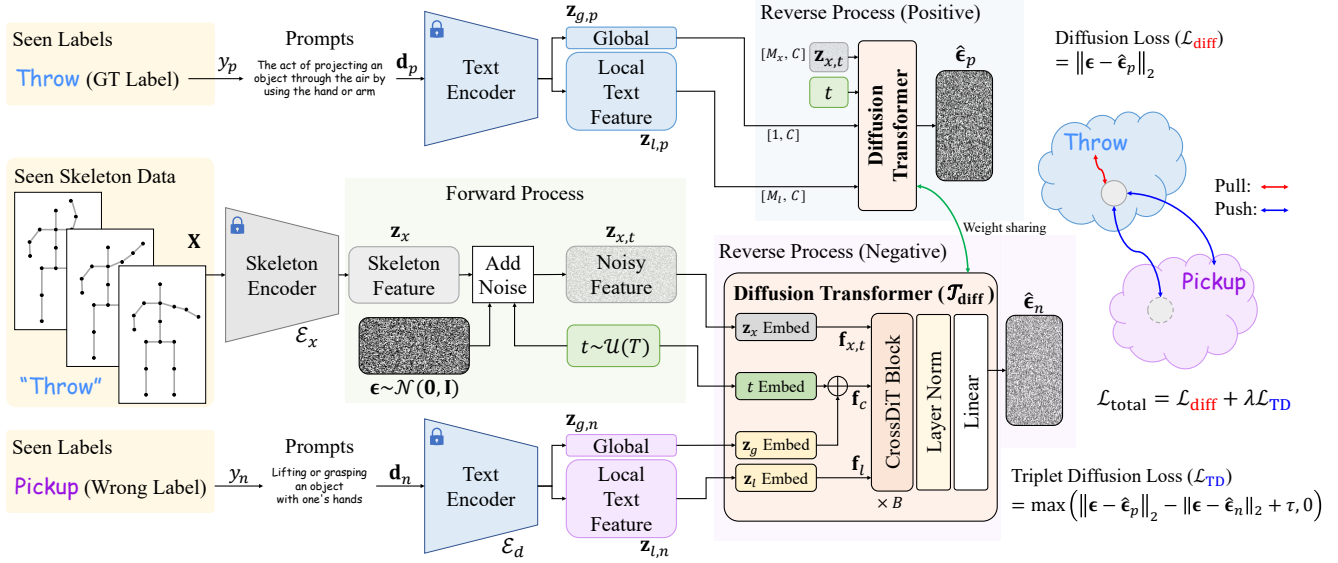


Figure 2. Training framework of our TDSM for zero-shot skeleton-based action recognition.

## 4. Methods

### 4.1. Overview of TDSM

In the training phase, we are given a dataset

$$\mathcal{D}_{\text{train}} = \{(\mathbf{X}_i, y_i)\}_{i=1}^N, y_i \in \mathcal{Y}, \quad (1)$$

where  $\mathbf{X}_i \in \mathbb{R}^{T \times V \times M \times C_{\text{in}}}$  represents a skeleton sequence, and  $y_i$  is the corresponding ground truth label. Each skeleton sequence  $\mathbf{X}_i$  consists of sequence length  $T$ , the number  $V$  of joints, the number  $M$  of actors, and the dimensionality  $C_{\text{in}}$  representing each joint. The label  $y_i$  belongs to the set of seen class labels  $\mathcal{Y}$ . Here,  $N$  denotes the total number of training samples in the seen dataset. In the inference phase, we are provided with a test dataset

$$\mathcal{D}_{\text{test}} = \{(\mathbf{X}_j^u, y_j^u)\}_{j=1}^{N_u}, y_j^u \in \mathcal{Y}_u, \quad (2)$$

where  $\mathbf{X}_j^u \in \mathbb{R}^{T \times V \times M \times C_{\text{in}}}$  denotes skeleton sequences from unseen classes, and  $y_j^u$  are their corresponding labels. In this phase,  $N_u$  represents the total number of test samples from unseen classes. In the zero-shot setting, the seen and unseen label sets are disjoint, i.e.,

$$\mathcal{Y} \cap \mathcal{Y}_u = \emptyset. \quad (3)$$

We train the TDSM using  $\mathcal{D}_{\text{train}}$  and enable it to generalize to unseen classes from  $\mathcal{D}_{\text{test}}$ . By learning a robust discriminative fusion of skeleton features and text descriptions, the model can predict the correct label  $\hat{y}^u \in \mathcal{Y}_u$  for an unseen skeleton sequence  $\mathbf{X}_j^u$  during inference.

Fig. 2 provides an overview of our training framework of TDSM. As detailed in Sec. 4.2, the pretrained skeleton encoder  $\mathcal{E}_x$  and text encoder  $\mathcal{E}_d$  embed the skeleton inputs

$\mathbf{X}$  and prompt input  $\mathbf{d}$  with an action label  $y$  into their respective feature spaces, producing the skeleton feature  $z_x$  and two types of text features: the global text feature  $z_g$  and the local text feature  $z_l$ . The skeleton feature  $z_x$  undergoes the forward process, where noise  $\epsilon$  is added to it. In the reverse process, as described in Sec. 4.3, the Diffusion Transformer  $\mathcal{T}_{\text{diff}}$  serves as the noise prediction network  $\epsilon_\theta$ , predicting the noise  $\hat{\epsilon}$ . The training objective function that ensures the TDSM to learn robust discriminative power is discussed in Sec. 4.3. Finally, Sec. 4.4 explains the strategy used during the inference phase to predict the correct label  $\hat{y}^u$  for unseen actions.

### 4.2. Embedding Skeleton and Prompt Input

Following the LDMs [15, 52], we perform the diffusion process in a compact latent space by projecting both skeleton data and prompt into their respective feature spaces. For skeleton data, we adopt GCNs as the architecture for the skeleton encoder  $\mathcal{E}_x$  that is trained on  $\mathcal{D}_{\text{train}}$  using a cross-entropy loss:

$$\mathcal{L}_{\text{CE}} = - \sum_{k=1}^{|\mathcal{Y}|} \mathbf{y}(k) \log \hat{\mathbf{y}}(k), \quad (4)$$

where  $\hat{\mathbf{y}} = \text{MLP}(\mathcal{E}_x(\mathbf{X}))$  is the predicted class label for a skeleton input  $\mathbf{X}_i$ ,  $|\mathcal{Y}|$  is the number of seen classes and  $\mathbf{y}$  is the one-hot vector of the ground truth label  $y$ . Once trained, the parameters of skeleton encoder  $\mathcal{E}_x$  are frozen and used to generate the skeleton latent space representation  $z_x = \mathcal{E}_x(\mathbf{X})$ . After reshaping the feature for the attention layer,  $z_x$  is represented in  $\mathbb{R}^{M_x \times C}$ , where  $M_x$  is the number of skeleton tokens, and  $C$  is the feature dimension.

For text encoder  $\mathcal{E}_d$ , we leverage the text prompts to capture rich semantic information about the action labels. Each

ground truth (GT) label  $y_p = y$  is associated with a prompt  $\mathbf{d}_p$ , while a randomly selected wrong label (negative sample)  $y_n \in \mathcal{Y} \setminus \{y_p\}$  is assigned a prompt  $\mathbf{d}_n$ . To encode these prompts, we utilize a pretrained text encoder, such as CLIP [27, 51], which provides two types of output features: a global text feature  $\mathbf{z}_g$  and a local text feature  $\mathbf{z}_l$ . The text encoder’s output for a given prompt  $\mathbf{d}$  can be expressed as:

$$[\mathbf{z}_g \mid \mathbf{z}_l] = \mathcal{E}_d(\mathbf{d}), \quad (5)$$

where  $[\cdot \mid \cdot]$  indicates token-wise concatenation,  $\mathbf{z}_g \in \mathbb{R}^{1 \times C}$  is a global text feature, and  $\mathbf{z}_l \in \mathbb{R}^{M_l \times C}$  is a local text feature, with  $M_l$  text tokens. For each GT label (positive sample), the text encoder  $\mathcal{E}_d$  extracts both the global and local text features, denoted as  $\mathbf{z}_{g,p}$  and  $\mathbf{z}_{l,p}$ , respectively. Similarly, for each wrong label (negative sample), the encoder extracts the features  $\mathbf{z}_{g,n}$  and  $\mathbf{z}_{l,n}$ . These four features later guide the diffusion process by conditioning the denoising of noisy skeleton features.

### 4.3. Diffusion Process

Our framework leverages a conditional denoising diffusion process, not to generate data but to learn a discriminative skeleton latent space by fusing skeleton features with text prompts through the reverse diffusion process. Our TDSM is trained to denoise skeleton features such that the resulting latent space becomes discriminative with respect to action labels. Guided by our triplet diffusion (TD) loss, the denoising process conditions on text prompts to strengthen the discriminative fusion of skeleton features and their corresponding prompts. The TD loss encourages correct skeleton-text pairs to be pulled closer in the fused skeleton-text latent space while pushing apart incorrect pairs, enhancing the model’s discriminative power.

**Forward process.** Random Gaussian noise is added to the skeleton feature  $\mathbf{z}_x$  at a random timestep  $t \sim \mathcal{U}(T)$  within total  $T$  steps. At each randomly selected step  $t$ , the noisy feature  $\mathbf{z}_{x,t}$  is generated as:

$$\mathbf{z}_{x,t} = \sqrt{\bar{\alpha}_t} \mathbf{z}_x + \sqrt{1 - \bar{\alpha}_t} \epsilon. \quad (6)$$

**Reverse process.** The Diffusion Transformer  $\mathcal{T}_{\text{diff}}$  predicts noise  $\hat{\epsilon}$  from noisy feature  $\mathbf{z}_{x,t}$ , conditioned on the global and local text features  $\mathbf{z}_g$  and  $\mathbf{z}_l$  at given timestep  $t$ :

$$\hat{\epsilon} = \mathcal{T}_{\text{diff}}(\mathbf{z}_{x,t}, t; \mathbf{z}_g, \mathbf{z}_l). \quad (7)$$

Using the shared weights in  $\mathcal{T}_{\text{diff}}$ , we predict  $\hat{\epsilon}_p$  for positive features ( $\mathbf{z}_{g,p}, \mathbf{z}_{l,p}$ ) and  $\hat{\epsilon}_n$  for negative features ( $\mathbf{z}_{g,n}, \mathbf{z}_{l,n}$ ).  $\mathcal{T}_{\text{diff}}$  builds upon DiT [48], which has been well-validated for cross-modality alignment in image-text tasks. We adopt it to the skeleton-text domain by: (i) reducing the number of blocks/channels to accommodate the relatively small-scale skeleton data, and (ii) incorporating both global and local

text embeddings to enhance skeleton-text alignment.  $\mathcal{T}_{\text{diff}}$  is very detailed in the *Supplementary Material*.

**Triplet diffusion (TD) loss.** The overall training objective combines a diffusion loss and our reformulated TD loss, which is inspired by the conventional triplet loss [23], to promote both effective noise prediction and discriminative alignment (fusion). The total loss is defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{diff}} + \lambda \mathcal{L}_{\text{TD}}, \quad (8)$$

where  $\mathcal{L}_{\text{diff}}$  ensures accurate denoising, and  $\mathcal{L}_{\text{TD}}$  enhances the ability to differentiate between correct and incorrect label predictions. The diffusion loss  $\mathcal{L}_{\text{diff}}$  is given by:

$$\mathcal{L}_{\text{diff}} = \|\epsilon - \hat{\epsilon}_p\|_2, \quad (9)$$

where  $\epsilon$  is a true noise, and  $\hat{\epsilon}_p$  is a predicted noise for the GT text feature. Our triplet diffusion loss  $\mathcal{L}_{\text{TD}}$  is defined as:

$$\mathcal{L}_{\text{TD}} = \max(\|\epsilon - \hat{\epsilon}_p\|_2 - \|\epsilon - \hat{\epsilon}_n\|_2 + \tau, 0), \quad (10)$$

where  $\hat{\epsilon}_n$  is the predicted noise for an incorrect (negative) text feature, and  $\tau$  is a margin parameter.  $\mathcal{L}_{\text{TD}}$  is simple but very effective, encouraging the model to minimize the distance ( $\|\epsilon - \hat{\epsilon}_p\|_2$ ) between true noise  $\epsilon$  and GT prediction  $\hat{\epsilon}_p$  while maximizing the distance ( $\|\epsilon - \hat{\epsilon}_n\|_2$ ) for negative predictions  $\hat{\epsilon}_n$ , which can ensure discriminative fusion of two modalities in the learned skeleton-text latent space.

### 4.4. Inference Phase

Our approach enhances discriminative fusion through the TD loss, which is designed to denoise GT skeleton-text pairs effectively while preventing the fusion of incorrect pairs within the seen dataset. This selective denoising process promotes a robust fusion of skeleton and text features, allowing the model to develop a discriminative feature space that can generalize to unseen action labels.

In inference (Fig. 3), each unseen skeleton sequence  $\mathbf{X}^u$  and its all candidate text prompts are inputted to our TDSM, and the resulting noises for the all candidate text prompts are compared with a fixed GT noise  $\epsilon_{\text{test}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ .  $\mathbf{X}^u$  is first encoded into the skeleton latent space through  $\mathcal{E}_x$  as:

$$\mathbf{z}_x^u = \mathcal{E}_x(\mathbf{X}^u). \quad (11)$$

Each candidate action label  $y_k^u \in \mathcal{Y}_u$  is associated with a prompt  $\mathbf{d}_k^u$  that is processed through  $\mathcal{E}_d$  to extract the global and local text features:

$$[\mathbf{z}_{g,k}^u \mid \mathbf{z}_{l,k}^u] = \mathcal{E}_d(\mathbf{d}_k^u). \quad (12)$$

Next, the forward process is performed using a fixed Gaussian noise  $\epsilon_{\text{test}}$  and a fixed timestep  $t_{\text{test}}$  to generate the noisy skeleton feature as:

$$\mathbf{z}_{x,t}^u = \sqrt{\bar{\alpha}_{t_{\text{test}}}} \mathbf{z}_x^u + \sqrt{1 - \bar{\alpha}_{t_{\text{test}}}} \epsilon_{\text{test}}. \quad (13)$$

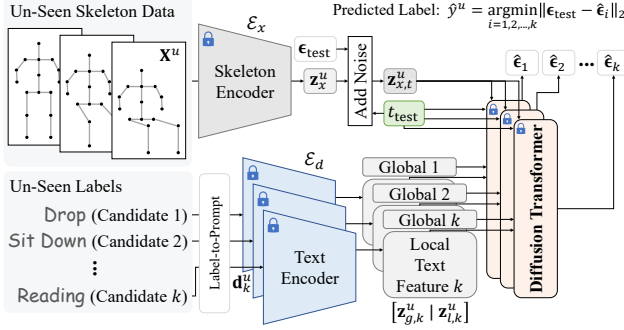


Figure 3. Inference framework of our TDSM for ZSAR.

For  $y_k^u$ , the Diffusion Transformer  $\mathcal{T}_{\text{diff}}$  predicts noise  $\hat{\epsilon}_k$  as:

$$\hat{\epsilon}_k = \mathcal{T}_{\text{diff}}(\mathbf{z}_{x,t}^u, t_{\text{test}}; \mathbf{z}_{g,k}^u, \mathbf{z}_{l,k}^u). \quad (14)$$

The score for  $y_k^u$  is computed as the  $\ell_2$ -norm between  $\epsilon_{\text{test}}$  and  $\hat{\epsilon}_k$ . The predicted label  $\hat{y}^u$  is then the one that minimizes this distance:

$$\hat{y}^u = \arg \min_k \|\epsilon_{\text{test}} - \hat{\epsilon}_k\|_2. \quad (15)$$

This process ensures that the model selects the action label whose text prompt well aligns with the skeleton sequence, enabling accurate zero-shot action recognition for unseen skeleton sequences with unseen action labels. Unlike the generative models that iteratively refine samples, our TDSM performs a *one-step inference* at a fixed timestep, making it efficient and well-suited for discriminative skeleton-text alignment.

## 5. Experiments

### 5.1. Datasets

**NTU RGB+D [56].** NTU RGB+D dataset, referred to as NTU-60, is one of the largest benchmarks for human action recognition, consisting of 56,880 action samples across 60 action classes. It captures 3D skeleton data, depth maps, and RGB videos using Kinect sensors [28], making it a standard for evaluating single- and multi-view action recognition models. The dataset provides a cross-subject split (X-sub) with 40 subjects in total, where 20 subjects are used for training and the remaining 20 for testing. In our experiments, we construct  $\mathcal{D}_{\text{train}}$  from the X-sub training set with seen labels and  $\mathcal{D}_{\text{test}}$  from the X-sub test set with unseen labels, ensuring a fully zero-shot action recognition setting. **NTU RGB+D 120 [42].** NTU RGB+D 120 dataset, referred to as NTU-120, extends the original NTU-60 dataset by adding 60 additional action classes, resulting in a total of 120 classes and 114,480 video samples. The X-sub setting for NTU-120 includes 106 subjects, with 53 used for training and the remaining 53 for testing. We apply the same experimental protocol as NTU-60, using the training set to

form  $\mathcal{D}_{\text{train}}$  with seen labels and the test set to construct  $\mathcal{D}_{\text{test}}$  with unseen labels.

**PKU-MMD [39].** PKU-MMD dataset is a large-scale dataset designed for multi-modality action recognition, offering 3D skeleton data and RGB+D recordings. The dataset contains a total of 66 subjects, where 57 subjects are used for training and the remaining 9 for testing. We follow the cross-subject setting to evaluate the generalization of our framework, constructing  $\mathcal{D}_{\text{train}}$  with seen labels and  $\mathcal{D}_{\text{test}}$  with unseen labels.

### 5.2. Experiment Setup

Our TDSM were implemented in PyTorch [47] and conducted on a single NVIDIA GeForce RTX 3090 GPU. Our model variants were trained for 50,000 iterations, with a warm-up period of 100 steps. We employed the AdamW optimizer [45] with a learning rate of  $1 \times 10^{-4}$  and a weight decay of 0.01. A cosine-annealing scheduler [44] was used to dynamically update the learning rate at each iteration. The batch size was set to 256, but for the TD loss computation, the batch was duplicated for positive and negative samples, resulting in an effective batch size of 512. Through empirical validation, the loss weight  $\lambda$  and margin  $\tau$  were both set to 1.0. The diffusion process was trained with a total timestep of  $T = 50$ . For inference, the optimal timestep  $t_{\text{test}} = 25$  was selected based on accuracy trends across datasets (Fig. 4). For the SynSE [20] and PURLS [79] seen and unseen split settings, we used the Shift-GCN [9] architecture as our skeleton encoder. In the SMIE [77] split setting, we employed the ST-GCN [71] structure. We adopted Shift-GCN and ST-GCN because the majority of previous ZSAR studies use them, ensuring a fair comparison without a confounding factor of potentially more powerful encoders. For fair comparison, we used the same text prompts employed in existing works. Across all tasks, we adopted the text encoder from CLIP [27, 51] to transform text prompts into latent representations.

### 5.3. Performance Evaluation

**Evaluation on SynSE [20] and PURLS [79] benchmarks.** Table 1 presents the performance comparison on the SynSE and PURLS benchmark splits across the NTU-60 and NTU-120 datasets. The SynSE benchmark focuses on standard settings, offering 55/5 and 48/12 splits on NTU-60, and 110/10 and 96/24 splits on NTU-120. These settings assess the model’s ability to generalize across typical seen-unseen splits. In contrast, the PURLS benchmark presents more extreme cases with 40/20 and 30/30 splits on NTU-60, and 80/40 and 60/60 splits on NTU-120, introducing higher levels of complexity by increasing the proportion of unseen labels in the test set. To account for the stochastic nature of noise during inference, we averaged the 10 runs with different Gaussian noise realizations. As shown

Methods	Publications	NTU-60 (Acc, %)				NTU-120 (Acc, %)			
		55/5 split	48/12 split	40/20 split	30/30 split	110/10 split	96/24 split	80/40 split	60/60 split
ReViSE [26]	ICCV 2017	53.91	17.49	24.26	14.81	55.04	32.38	19.47	8.27
JPoSE [67]	ICCV 2019	64.82	28.75	20.05	12.39	51.93	32.44	13.71	7.65
CADA-VAE [54]	CVPR 2019	76.84	28.96	16.21	11.51	59.53	35.77	10.55	5.67
SynSE [20]	ICIP 2021	75.81	33.30	19.85	12.00	62.69	38.70	13.64	7.73
SMIE [77]	ACM MM 2023	77.98	40.18	-	-	65.74	45.30	-	-
PURLS [79]	CVPR 2024	79.23	40.99	<u>31.05</u>	<u>23.52</u>	<u>71.95</u>	<u>52.01</u>	<u>28.38</u>	<u>19.63</u>
SA-DVAE [38]	ECCV 2024	<u>82.37</u>	41.38	-	-	68.77	46.12	-	-
STAR [8]	ACM MM 2024	81.40	<u>45.10</u>	-	-	63.30	44.30	-	-
<b>TDSM (Ours)</b>	-	<b>86.49</b>	<b>56.03</b>	<b>36.09</b>	<b>25.88</b>	<b>74.15</b>	<b>65.06</b>	<b>36.95</b>	<b>27.21</b>

Table 1. Top-1 accuracy results of various zero-shot skeleton-based action recognition (ZSAR) methods evaluated on the SynSE and PURLS benchmarks for the NTU-60 and NTU-120 datasets. Each split is denoted as X/Y, where X represents the number of seen classes and Y the number of unseen classes. The results in **red** highlight the best-performing model, while those in **blue** indicate the second-best. For our TDSM framework, the reported accuracy is the average value obtained from 10 trials, each with different Gaussian noise.

Methods	NTU-60 (Acc, %)	NTU-120 (Acc, %)	PKU-MMD (Acc, %)
	55/5 split	110/10 split	46/5 split
ReViSE [26]	60.94	44.90	59.34
JPoSE [67]	59.44	46.69	57.17
CADA-VAE [54]	61.84	45.15	60.74
SynSE [20]	64.19	47.28	53.85
SMIE [77]	65.08	46.40	60.83
SA-DVAE [38]	<u>84.20</u>	<u>50.67</u>	66.54
STAR [8]	77.50	-	<u>70.60</u>
<b>TDSM (Ours)</b>	<b>88.88</b>	<b>69.47</b>	<b>70.76</b>

Table 2. Top-1 accuracy results of various ZSAR methods evaluated on the NTU-60, NTU-120, and PKU-MMD datasets under the SMIE benchmark. The reported values are the average performance across three splits.

$\mathcal{L}_{diff}$	$\mathcal{L}_{TD}$	NTU-60 (Acc, %)		NTU-120 (Acc, %)	
		55/5 split	48/12 split	110/10 split	96/24 split
✓		79.87	53.03	72.44	57.65
	✓	80.90	54.36	70.73	60.95
✓	✓	<b>86.49</b>	<b>56.03</b>	<b>74.15</b>	<b>65.06</b>

Table 3. Ablation study on loss function configurations. The results compare models trained with only the diffusion loss  $\mathcal{L}_{diff}$ , only the triplet diffusion loss  $\mathcal{L}_{TD}$ , and their combination.

in Table 1, our TDSM significantly outperforms the very recent state-of-the-art results across all benchmark splits, demonstrating superior generalization and robustness for various splits. Specifically, TDSM outperforms the existing methods on both standard (SynSE) and extreme (PURLS) settings, with 4.12%-point, 9.93%-point, 5.04%-point, and 2.36%-point improvements in top-1 accuracy on the NTU-60 55/5, 48/12, 40/20, and 30/30 splits, respectively, compared to the second best models. Also compared to the second best models, our TDSM attains 2.20%-point, 13.05%-point, 8.57%-point, and 7.58%-point accuracy improvements on the NTU-120 splits, further validating its scalability to larger datasets and more complex unseen classes.

**Evaluation on SMIE [77] benchmark.** The SMIE benchmark provides three distinct splits to evaluate the generalization capability of models across different sets of unseen labels. Each split ensures that unseen labels do not over-

Global $\mathbf{z}_g$	Local $\mathbf{z}_l$	NTU-60 (Acc, %)		NTU-120 (Acc, %)	
		55/5 split	48/12 split	110/10 split	96/24 split
✓		83.41	51.50	70.14	61.90
	✓	83.33	52.63	69.95	62.10
✓	✓	<b>86.49</b>	<b>56.03</b>	<b>74.15</b>	<b>65.06</b>

Table 4. Ablation study on text feature types. The results compare models trained with only  $\mathbf{z}_g$ , only  $\mathbf{z}_l$ , and their combination.

lap with seen ones, thereby rigorously testing the model’s ability to recognize new classes without prior exposure. For fair comparison, the reported performance is the average of the three splits. In this benchmark, as shown in Table 2, our TDSM outperforms the other methods, demonstrating strong generalization across all evaluated datasets.

## 5.4. Ablation Studies

**Effect of loss function design.** The ablation study shown in Table 3 evaluates the impact of combining the diffusion loss  $\mathcal{L}_{diff}$  and the triplet diffusion (TD) loss  $\mathcal{L}_{TD}$  on the model’s performance. The results demonstrate that leveraging both losses yields superior performance compared to using either one alone. Specifically, when only  $\mathcal{L}_{diff}$  is employed, the model focuses on accurately denoising the skeleton features conditioned on the text prompts but may lack sufficient discriminative power between similar actions. Conversely, applying only  $\mathcal{L}_{TD}$  enhances discriminative fusion but without ensuring optimal noise prediction. The combination of both losses strikes a balance, ensuring discriminative fusion, which resulting in the highest performance across all evaluated splits.

**Contribution of global and local text features.** Our TDSM framework employs two types of text features for skeleton-text matching: a global text feature  $\mathbf{z}_g$  that encodes the entire sentence as a single token, and local text features  $\mathbf{z}_l$  that preserve token-level details for each word in the sentence. As shown in Table 4, combining the global and local text features achieves the best performance, outperforming the models that use either feature independently. The global text feature provides overall discriminative power




Figure 4. Effect of varying inference timesteps  $t_{\text{test}}$  across multiple datasets. Each plot shows the top-1 accuracy trend on the NTU-60 and NTU-120 datasets under different splits. The solid red line represents the average accuracy of our method, with the shaded orange area indicating the variation in accuracy across 10 different random Gaussian noise instances. Dashed blue line corresponds to the second-best method in each benchmark.

by capturing high-level semantics of the action description, enabling robust matching across diverse action categories. Meanwhile, the local text features retain finer details that are effective for distinguishing subtle differences between semantically similar actions.

**Impact of total timesteps  $T$ .** The results of the ablation study on diffusion timesteps are shown in Table 5. We observe that the choice of  $T$  has a significant impact on performance across all datasets and splits. When  $T$  is too small, the model tends to overfit, as the problem becomes too simple, limiting the diversity in noise added to the skeleton features during training. On the other hand, too large  $T$  values introduces diverse noise strengths, making it challenging for the model to denoise effectively, which deteriorates performance. The best  $T$  is found empirically with  $T = 50$ , striking a balance between maintaining a challenging task and avoiding overfitting.

**Effect of random Gaussian noise.** To examine the role of Gaussian noise in training, we conducted ablation studies by using fixed Gaussian noise during both training and inference, instead of introducing new random noise at each training step. As shown in Table 6, using fixed Gaussian noise overly simplifies the learning process, causing the network to overfit specific noise patterns and reducing its generalization ability. In contrast, introducing random Gaussian noise at each step increases variability in the learning process, acting as a regularization mechanism that prevents overfitting. This enhances model robustness and improves the alignment between skeleton features and text prompts.

**Impact of timestep  $t_{\text{test}}$  and noise  $\epsilon_{\text{test}}$  in inference.** We conducted experiments across different test timesteps  $t_{\text{test}} \in [0, 50]$  and observed the accuracy trends on multiple datasets, as illustrated in Fig. 4. Based on these observations, we set  $t_{\text{test}} = 25$  for all experiments. To examine the impact of noise during inference, we repeated experiments with 10 different random Gaussian noise samples. In Fig. 4, the shaded orange regions in the graphs depict the variances in top-1 accuracy due to noise, while the red lines represent the average accuracy. The blue dashed lines in-

Total $T$	NTU-60 (Acc, %)		NTU-120 (Acc, %)	
	55/5 split	48/12 split	110/10 split	96/24 split
1	85.03	44.10	69.91	60.35
10	84.51	50.89	69.97	62.04
50	<b>86.49</b>	<b>56.03</b>	<b>74.15</b>	<b>65.06</b>
100	83.48	56.27	71.05	64.57
500	81.34	53.43	71.93	60.81

Table 5. Ablation study on the impact of total timesteps  $T$  in the training of the diffusion process.

Gaussian noise $\epsilon$	NTU-60 (Acc, %)		NTU-120 (Acc, %)	
	55/5 split	48/12 split	110/10 split	96/24 split
Fixed	76.40	44.25	64.01	52.21
Random	<b>86.49</b>	<b>56.03</b>	<b>74.15</b>	<b>65.06</b>

Table 6. Ablation study on the effect of noise  $\epsilon$  during training.

dicating the second-best method’s accuracy for comparison. Our analysis shows that while noise variations can cause up to a  $\pm 2.5\%$ -point changes in top-1 accuracy at  $t_{\text{test}} = 25$ , our TDSM shows consistently outperforming the state-of-the-art methods regardless of noise levels.

## 6. Conclusion

Our TDSM is the first framework to apply diffusion models to zero-shot skeleton-based action recognition. The selective denoising process promotes a robust fusion of skeleton and text features, allowing the model to develop a discriminative feature space that can generalize to unseen action labels. Also, our approach enhances discriminative fusion through the TD loss which is designed to denoise GT skeleton-text pairs effectively while preventing the fusion of incorrect pairs within the seen dataset. Extensive experiments show that our TDSM significantly outperforms the very recent SOTA models with large margins for various benchmark datasets.

**Acknowledgements.** This work was supported by IITP grant funded by the Korea government(MSIT) (No.RS2022-00144444, Deep Learning Based Visual Representational Learning and Rendering of Static and Dynamic Scenes).

## References

- [1] Tom B Brown. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020. 2, 12
- [2] Ryan Burgert, Kanchana Ranasinghe, Xiang Li, and Michael S Ryoo. Peekaboo: Text to image diffusion models are zero-shot segmentors. *arXiv preprint arXiv:2211.13224*, 2022. 3
- [3] Dongqi Cai, Yangyuxuan Kang, Anbang Yao, and Yurong Chen. Ske2grid: Skeleton-to-grid representation learning for action recognition. In *International Conference on Machine Learning*, 2023. 13
- [4] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. 1
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 2
- [6] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18000–18010, 2023. 11
- [7] Yuxin Chen, Ziqi Zhang, Chunfeng Yuan, Bing Li, Ying Deng, and Weiming Hu. Channel-wise topology refinement graph convolution for skeleton-based action recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 13359–13368, 2021. 1, 13
- [8] Yang Chen, Jingcai Guo, Tian He, and Ling Wang. Fine-grained side information guided dual-prompts for zero-shot skeleton action recognition. *arXiv preprint arXiv:2404.07487*, 2024. 1, 2, 7
- [9] Ke Cheng, Yifan Zhang, Xiangyu He, Weihang Chen, Jian Cheng, and Hanqing Lu. Skeleton-based action recognition with shift graph convolutional network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 183–192, 2020. 1, 6, 12, 13
- [10] Hyung-gun Chi, Myoung Hoon Ha, Seunggeun Chi, Sang Wan Lee, Qixing Huang, and Karthik Ramani. Infogcn: Representation learning for human skeleton-based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20186–20196, 2022. 1, 13
- [11] Kevin Clark and Priyank Jaini. Text-to-image diffusion models are zero shot classifiers. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [12] Jeonghyeok Do and Munchurl Kim. Skateformer: Skeletal-temporal transformer for human action recognition. *arXiv preprint arXiv:2403.09508*, 2024. 1, 13
- [13] Haodong Duan, Yue Zhao, Kai Chen, Dahua Lin, and Bo Dai. Revisiting skeleton-based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2969–2978, 2022. 1, 13
- [14] Haodong Duan, Mingze Xu, Bing Shuai, Davide Modolo, Zhuowen Tu, Joseph Tighe, and Alessandro Bergamo. Skeletr: Towards skeleton-based action recognition in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13634–13644, 2023. 13
- [15] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. 2, 3, 4, 11
- [16] Valter Estevam, Helio Pedrini, and David Menotti. Zero-shot action recognition in videos: A survey. *Neurocomputing*, 439:159–175, 2021. 1
- [17] Lin Geng Foo, Tianjiao Li, Hossein Rahmani, and Jun Liu. Action detection via an image diffusion process. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18351–18361, 2024. 11
- [18] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems*, 26, 2013. 10
- [19] Jia Gong, Lin Geng Foo, Zhipeng Fan, Qihong Ke, Hossein Rahmani, and Jun Liu. Diffpose: Toward more reliable 3d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13041–13051, 2023. 11
- [20] Pranay Gupta, Divyanshu Sharma, and Ravi Kiran Sarvadevabhatla. Syntactically guided generative embeddings for zero-shot skeleton action recognition. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 439–443. IEEE, 2021. 1, 2, 6, 7
- [21] Jing He, Haodong Li, Wei Yin, Yixun Liang, Leheng Li, Kaiqiang Zhou, Hongbo Liu, Bingbing Liu, and Yingcong Chen. Lotus: Diffusion-based visual foundation model for high-quality dense prediction. *arXiv preprint arXiv:2409.18124*, 2024. 9
- [22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3
- [23] Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In *Similarity-based pattern recognition: third international workshop, SIMBAD 2015, Copenhagen, Denmark, October 12-14, 2015. Proceedings 3*, pages 84–92. Springer, 2015. 5
- [24] Karl Holmquist and Bastian Wandt. Diffpose: Multi-hypothesis human pose estimation using diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15977–15987, 2023. 11
- [25] Yiheng Huang, Hui Yang, Chuanchen Luo, Yuxi Wang, Shibiao Xu, Zhaoxiang Zhang, Man Zhang, and Junran Peng. Stablemofusion: Towards robust and efficient diffusion-based motion generation framework. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 224–232, 2024. 11
- [26] Yao-Hung Hubert Tsai, Liang-Kang Huang, and Ruslan Salakhutdinov. Learning robust visual-semantic embeddings. In *Proceedings of the IEEE International conference on Computer Vision*, pages 3571–3580, 2017. 7, 10
- [27] Gabriel Ilharco, Mitchell Wortsman, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok

- Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Open clip, 2021. 5, 6, 12
- [28] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 1, 6, 9
- [29] Qihong Ke, Mohammed Bennamoun, Senjian An, Ferdous Sohel, and Farid Boussaid. Learning clip representations for skeleton-based 3d action recognition. *IEEE Transactions on Image Processing*, 27(6):2842–2855, 2018. 13
- [30] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [31] Yu Kong and Yun Fu. Human action recognition and prediction: A survey. *International Journal of Computer Vision*, 130(5):1366–1401, 2022. 1
- [32] Jidong Kuang, Hongsong Wang, Chaolei Han, and Jie Gui. Zero-shot skeleton-based action recognition with dual visual-text alignment. *arXiv preprint arXiv:2409.14336*, 2024. 1, 2
- [33] Jungho Lee, Minhyeok Lee, Suhwan Cho, Sungmin Woo, Sungjun Jang, and Sangyoun Lee. Leveraging spatio-temporal dependency for skeleton-based action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10255–10264, 2023. 13
- [34] Jungho Lee, Minhyeok Lee, Dogyoon Lee, and Sangyoun Lee. Hierarchically decomposed graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10444–10453, 2023. 13
- [35] Alexander C Li, Mihir Prabhudesai, Shivam Duggal, Ellis Brown, and Deepak Pathak. Your diffusion model is secretly a zero-shot classifier. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2206–2217, 2023. 3
- [36] Ming-Zhe Li, Zhen Jia, Zhang Zhang, Zhanyu Ma, and Liang Wang. Multi-semantic fusion model for generalized zero-shot skeleton-based action recognition. In *International Conference on Image and Graphics*, pages 68–80. Springer, 2023. 1, 2
- [37] Shuai Li, Wanqing Li, Chris Cook, Ce Zhu, and Yanbo Gao. Independently recurrent neural network (indrnn): Building a longer and deeper rnn. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5457–5466, 2018. 13
- [38] Sheng-Wei Li, Zi-Xiang Wei, Wei-Jie Chen, Yi-Hsin Yu, Chih-Yuan Yang, and Jane Yung-jen Hsu. Sa-dvae: Improving zero-shot skeleton-based action recognition by disentangled variational autoencoders. *arXiv preprint arXiv:2407.13460*, 2024. 1, 2, 7
- [39] Chunhui Liu, Yueyu Hu, Yanghao Li, Sijie Song, and Jiaying Liu. Pku-mmd: A large scale benchmark for continuous multi-modal human action understanding. *arXiv preprint arXiv:1703.07475*, 2017. 6, 12
- [40] Hongjie Liu, Yingchun Niu, Kun Zeng, Chun Liu, Mengjie Hu, and Qing Song. Beyond-skeleton: Zero-shot skeleton action recognition enhanced by supplementary rgb visual information. *Expert Systems with Applications*, 273:126814, 2025. 9, 10
- [41] Jun Liu, Amir Shahroudy, Dong Xu, and Gang Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, pages 816–833. Springer, 2016. 13
- [42] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C Kot. Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2684–2701, 2019. 6
- [43] Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. Disentangling and unifying graph convolutions for skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 143–152, 2020. 13
- [44] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 6
- [45] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [46] Yunsheng Pang, Qihong Ke, Hossein Rahmani, James Bailey, and Jun Liu. Igformer: Interaction graph transformer for skeleton-based human interaction recognition. In *European Conference on Computer Vision*, pages 605–622. Springer, 2022. 13
- [47] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 6
- [48] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. 3, 5, 11
- [49] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. 11
- [50] Liliana Lo Presti and Marco La Cascia. 3d skeleton-based human action classification: A survey. *Pattern Recognition*, 53:130–147, 2016. 1
- [51] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 5, 6, 12
- [52] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2, 3, 4
- [53] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted*

- intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015. 3, 11, 12
- [54] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero-and few-shot learning via aligned variational autoencoders. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8247–8255, 2019. 1, 2, 7
- [55] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 3
- [56] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1010–1019, 2016. 6
- [57] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5693–5703, 2019. 1
- [58] Zehua Sun, Qihong Ke, Hossein Rahmani, Mohammed Bennamoun, Gang Wang, and Jun Liu. Human action recognition from various data modalities: A review. *IEEE transactions on pattern analysis and machine intelligence*, 2022. 1
- [59] Junjiao Tian, Lavisha Aggarwal, Andrea Colaco, Zsolt Kira, and Mar Gonzalez-Franco. Diffuse attend and segment: Un-supervised zero-shot segmentation using stable diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3554–3563, 2024. 3
- [60] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 11
- [61] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 11
- [62] Lei Wang and Piotr Koniusz. 3mformer: Multi-order multi-mode transformer for skeletal action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5620–5631, 2023. 13
- [63] Lei Wang, Du Q Huynh, and Piotr Koniusz. A comparative review of recent kinect-based action recognition algorithms. *IEEE Transactions on Image Processing*, 29:15–28, 2019. 1
- [64] Wei Wang, Vincent W Zheng, Han Yu, and Chunyan Miao. A survey of zero-shot learning: Settings, methods, and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–37, 2019. 1
- [65] Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Recovering realistic texture in image super-resolution by deep spatial feature transform. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 606–615, 2018. 11
- [66] Mingjie Wei, Xuemei Xie, Yutong Zhong, and Guangming Shi. Learning pyramid-structured long-range dependencies for 3d human pose estimation. *IEEE Transactions on Multimedia*, 2025. 11
- [67] Michael Wray, Diane Larlus, Gabriela Csurka, and Dima Damen. Fine-grained action retrieval through multiple parts-of-speech embeddings. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 450–459, 2019. 7
- [68] Wangmeng Xiang, Chao Li, Yuxuan Zhou, Biao Wang, and Lei Zhang. Generative action description prompts for skeleton-based action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10276–10285, 2023. 13
- [69] Haojun Xu, Yan Gao, Jie Li, and Xinbo Gao. An information compensation framework for zero-shot skeleton-based action recognition. *arXiv preprint arXiv:2406.00639*, 2024. 1, 2, 3
- [70] Kailin Xu, Fanfan Ye, Qiaoyong Zhong, and Di Xie. Topology-aware convolutional neural network for efficient skeleton-based action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2866–2874, 2022. 13
- [71] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. 1, 6, 12, 13
- [72] Chenglin Yang, Siyuan Qiao, Yuan Cao, Yu Zhang, Tao Zhu, Alan Yuille, and Jiahui Yu. Ig captioner: Information gain captioners are strong zero-shot classifiers. *arXiv preprint arXiv:2311.17072*, 2023. 3
- [73] Junyi Zhang, Charles Herrmann, Junhwa Hur, Luisa Polania Cabrera, Varun Jampani, Deqing Sun, and Ming-Hsuan Yang. A tale of two features: Stable diffusion complements dino for zero-shot semantic correspondence. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [74] Pengfei Zhang, Cuiling Lan, Junliang Xing, Wenjun Zeng, Jianru Xue, and Nanning Zheng. View adaptive recurrent neural networks for high performance human action recognition from skeleton data. In *Proceedings of the IEEE international conference on computer vision*, pages 2117–2126, 2017. 13
- [75] Huanyu Zhou, Qingjie Liu, and Yunhong Wang. Learning discriminative representations for skeleton based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10608–10617, 2023. 1, 13
- [76] Yuxuan Zhou, Chao Li, Zhi-Qi Cheng, Yifeng Geng, Xuansong Xie, and Margret Keuper. Hypergraph transformer for skeleton-based action recognition. *arXiv preprint arXiv:2211.09590*, 2022. 13
- [77] Yujie Zhou, Wenwen Qiang, Anyi Rao, Ning Lin, Bing Su, and Jiaqi Wang. Zero-shot skeleton-based action recognition via mutual information estimation and maximization. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 5302–5310, 2023. 1, 2, 6, 7, 9, 10
- [78] Yuxuan Zhou, Xudong Yan, Zhi-Qi Cheng, Yan Yan, Qi Dai, and Xian-Sheng Hua. Blockgc: Redefine topology awareness for skeleton-based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2049–2058, 2024. 1, 13

- [79] Anqi Zhu, Qihong Ke, Mingming Gong, and James Bailey. Part-aware unified representation of language and skeleton for zero-shot action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18761–18770, 2024. [1](#), [2](#), [6](#), [7](#), [9](#), [10](#)
- [80] Wentao Zhu, Cuiling Lan, Junliang Xing, Wenjun Zeng, Yanghao Li, Li Shen, and Xiaohui Xie. Co-occurrence feature learning for skeleton based action recognition using regularized deep lstm networks. In *Proceedings of the AAAI conference on artificial intelligence*, 2016. [13](#)