

# Dual Conditioned Motion Diffusion for Pose-Based Video Anomaly Detection

Hongsong Wang<sup>1,2</sup>, Andi Xu<sup>3</sup>, Pinle Ding<sup>3</sup>, Jie Gui<sup>3,4,5</sup>\*

<sup>1</sup>School of Computer Science and Engineering, Southeast University, Nanjing 210096, China

<sup>2</sup>Key Laboratory of New Generation Artificial Intelligence Technology and Its Interdisciplinary Applications (Southeast University), Ministry of Education, China

<sup>3</sup>School of Cyber Science and Engineering, Southeast University, Nanjing 210096, China

<sup>4</sup>Engineering Research Center of Blockchain Application, Supervision And Management (Southeast University), Ministry of Education, China

<sup>5</sup> Purple Mountain Laboratories, Nanjing 210000, China  
{hongsongwang, andixu, pinleding, guijie}@seu.edu.cn

## Abstract

Video Anomaly Detection (VAD) is essential for computer vision research. Existing VAD methods utilize either reconstruction-based or prediction-based frameworks. The former excels at detecting irregular patterns or structures, whereas the latter is capable of spotting abnormal deviations or trends. We address pose-based video anomaly detection and introduce a novel framework called Dual Conditioned Motion Diffusion (DCMD), which enjoys the advantages of both approaches. The DCMD integrates conditioned motion and conditioned embedding to comprehensively utilize the pose characteristics and latent semantics of observed movements, respectively. In the reverse diffusion process, a motion transformer is proposed to capture potential correlations from multi-layered characteristics within the spectrum space of human motion. To enhance the discriminability between normal and abnormal instances, we design a novel United Association Discrepancy (UAD) regularization that primarily relies on a Gaussian kernel-based time association and a self-attention-based global association. Finally, a mask completion strategy is introduced during the inference stage of the reverse diffusion process to enhance the utilization of conditioned motion for the prediction branch of anomaly detection. Extensive experiments on four datasets demonstrate that our method dramatically outperforms state-of-the-art methods and exhibits superior generalization performance.

**Code** — <https://github.com/guijie/DCMD-main>

## Introduction

Video Anomaly Detection (VAD) is an essential topic in computer vision and security applications. Different from human action understanding (2018; 2023; 2024; 2025), VAD enables prompt identification of abnormal occurrences, such as human actions, accidents, and illnesses. Anomalies are generally characterized as uncommon, unexpected, or unusual phenomena that exhibit significant deviations from normality. Conversely, normality is defined as what is expected and regularly encountered. Detecting video anomalies can be challenging as these events occur infrequently and belong to various categories. Labeling data is

\*Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

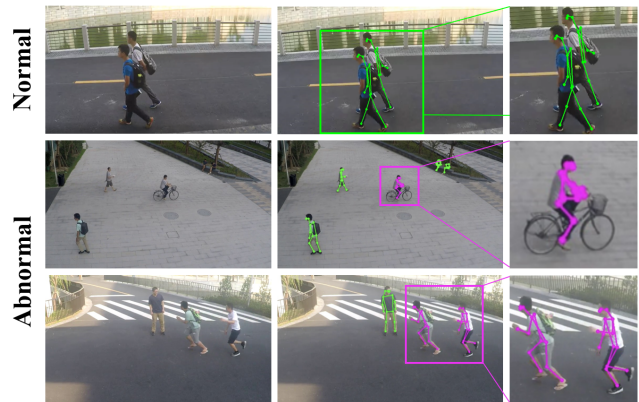


Figure 1: Illustration of video anomaly detection with normal instances (green) and abnormal instances (red).

costly and time-consuming, making collecting all possible abnormal samples impractical for fully supervised learning approaches. Consequently, anomaly detection problems are typically treated as One Class Classification (OCC) (2020). Only normal data is utilized during model training, and data that deviates significantly from the normal pattern is identified as an anomaly during testing (2021).

There has been a growing trend toward detecting abnormal events or behaviors by analyzing human poses extracted from video frames (2023), as illustrated in Figure 1. Utilizing skeletons to depict human motions in videos is a highly effective approach to protecting privacy and circumventing the limitations of appearance-based attributes. Besides privacy protection, skeletal features are compact, well-structured, and highly descriptive of human motion (2019).

Existing VAD works are mainly categorized into reconstruction-based and prediction-based methods. Reconstruction-based methods (2021; 2020) first compress the input data into a low-dimensional representation and then recover the original data from this representation. By comparing the reconstruction error between the original data and the recovered data, anomalies such as irregular patterns or structures in the data can be detected. Prediction-based methods (2017) focus on predicting future frames or events

based on history data, and anomalies are indicated when future data deviates from the predicted trends. These methods are effective in modeling and uncovering the temporal connections between successive frames. Nevertheless, they may be prone to noise (2020).

Recently, diffusion models have shown great success in generating high-quality samples (Ho, Jain, and Abbeel 2020). One of the outstanding capabilities of diffusion models is the flexibility to handle a wide range of input conditions that can encompass different types of data, enabling customized generation tasks (Ho and Salimans 2021). Due to the inherent multimodality of human motion and the diversity of both normal and abnormal patterns, diffusion models are naturally well-suited for modeling human motion and detecting anomalies. However, simply applying diffusion models for VAD suffers inherent limitations of reconstruction-based or prediction-based approaches.

To address the above issues, we present a unified framework that seamlessly combines the advantages of both reconstruction-based and prediction-based approaches for pose-based VAD. The reconstruction is based on an auto-encoder architecture, whereas the prediction utilizes a diffusion model. We propose a novel Dual Conditioned Motion Diffusion (DCMD) that seamlessly integrates conditioned motion and conditioned embedding, enabling the exploitation of both pose characteristics and latent semantics of observed movements. During the reverse diffusion process, we introduce a motion transformer specifically designed to extract potential correlations from multi-layered spectrum features of human motion. In addition to reconstruction and prediction losses, we devise a United Association Discrepancy (UAD) regularization that leverages Gaussian kernel-based time association and self-attention-based global association. Furthermore, during the inference stage, we employ a mask completion strategy to bolster the utilization of conditioned motion for prediction-based anomaly detection. Experiments conducted on popular human-related anomaly detection datasets demonstrate the superior performance of our proposed method compared to state-of-the-art approaches.

Our main contributions are summarised below:

- We introduce a novel framework that seamlessly integrates reconstruction-based and prediction-based methods for video anomaly detection, leveraging the strengths of both approaches.
- We propose a Dual Conditioned Motion Diffusion (DCMD), which incorporates both conditioned motion and conditioned embedding in a diffusion-based model.
- We propose a motion transformer for anomaly detection with a novel regularization that uses both time association and global association to improve the discriminability between normal and abnormal instances.
- We present a mask completion method during the denoising process of diffusion, enabling more effective utilization of observed motions while predicting future motions for anomaly detection.

## Related Work

We summarize previous work on Video Anomaly Detection (VAD) from three aspects: reconstruction-based, prediction-based, and pose-based.

**Reconstruction-Based Methods:** Reconstruction-based methods comprise two components: an encoder and a decoder. The encoder compresses the input frame into a low-dimensional feature while the decoder reconstructs the output from this compressed representation. Reconstruction error is a criterion for distinguishing between normal and abnormal events. Luo et al. (2019) introduce a novel deep neural network architecture that leverages sparse coding, incorporates a temporal coherence term to preserve similarity between similar frames, and employs a stacked recurrent neural network to optimize sparse coefficients for achieving real-time anomaly detection. Li et al. (2021) present a two-stream network designed to capture the visual and motion characteristics of typical events in videos. To encode the scene, objects as well as motion information, Chang et al. (2020) propose a novel deep k-mean clustered convolutional self-encoder architecture. There are also some methods utilizing architectures other than convolutional neural networks. Gong et al. (2019) introduce a memory-enhanced autoencoder that can distinguish new test samples by memorizing prototypical normal data elements during training. Doshi et al. (2022) propose a dual-stage approach combining deep learning with a kNN-based RNN to overcome forgetting in end-to-end models, enabling efficient continual learning. Zhong et al. (2022b; 2022c) design a cascade reconstruction model and spatio-temporal autoencoder.

**Prediction-Based Methods:** In prediction-based methods, the models are often employed to forecast the next frame by employing a sequence of preceding frames as inputs. Prediction-based methods are more effective in analyzing spatio-temporal patterns between frames than reconstruction-based methods. Liu et al. (2017) pioneer prediction-based VAD, using U-Net to predict future frames and incorporating motion constraints for consistent optical flow. Zhou et al. (2020) adopt a similar network architecture and propose an attention-driven loss algorithm to address the challenge of imbalanced foreground targets and static backgrounds in anomaly detection videos. Doshi et al. (2021) present an online anomaly detection method that consists of a feature extraction module and a statistical decision-making module. To consider both spatio-temporal characteristics, Lee et al. (2018) utilize a spatio-temporal generator that synthesizes an inter-frame with bidirectional ConvLSTM. Different from the memory module in (Gong et al. 2019), Park et al. (2020) introduce a novel memory module to capture normal prototype features and incorporate a wider range of patterns. There are also two-stream approaches that separately learn spatio-temporal normality patterns. Chang et al. (2022) design a two-stream model consisting of a spatial autocoder and a temporal autocoder based on deep K-mean clustering. Hao et al. (2022) leverage a 3D CNN-based encoder and a 2D CNN-based decoder to improve the consistency of generated results in the spatio-temporal domain. Cai et al. (2021) utilize the prior knowl-

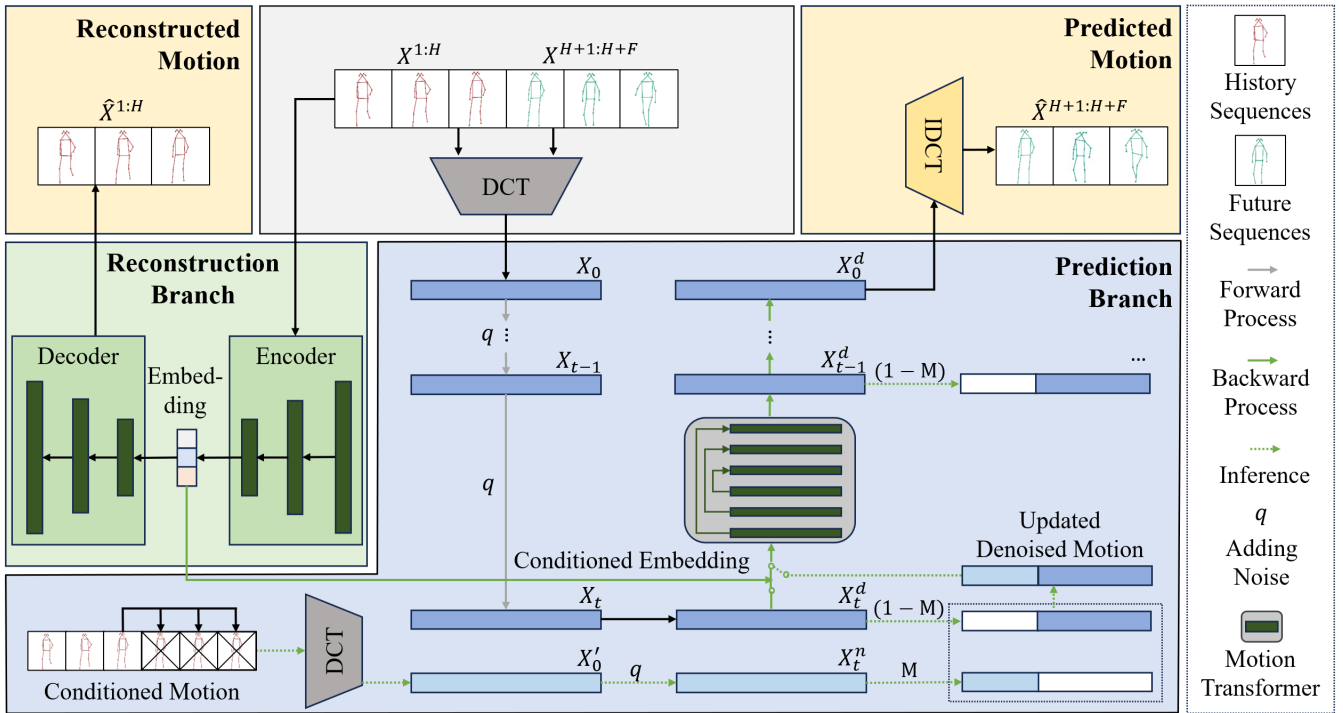


Figure 2: Overall architecture of the proposed method. The  $H + F$  skeletal motion sequences are split into history motion sequences (red skeletal motion sequences) and future motion sequences (green skeletal motion sequences). The key point of our model is the dual conditioned motion diffusion, i.e., the hidden representation of the observed sequence obtained by the encoder and the complete sequence of observed sequences with added noise connected to the predicted future motion sequence.

edge of appearance and motion signals to explicitly capture their correspondence in the high-level feature space. A bidirectional spatio-temporal feature learning framework is also proposed (2022a). Although the aforementioned methods make numerous attempts to acquire spatio-temporal representations and achieve improved results, prediction-based methods are still limited in certain aspects. A potential solution to this challenge lies in combining the strengths of both prediction-based and reconstruction-based methods.

**Pose-Based Methods:** Instead of using RGB video, pose-based methods utilize low-dimensional semantically skeleton data which simulates the dynamics of human joints over time. Morais et al. (2019) capture the overall dynamics of the body in motion and the spatial relationships of the skeleton. Markovitz et al. (2019) use embedded pose graphs and a Dirichlet process mixture for pose-based anomaly detection and introduce a coarse-grained setting aiming to detect abnormal variations of an action. Luo et al. (2020) propose a novel technique using a spatio-temporal Graph Convolutional Network (GCN). Yu et al. (2023) introduce a motion prior regularity learner to enhance dynamic representation. Jain et al. (2021) recently introduce an innovative strategy of the Conditional Variational Auto-Encoder (CVAE) framework, which employs a hybrid training strategy that combines self-supervised and unsupervised learning. Flaborea et al. (2024) utilize a graph convolutional network to represent human skeletal motion and learn to encode skeletal kinematics onto a minimum volume of the potential hypersphere.

Unlike the aforementioned approaches, we investigate the roles of conditions in the diffusion model for VAD, along with the examination of training and inference strategies.

## Method

Consider  $X = [x^1, \dots, x^{H+F}] \in \mathbb{R}^{(H+F) \times J \times C}$  as a series of  $H + F$  continuous motion sequences belonging to a participant. Here,  $x^t \in \mathbb{R}^{J \times C}$  refers to the joint coordinates in frame  $t$ ,  $J$  represents the number of joints, and  $C$  denotes the pose's dimension. We divide  $X$  into two parts: history motion sequence  $X^{1:H}$  and future motion sequence  $X^{H+1:H+F}$ . The architecture of the proposed model is illustrated in Figure 2. The model comprises two branches, namely the prediction branch and the reconstruction branch. The former involves a diffusion process and a reverse process to make predictions, whereas the latter employs an encoder-decoder structure to reconstruct the history motion.

### Dual Conditioned Motion Diffusion

We construct a new Dual Conditioned Motion Diffusion (DCMD) for human motion synthesis. The key idea is to consider conditioned embedding and conditioned motion as the dual conditions in the reverse process to gradually predict the future sequence from noisy variable distribution. We assume that the network of noise prediction is parameterized with  $\theta$ . Let  $E$  denote the reconstruction branch encoder. In the reverse process of the proposed, the recovered human

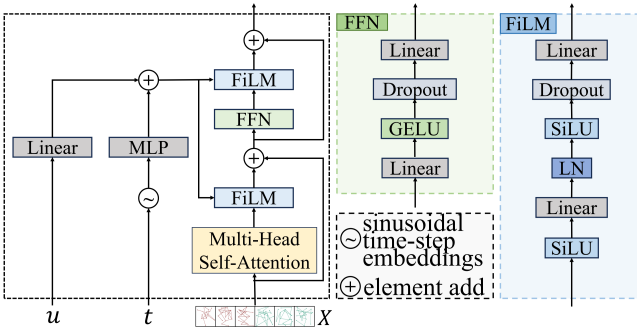


Figure 3: The architecture of the denoising network Motion Transformer. The green box on the right describes the details of the FFN module, and the blue box describes the details of the FiLM module.

motion is

$$X_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( X_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(X_t, t, u) \right) + \sigma_t z, \quad (1)$$

where  $u$  denotes motion encoding, which contains content information of the history motion sequence, i.e.,  $u = E(X^{1:H})$ , we name  $u$  as **conditioned embedding**.

Instead of recovering human motions from raw random noise, we add noise to the history motion sequence, concatenate it with the predicted future motion sequence, and subsequently send the whole sequence into the reverse process. We name the noised history sequence as **conditioned motion**. Details are shown in the Inference Section.

The denoising network in the reverse process is denoted as **motion transformer**, which is characterized by allowing for acquiring potential correlations from multi-layered characteristics.  $L$  motion transformer blocks incorporate skip connection stacking. Each block contains two linear-based FiLM modules inspired by (Perez et al. 2017). The FiLM modules are modulated by the diffusion time-step  $t$  embedding and conditioned embedding to establish temporal relationships. Its overall structure is shown in Figure 3. Assuming the input motion sequence is  $X$ , the overall equation can be formalized as follows:

$$\begin{aligned} Z &= \text{FiLM}(\text{Attention}(X) + \text{TE}(t) + u) + X, \\ Y &= \text{FiLM}(\text{FFN}(Z) + \text{TE}(t) + u) + Z, \end{aligned} \quad (2)$$

where  $Y \in \mathbb{R}^{(H+F) \times D}$  denotes an output with channel  $D$ .  $Z \in \mathbb{R}^{(H+F) \times D}$  is the hidden representation.  $\text{TE}(\cdot)$  denotes sinusoidal time-step embedding, and  $\text{Attention}(\cdot)$  denotes multi-head self-attention which is calculated by

$$\text{Attention}(X) = \text{Softmax}\left(\frac{Q_m K_m^T}{\sqrt{2D}}\right) V_m, \quad (3)$$

where  $Q_m = XW_Q, K_m = XW_K, V_m = XW_V \in \mathbb{R}^{(H+F) \times \frac{D}{h}}$  is the query, key, and value of the  $m$ -th head self-attention, respectively.  $\text{Softmax}(\cdot)$  normalizes the attention graph along the last dimension.

Finally, motion transformer concatenates the outputs of multi-heads  $\left\{ Y_m \in \mathbb{R}^{(H+F) \times \frac{D}{h}} \right\}_{1 \leq m \leq h}$  to obtain the final result  $Y$ .

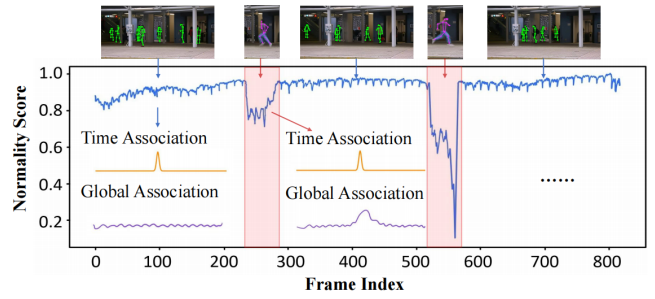


Figure 4: Graphical illustration of time association and global association. The blue curve is the normality score for a test segment of the CHUK dataset, and the red area indicates the time period in which the abnormal event occurred. Yellow curves indicate time association, and purple curves indicate global association.

Algorithm 1: Training procedure of the proposed framework.

**Require:** noising steps  $T$ , maximum iterations  $I_{max}$ .

**Input:** motion sequence  $X \in \mathbb{R}^{(H+F) \times 2J}$ .

**Output:** the noise prediction network  $\epsilon_\theta$ .

- 1: **for**  $I \leftarrow 0$  to  $I_{max}$  **do**
- 2: Divide  $X$  into history motion sequence  $X^{1:H}$  and future motion sequence  $X^{H+1:H+F}$ .
- 3: Encode the history motion sequence as  $u$ .
- 4: Decode  $u$  as reconstruction motion sequence  $\hat{X}^{1:H}$ .
- 5: Calculate the reconstruction loss  $\mathcal{L}_{rec}$ .
- 6: Perform DCT transformation on  $X$  yields  $X_0$ .
- 7: Sample the time steps  $t = \text{Uniform}(\{1, 2, \dots, T\})$ .
- 8: Add  $t$ -step noise to  $X_0$  using pre-defined variance parameters  $\alpha_t$  and noise  $\epsilon \in \mathcal{N}(0, I)$  yields  $X_t$ .
- 9: Calculate the prediction loss  $\mathcal{L}_{pred}$  in Eq. (6).
- 10: Calculate the reconstruction loss  $\mathcal{L}_{rec}$  in Eq. (7).
- 11: Calculate the additional loss UAD in Eq. (10).
- 12: Calculate the total loss  $\mathcal{L}_{total}$  in Eq. (11).
- 13: Update parameters of the network  $\theta = \theta - \nabla_{\theta} \mathcal{L}_{total}$ .
- 14: **end for**

### Training with UAD

Inspired by (Mao et al. 2019), during the training stage, we perform a Discrete Cosine Transform (DCT) operation on the complete motion sequence  $X$  to acquire the spectrum,

$$X_0 = \text{DCT}(X), \quad (4)$$

where  $\text{DCT}(\cdot)$  denotes the DCT operation, and  $X_0$  is the DCT coefficients.

In the forward process  $q$ , we can compute the noisy DCT spectrum  $X_t$  at a given time step  $t$  by employing the reparameterization trick.

$$X_t = \sqrt{\bar{\alpha}_t} X_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad (5)$$

where  $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ ,  $\alpha_i \in [0, 1]$  are pre-defined variance parameters, and  $\epsilon \sim \mathcal{N}(0, I)$ .

In the reverse process, we optimize the denoising model parameter  $\theta$  using the noise prediction loss function,

$$\mathcal{L}_{pred} = \text{smooth}_{L_1}(\mathbb{E}_{\epsilon, t} [\|\epsilon - \epsilon_\theta(X_t, t, u)\|^2]), \quad (6)$$

where  $\text{smooth}_{L_1}(\cdot)$  denotes the Smooth L1 Loss.

Considering both the temporal and spatial dimensions of the input motion sequence, we use the Space-Time-Separable Autoencoder (STSAE) (Sofianos et al. 2021; Flaborea et al. 2023), which relies on a U-Net-like architecture (Wyatt et al. 2022) to shrink the skeletal motion network step-wise and expand the spatial dimensions of the input motion sequence. The encoder E encodes the history motion sequence as  $u = E(X^{1:H})$ , while the decoder D reconstructs them into  $\hat{X}^{1:H}$  by using the reconstruction loss,

$$\mathcal{L}_{rec} = \left\| \hat{X}^{1:H} - X^{1:H} \right\|_2^2. \quad (7)$$

In order to enhance the distinction between normal and abnormal frames, we leverage the **United Association Discrepancy (UAD)** as an additional loss function. The primary components of the UAD loss consist of **time association** and **global association**, both aimed at serving the subsequent United Association Discrepancy. The two components are illustrated in Figure 4.

Time association is only considered from the dimension of time distance. Frames close to each other should have a stronger association, while those farther apart should have a weaker one. This pattern follows the normal distribution curve, resulting in a single-peaked shape. To establish time association, we employ a learnable Gaussian kernel (Xu et al. 2022) to determine relative time distance beforehand. Besides, we incorporate a learnable scale parameter to the Gaussian kernel, which enables us to concentrate on the neighboring regions for time association and accommodate various time-series patterns, including anomalies of varying durations. The formulation is as follows:

$$\mathcal{T} = \text{Rescale} \left( \left[ \frac{1}{\sqrt{2\pi}\sigma_i} \exp \left( -\frac{|j-i|^2}{2\sigma_i^2} \right) \right] \right), \quad (8)$$

where  $\sigma \in \mathbb{R}^{(H+F) \times h}$  denotes the learnable scale parameter for  $h$  heads and  $i, j \in \{1, \dots, (H+F)\}$ . To transform the association weights into a discrete distribution, we utilize  $\text{Rescale}(\cdot)$  by dividing the row sum.

Global association is considered from the numerical dimension. Frames with abnormal behavior should have a stronger association with the frames close to each other in the time dimension and a weaker association with the frames farther away, showing a single-peaked pattern. Similarly, frames that behave normally have a strong association in the time dimension with frames that are either far or close in distance. Global association learns associations from the original motion sequence and subsequently adaptively identifies the most efficient associations. Global association is part of self-attention, which is calculated as:

$$\mathcal{G} = \text{Softmax} \left( \frac{Q_m K_m^T}{\sqrt{2D}} \right). \quad (9)$$

---

Algorithm 2: Inference procedure of the proposed framework.

---

**Require:** noising steps  $T$ , the mask of the observation  $M$ , the trained noise prediction network  $\epsilon_\theta$ .

**Input:** observed motion sequence  $X^{1:H} \in \mathbb{R}^{H \times 2J}$ .

**Output:** completed motion sequence  $X \in \mathbb{R}^{(H+F) \times 2J}$ .

- 1: Encode the  $X^{1:H}$  sequence as  $u$ .
  - 2: Sample random noise  $X_T^d \in \mathcal{N}(0, I)$ .
  - 3: Pad the observed motion sequence  $X' \in \mathbb{R}^{(H+F) \times 2J}$ .
  - 4: Perform DCT transformation on  $X'$  yields  $X'_0$ .
  - 5: **for**  $t \leftarrow T - 1$  to 0 **do**
  - 6:   Select Gaussian noise  $z \in \mathcal{N}(0, I)$  if  $t > 0$ , else set  $z = 0$ .
  - 7:   Add  $t$ -step noise to  $X'_0$  yields noised history sequence  $X_t^n$ .
  - 8:   Select denoised prediction sequence  $X_t^d$ .
  - 9:   Perform mask completion for predicted motions  $\text{DCT} [M \odot \text{iDCT}(X_t^n) + (1 - M) \odot \text{iDCT}(X_t^d)]$ .
  - 10: **end for**
  - 11: Perform iDCT transformation on  $X_0$ .
  - 12: Select the following  $F$  frames as predicted motion sequence  $\hat{X}^{H+1:H+F}$ .
  - 13: Decode  $u$  as reconstructed motion sequence  $\hat{X}^{1:H}$ .
  - 14: Concatenate the reconstructed motion and the predicted motion to obtain the completed motion  $X$ .
- 

Abnormal frames exhibit a more significant similarity between time association and global association, resulting in a smaller KL value. Conversely, normal frames display lower similarity between these associations, yielding a larger KL value. As a result, we define the United Association Discrepancy as a symmetric KL divergence between time association and global association (Xu et al. 2022). UAD is computed by averaging over multiple layers, thereby consolidating the associations of a range of feature layers into a more informative and robust metric,

$$\text{UAD}(\mathcal{T}, \mathcal{G}; X) = \left[ \frac{1}{L} \sum_{l=1}^N (\text{KL}(\mathcal{T}_{t,:} || \mathcal{G}_{t,:}) + \text{KL}(\mathcal{G}_{t,:} || \mathcal{T}_{t,:})) \right]_{t=1, \dots, H+F}, \quad (10)$$

where the KL divergence is computed between the two discrete distributions that correspond to each row of  $\mathcal{T}$  and  $\mathcal{G}$ . We can see that abnormal frames have a smaller UAD than normal frames, making it an inherently distinguishable factor. Therefore, the total loss becomes:

$$\mathcal{L}_{total} = \mathcal{L}_{rec} + \mathcal{L}_{pred} - \lambda \times \|\text{UAD}(\mathcal{T}, \mathcal{G}; X)\|_1, \quad (11)$$

where  $\lambda$  is the contribution of additional loss.

The primary objective of our study is to minimize the total loss. In cases where the parameter  $\lambda > 0$ , we aim to optimize the additional loss by enhancing the UAD. However, directly maximizing the UAD is problematic because the abnormal frames are few and can be easily overlooked. Meanwhile, the scale parameter of the Gaussian kernel is significantly reduced, rendering the time association irrelevant. To better

		HR-STC	HR-Avenue	HR-UBnormal	UBnormal
MPED-RNN (Morais et al. 2019)	<i>CVPR'19</i>	75.4	86.3	61.2	60.0
GEPC (Markovitz et al. 2019)	<i>CVPR'20</i>	74.8	58.1	55.2	53.4
PoseCVAE (Jain et al. 2021)	<i>ICPR'21</i>	75.7	87.8	–	–
MoCoDAD (Flaborea et al. 2023)	<i>ICCV'23</i>	77.6	89.0	68.4	68.3
COSKAD (Flaborea et al. 2024)	<i>PR'24</i>	77.1	87.8	65.5	65.0
TrajREC (Stergiou and De Weerd 2024)	<i>WACV'24</i>	77.9	89.4	68.2	68.0
Reconstruction Only		77.1	87.4	66.4	66.2
Prediction Only		77.5	88.7	67.3	67.2
<b>DCMD (Ours)</b>		<b>78.6</b>	<b>90.0</b>	<b>69.0</b>	<b>69.0</b>

Table 1: Comparison of our proposed method with the state-of-the-art methods based on pose data for the AUC score (%).

control united learning, we utilize the minimax strategy for UAD, which employs a specially designed stopping gradient mechanism to constrain time association and global association for more distinguishable united association discrepancy. Specifically, during the UAD minimization phase, we aim to make  $\mathcal{T}$  converge towards  $\mathcal{G}$ , which is learned from the original motion sequence.  $\mathcal{G}$  stops gradient backpropagation, and this convergence process helps  $\mathcal{T}$  adjust different temporal patterns. In the UAD maximization phase, we focus on optimizing  $\mathcal{G}$  by stopping the gradient backpropagation of  $\mathcal{T}$  in order to increase the UAD, thereby compelling  $\mathcal{G}$  to place more emphasis on non-adjacent frames. The training process of the proposed method is described in Algorithm 1.

### Inference with Mask Completion

During the inference stage, the poses of all participants in all frames of the input motion sequences window  $\mathcal{W}$  are first extracted to obtain the set  $\mathcal{A}$ . The reconstruction branch generates  $m$  history motion sequence by utilizing the encoder in conjunction with the symmetric decoder. Meanwhile, the prediction branch generates  $m$  future motion sequence via mask completion. Specifically, the last frame of the observation sequence is first repeated until it matches the length of the complete sequence. The padded sequence is then transformed by DCT to get compact history information, denoted as  $X'_0$ . The noise is added to  $X'_0$  to derive the noise spectrum of the observation at time step  $t$ , which is denoted as  $X_t^n$ .

Unlike previous methods, we aggregate the noised history sequence and denoised prediction sequence because it can be observed that the noise observation spectrum,  $X_t^n$ , and the denoised prediction spectrum,  $X_t^d$ , are approximately distributed similarly. The iDCT operation is employed to transform the noised and denoised spectrum into the time domain. Then, a mask mechanism is employed to combine the two spectrums,

$$X_t = \text{DCT} \left[ \text{M} \odot \text{iDCT} (X_t^n) + (1 - \text{M}) \odot \text{iDCT} (X_t^d) \right], \quad (12)$$

where  $\text{M}$  is the mask with the first  $H$  values being 1 and the other values being 0, and  $\odot$  denotes the Hadamard product.

Subsequently, this updated denoised motion and the conditioned embedding are treated as dual conditions and fed into the denoising network. Therefore, we can complete the

predicted motion for each reverse diffusion step. The inference process is described in Algorithm 2.

## Experiments

### Implementation Details

We conduct experiments on four popular benchmarks: Human-related ShanghaiTech Campus (**HR-STC**), Human-related CUHK Avenue (**HR-Avenue**), **HR-UBnormal**, and **UBnormal**. To avoid influences caused by incorrect skeleton detection in video frames, we followed the setting (Morais et al. 2019) that removes segments where skeletons cannot be detected using pose estimation algorithms. The Area Under Curve (AUC) of the Receiver Operating Characteristic (ROC) curve is used as the evaluation metric.

Human motion is represented using a 17-joint skeleton. For extracting motion sequences, we employ a window size of 7 frames, where the first 3 frames comprise the historical motion sequences, and the subsequent 4 frames represent the future motion sequences. We train the network end-to-end using the Adam optimizer with a learning rate of  $1e - 4$  that is decayed every 36 epochs. The diffusion process employs cosine variance scheduling with  $\beta_1 = 1e - 4$ ,  $\beta_T = 2e - 2$ , and  $T = 10$ . We set  $\lambda = 0.01$ . The hidden sizes of the encoder for the reconstruction branch are (512, 256), and the dimension of the hidden embedding is 256. The noise prediction network consisted of 6 layers of motion transformer blocks, where the number of heads is 8, and the hidden dimension is 512. The experiments are conducted on an NVIDIA GeForce RTX 4090 GPU. The batch size is set to 4096 for HR-STC and 1024 for HR-Avenue. The training process lasts approximately 6 hours.

### Comparison with the State-of-the-Arts

Our approach is compared with recent state-of-the-art methods such as PoseCVAE (Jain et al. 2021), COSKAD (Flaborea et al. 2024), and MoCoDAD (Flaborea et al. 2023). Specifically, PoseCVAE (Jain et al. 2021) and COSKAD (Flaborea et al. 2024) are reconstruction-based approaches. Specifically, MoCoDAD (Flaborea et al. 2023) applies the diffusion model to generate future motions based on past motions.

Table 1 shows the comparison of our approach with state-of-the-art methods on the HR-STC, the HR-Avenue, the

	UAD	CE	MC	HR-Avenue (AUC)
Ours	✓	✓	✓	<b>90.0</b>
①		✓	✓	89.1
②	✓		✓	88.7
③	✓	✓		89.1
④			✓	87.2
⑤		✓		87.8
⑥	✓			87.0
⑦				86.7
⑧	w/o DCT			89.2

Table 2: Ablation Studies on the HR-Avenue dataset (%).

HR-UBnormal, and the UBnormal. Our approach consistently beats the other state-of-the-art methods on the four benchmarks. Our approach outperforms the recent diffusion-based method MoCoDAD (Flaborea et al. 2023) by 1.0%, 1.0%, 0.6%, and 0.7% in AUC scores on the HR-STC, HR-Avenue, HR-UBnormal, and UBnormal datasets, respectively. Furthermore, our method significantly surpasses the reconstruction-based methods, specifically outperforming COSKAD (Flaborea et al. 2024) by 3.5% and 4.0% in AUC scores on the HR-UBnormal and UBnormal datasets, respectively. The results clearly demonstrate the advantages of our approach for video anomaly detection.

As our approach performs video anomaly detection by combining reconstructed and predicted errors of both history and future frames, respectively, we also present the results of anomaly detection solely based on either the reconstruction or the prediction branch. Our results surpass those of any single branch, demonstrating the complementarity of reconstruction and prediction for video anomaly detection.

### Ablation Studies

In order to more comprehensively demonstrate the efficacy of our proposed framework, we undertake ablation studies, examining factors such as DCT, United Association Discrepancy (UAD), Conditioned Embedding (CE), and Mask Completion (MC). We report the AUC scores for our model and its variants on HR-Avenue (Morais et al. 2019) in Table 2. Specifically, ⑦ indicates the baseline of the diffusion-based model that generates motion only from pure noise.

After removing the DCT, the result on the HR-Avenue dataset decreases 0.8%, it is interpreted that incorporating DCT enhances the accuracy of the motion predictions, thus improving the prediction branch for VAD. By removing the UAD module, we find that the AUC score decreases by 0.9%, which shows that UAD enhances the distinction between normal and abnormal frames. When we remove the CE module, the whole framework amounts to prediction only, and we notice a 1.3% decrease in the AUC score, demonstrating the necessity of combining reconstruction and prediction methods. Removing the MC module results in a decrease in the AUC score of 0.9%, suggesting that the mask completion operation can effectively utilize the observed frames to generate more controlled predictions. In

Parameter	Value	HR-STC	HR-Avenue
History	2	77.0	86.7
	<b>3</b>	<b>78.6</b>	<b>89.9</b>
	4	77.8	88.0
	5	77.7	87.1
Future	3	77.9	89.9
	<b>4</b>	<b>78.1</b>	<b>90.0</b>
$\lambda$	5	77.8	88.5
	0	77.6	89.1
	<b>0.01</b>	<b>77.9</b>	<b>90.0</b>
	0.02	77.8	88.6
	0.05	75.6	87.0

Table 3: Parameter analysis of the proposed method.

addition, we further demonstrate the necessity of each module by removing both UAD and CE, UAD and MC, and CE and MC, the AUC scores decrease by 2.8%, 2.2%, and 3.0%. All results show the necessity of each module, and by combining each module, our model obtained the highest AUC score.

### Parameter Analysis

We explore the impacts of the length of history and future motion sequence, and the value of the additional loss parameter on the HR-STC and HR-Avenue datasets, respectively. Our best method is highlighted in bold in Table 3. Since we employ the sliding window technique, the length of each of our windows should not be too large to enhance the accuracy of our computations. Specifically, for both datasets, we select a history motion sequence length of 3 and a future motion sequence length of 4. Besides, we present a reasonable selection of different values of additional loss parameters, and the most reasonable value is 0.01. When the value increases, the result decreases significantly.

### Conclusion

In this work, we propose a Dual Conditioned Motion Diffusion (DCMD) for Video Anomaly Detection (VAD), which combines the advantages of both reconstruction-based and prediction-based methods. The DCMD incorporates conditioned motion and conditioned embedding in the diffusion-based network of future motion prediction. During training, a United Association Discrepancy (UAD) regularization is introduced, and during inference, a mask completion method is employed. Our approach consistently achieves state-of-the-art performance on four VAD datasets. Detailed ablated experiments demonstrate the effectiveness of different components of the DCMD. Further experimental analysis demonstrates that the reconstruction and prediction branches are very complementary in VAD. Parameter analysis also suggests that the proposed framework is not sensitive to parameters. We hope this research can bring a new perspective on VAD by combining the advantages of both reconstruction-based and prediction-based methods.

## Acknowledgments

This work was supported by National Science Foundation of China (62172090, 62302093, 52441503), Jiangsu Province Natural Science Fund (BK20230833), Start-up Research Fund of Southeast University (RF1028623097), and Big Data Computing Center of Southeast University.

## References

- Cai, R.; Zhang, H.; Liu, W.; Gao, S.; and Hao, Z. 2021. Appearance-Motion Memory Consistency Network for Video Anomaly Detection. In *AAAI Conference on Artificial Intelligence*.
- Chang, Y.; Tu, Z.; Xie, W.; Luo, B.; Zhang, S.; Sui, H.; and Yuan, J. 2022. Video anomaly detection with spatio-temporal dissociation. *Pattern Recognit.*, 122: 108213.
- Chang, Y.; Tu, Z.; Xie, W.; and Yuan, J. 2020. Clustering Driven Deep Autoencoder for Video Anomaly Detection. In *European Conference on Computer Vision*.
- Chen, Z.; Wang, H.; and Gui, J. 2023. Occluded Skeleton-Based Human Action Recognition with Dual Inhibition Training. In *Proceedings of the ACM International Conference on Multimedia*, 2625–2634.
- Doshi, K.; and Yilmaz, Y. 2021. Online anomaly detection in surveillance videos with asymptotic bound on false alarm rate. *Pattern Recognition*, 114: 107865.
- Doshi, K.; and Yilmaz, Y. 2022. Rethinking Video Anomaly Detection - A Continual Learning Approach. *IEEE/CVF Winter Conference on Applications of Computer Vision*, 3036–3045.
- Flaborea, A.; Collorone, L.; D’Amely, G.; D’Arrigo, S.; Prenkaj, B.; and Galasso, F. 2023. Multimodal Motion Conditioned Diffusion Model for Skeleton-based Video Anomaly Detection. *IEEE/CVF International Conference on Computer Vision*, 10284–10295.
- Flaborea, A.; di Melendugno, G. M. D.; D’arrigo, S.; Sterpa, M. A.; Sampieri, A.; and Galasso, F. 2024. Contracting skeletal kinematics for human-related video anomaly detection. *Pattern Recognition*, 110817.
- Gong, D.; Liu, L.; Le, V.; Saha, B.; Mansour, M. R.; Venkatesh, S.; and van den Hengel, A. 2019. Memorizing Normality to Detect Anomaly: Memory-Augmented Deep Autoencoder for Unsupervised Anomaly Detection. *IEEE/CVF International Conference on Computer Vision*, 1705–1714.
- Hao, Y.; Li, J.; Wang, N.; Wang, X.; and Gao, X. 2022. Spatiotemporal consistency-enhanced network for video anomaly detection. *Pattern Recognit.*, 121: 108232.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.
- Ho, J.; and Salimans, T. 2021. Classifier-Free Diffusion Guidance. In *NeurIPS Workshop on Deep Generative Models and Downstream Applications*.
- Jain, Y.; Sharma, A. K.; Velmurugan, R.; and Banerjee, B. 2021. PoseCVAE: Anomalous Human Activity Detection. *International Conference on Pattern Recognition*, 2927–2934.
- Lee, S.; Kim, H. G.; and Ro, Y. M. 2018. Stan: Spatio-Temporal Adversarial Networks for Abnormal Event Detection. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 1323–1327.
- Li, N.; Chang, F.; and Liu, C. 2021. Spatial-Temporal Cascade Autoencoder for Video Anomaly Detection in Crowded Scenes. *IEEE Transactions on Multimedia*, 23: 203–215.
- Liu, W.; Luo, W.; Lian, D.; and Gao, S. 2017. Future Frame Prediction for Anomaly Detection - A New Baseline. *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6536–6545.
- Luo, W.; Liu, W.; and Gao, S. 2020. Normal graph: Spatial temporal graph convolutional networks based prediction network for skeleton based video anomaly detection. *Neurocomputing*, 444: 332–337.
- Luo, W.; Liu, W.; Lian, D.; Tang, J.; Duan, L.; Peng, X.; and Gao, S. 2019. Video Anomaly Detection with Sparse Coding Inspired Deep Neural Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43: 1070–1084.
- Mao, W.; Liu, M.; Salzmann, M.; and Li, H. 2019. Learning Trajectory Dependencies for Human Motion Prediction. *IEEE/CVF International Conference on Computer Vision*, 9488–9496.
- Markovitz, A.; Sharir, G.; Friedman, I.; Zelnik-Manor, L.; and Avidan, S. 2019. Graph Embedded Pose Clustering for Anomaly Detection. *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10536–10544.
- Morais, R.; Le, V.; Tran, T.; Saha, B.; Mansour, M. R.; and Venkatesh, S. 2019. Learning Regularity in Skeleton Trajectories for Anomaly Detection in Videos. *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11988–11996.
- Park, H.; Noh, J.; and Ham, B. 2020. Learning Memory-Guided Normality for Anomaly Detection. *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14360–14369.
- Perez, E.; Strub, F.; de Vries, H.; Dumoulin, V.; and Courville, A. C. 2017. FiLM: Visual Reasoning with a General Conditioning Layer. In *AAAI Conference on Artificial Intelligence*.
- Sofianos, T.; Sampieri, A.; Franco, L.; and Galasso, F. 2021. Space-Time-Separable Graph Convolutional Network for Pose Forecasting. *IEEE/CVF International Conference on Computer Vision*, 11189–11198.
- Stergiou, A.; and De Weerd, B. e. a. 2024. Holistic representation learning for multitask trajectory anomaly detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 6729–6739.
- Tang, Y.; Zhao, L.; Zhang, S.; Gong, C.; Li, G.; and Yang, J. 2020. Integrating prediction and reconstruction for anomaly detection. *Pattern Recognit. Lett.*, 129: 123–130.
- Wang, H.; and Wang, L. 2018. Beyond joints: Learning representations from primitive geometries for skeleton-based action recognition and detection. *IEEE Transactions on Image Processing*, 27(9): 4382–4394.

Wang, H.; Zhao, J.; and Gui, J. 2024. Region-aware image-based human action retrieval with transformers. *Computer Vision and Image Understanding*, 249: 104202.

Weng, W.; Wang, H.; He, J.; He, L.; and Xie, G. 2025. US-DRL: Unified Skeleton-Based Dense Representation Learning with Multi-Grained Feature Decorrelation. In *AAAI Conference on Artificial Intelligence*.

Wyatt, J.; Leach, A.; Schmon, S. M.; and Willcocks, C. G. 2022. AnoDDPM: Anomaly Detection with Denoising Diffusion Probabilistic Models using Simplex Noise. *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 649–655.

Xu, J.; Wu, H.; Wang, J.; and Long, M. 2022. Anomaly Transformer: Time Series Anomaly Detection with Association Discrepancy. In *International Conference on Learning Representations*.

Yu, S.; Zhao, Z.; Fang, H.; Deng, A.; Su, H.; Wang, D.; Gan, W.; Lu, C.; and Wu, W. 2023. Regularity learning via explicit distribution modeling for skeletal video anomaly detection. *IEEE Transactions on Circuits and Systems for Video Technology*.

Zaheer, M.; ha Lee, J.; Astrid, M.; and Lee, S.-I. 2020. Old Is Gold: Redefining the Adversarially Learned One-Class Classifier Training Paradigm. *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14171–14181.

Zhong, Y.; Chen, X.; Hu, Y.; Tang, P.; and Ren, F. 2022a. Bidirectional spatio-temporal feature learning with multi-scale evaluation for video anomaly detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(12): 8285–8296.

Zhong, Y.; Chen, X.; Jiang, J.; and Ren, F. 2022b. A cascade reconstruction model with generalization ability evaluation for anomaly detection in videos. *Pattern Recognition*, 122: 108336.

Zhong, Y.; Chen, X.; Jiang, J.; and Ren, F. 2022c. Reverse erasure guided spatio-temporal autoencoder with compact feature representation for video anomaly detection. *Sci. China Inf. Sci.*, 65(9): 1–3.

Zhou, J. T.; Zhang, L.; Fang, Z.; Du, J.; Peng, X.; and Yang, X. 2020. Attention-Driven Loss for Anomaly Detection in Video Surveillance. *IEEE Transactions on Circuits and Systems for Video Technology*, 30: 4639–4647.