

Diffusion Action Segmentation

Daochang Liu¹, Qiyue Li^{2,4}, Anh-Dung Dinh¹, Tingting Jiang², Mubarak Shah³, Chang Xu¹

¹School of Computer Science, Faculty of Engineering, The University of Sydney

²NERCVT, NKLMIP, School of Computer Science, ⁴School of Mathematical Sciences, Peking University

³Center for Research in Computer Vision, University of Central Florida

{daochang.liu, c.xu}@sydney.edu.au shah@crcv.ucf.edu

Abstract

Temporal action segmentation is crucial for understanding long-form videos. Previous works on this task commonly adopt an iterative refinement paradigm by using multi-stage models. We propose a novel framework via denoising diffusion models, which nonetheless shares the same inherent spirit of such iterative refinement. In this framework, action predictions are iteratively generated from random noise with input video features as conditions. To enhance the modeling of three striking characteristics of human actions, including the position prior, the boundary ambiguity, and the relational dependency, we devise a unified masking strategy for the conditioning inputs in our framework. Extensive experiments on three benchmark datasets, i.e., GTEA, 50Salads, and Breakfast, are performed and the proposed method achieves superior or comparable results to state-of-the-art methods, showing the effectiveness of a generative approach for action segmentation. Code is at tinyurl.com/DiffAct.

1. Introduction

Temporal action segmentation is a key task for understanding and analyzing human activities in complex long videos, with a wide range of applications from video surveillance [62], video summarization [3] to skill assessment [43]. The goal of temporal action segmentation is to take as input an untrimmed video and output an action sequence indicating the class label for each frame.

This task has witnessed remarkable progresses in recent years with the development of multi-stage models [20, 41, 68, 72]. The core idea of multi-stage models is to stack several stages where the first stage produces an initial prediction and later stages adjust the prediction from the preceding stage. Researchers have actively explored various architectures to implement the multi-stage model, such as the MS-TCN [20, 41] relying on dilated temporal convolutional lay-

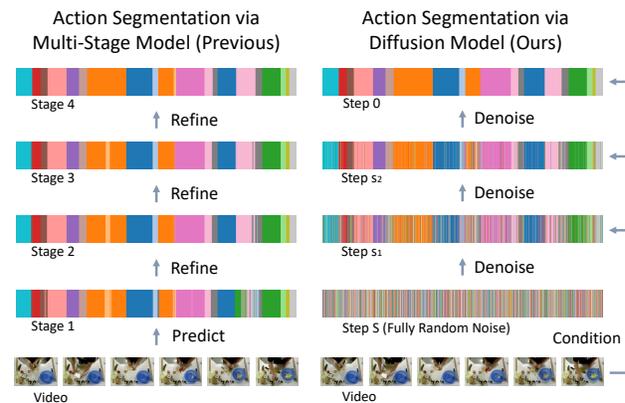


Figure 1. **Multi-stage model vs. diffusion model for action segmentation.** They both follow an iterative refinement paradigm. *Left:* Many previous methods utilize a multi-stage framework to refine the initial prediction. *Right:* We formulate action segmentation as a frame-wise action sequence generation problem and obtain the refined prediction by an iterative denoising process. Colors in the barcodes represent different actions.

ers and the ASFormer [72] with attention mechanisms. The success of multi-stage models could be largely attributed to the underlying iterative refinement paradigm that properly captures temporal dynamics of actions and significantly reduces over-segmentation errors [15].

In this paper, we propose an action segmentation method following the same philosophy of iterative refinement but in an essentially new generative approach, which incorporates the denoising diffusion model. Favored for its simple training recipe and high generation quality, the diffusion model [13, 54, 14, 26] has become a rapidly emerging category of generative models. A forward process in the diffusion model corrupts the data by gradually adding noise, while a corresponding reverse process removes the noise step by step so that new samples can be generated from the data distribution starting from fully random noise. Such iterative denoising in the reverse process coincides with the iterative refinement paradigm for action segmentation. This

motivates us to rethink action segmentation from a generative view employing the diffusion model as in Fig. 1, where we can formulate action segmentation as an action sequence generation problem conditioned on the input video. One distinct advantage of such diffusion-based action segmentation is that it not only learns the discriminative mapping from video frames to actions but also implicitly captures the prior distribution of human actions through generative modeling. This prior modeling can be further explicitly enhanced in line with three prominent characteristics of human actions. The first characteristic is the temporal *position prior*, which means certain actions are more likely to occur at particular time locations in the video. Taking a video of making salads as an example, actions of cutting vegetables tend to appear in the middle of the video, while serving salads onto the plate is mostly located at the end. The second characteristic is the *boundary prior*, which reflects that transitions between actions are visually gradual and thus lead to ambiguous features around action boundaries. The third characteristic is the *relation prior* that represents human actions usually adhere to some intrinsic temporal ordering, *e.g.*, cutting a cucumber typically follows behind peeling the cucumber. This relation prior differs from the position prior since it focuses on the arrangements relative to other actions. To jointly exploit these priors of human actions, we devise a condition masking strategy, which naturally fits into the newly proposed framework.

In our method, dubbed *DiffAct*, we formulate action segmentation as a conditional generation problem of the frame-wise action label sequence, leveraging the input video as the condition. During training, the model is provided with the input video features as well as a degraded temporal action segmentation sequence obtained from the ground truth with varying levels of injected noise. The model then learns to denoise the sequence to restore the original ground truth. To achieve this, we impose three loss functions between the denoised sequence and the ground-truth sequence, including a standard cross-entropy loss, a temporal smoothness loss, and a boundary alignment loss. At inference, following the reversed diffusion process, the model refines a random sequence in an iterative manner to generate the action prediction sequence. On the other hand, to boost the modeling of the three aforementioned priors of human actions, the conditional information in our framework is controlled through a condition masking strategy during training, which encourages the model to reason over information other than visual features, *e.g.*, time locations, action durations, and the temporal context. The masked conditions convert the generative learning of action sequences from a basic conditional one to a combination of fully conditional, partially conditional, and unconditional ones to enhance the three action priors simultaneously.

The effectiveness of our diffusion-based temporal action

segmentation is demonstrated by the experiments on three datasets, GTEA [21], 50Salads [59], and Breakfast [34], on which our model performs better or on par compared to state-of-the-art methods. In summary, our contributions are three-fold: 1) temporal action segmentation is formulated as a conditional generation task; 2) a new iterative refinement framework is proposed based on the denoising diffusion process; 3) a condition masking strategy is designed to further exploit the priors of human actions.

2. Related Works

Temporal action segmentation [15] is a complex video understanding task segmenting videos that can span for minutes. To model the long-range dependencies among actions, a rich variety of temporal models have been employed in the literature, evolving from early recurrent neural networks [16, 51], to temporal convolutional networks [37, 36, 38, 20, 46, 65, 25, 41, 64, 58, 23, 52, 42, 47, 75], graph neural networks [29, 74], and recent transformers [72, 4, 66, 18, 19, 61, 44]. Multi-stage models [72, 4, 41, 64, 20, 47, 68], which employ incremental refining, are especially notable given their superiority in capturing temporal context and mitigating over-segmentation errors [15]. Apart from architecture designs, another line of works focuses on acquiring more accurate and robust frame-level features through representation learning [1, 50, 40] or domain adaptation [10, 11].

For action segmentation, the position prior, the boundary ambiguity, and the relation prior are three beneficial inductive biases. The boundary ambiguity and relation prior have drawn attention from researchers while the position prior has been less explored. To cope with boundary ambiguity, in the boundary-aware cascade network [68], a local barrier pooling is presented to assign aggregation weights adaptive to boundaries. Ishikawa *et al.* integrated a complementary network branch to regress boundary locations [31], and Chen *et al.* estimated the uncertainty due to ambiguous boundaries by Monte-Carlo sampling [9]. Another strategy is to smooth the annotation into a soft version where action probabilities decrease around boundaries [33]. As for the relation prior, GTRM [29] and UVAST [6] respectively leveraged graph convolutional networks and sequence-to-sequence translation to facilitate relational reasoning at the segment level. For contextual relations between adjacent actions, Br-Prompt [40] resorted to multi-modal learning by using text prompts as supervision. Recently, DTL [69] mined relational constraints from annotations and enforced the constraints by a temporal logic loss during training. In this paper, we simply use condition masking to take care of the position prior, boundary ambiguity, and relation prior simultaneously, without additional designs for each of them.

One related work [22] also employs generative learning for action segmentation. This method synthesized interme-

diate action codes per frame to aid recognition using GANs. In contrast, our method generates action label sequences via the diffusion model.

Diffusion models [53, 13, 54, 26], which have been theoretically unified with the score-based models [55, 56, 57], are known due to their stable training procedure and do not require any adversarial mechanism for generative learning. The diffusion-based generation has achieved impressive results in image generation [14, 49, 7, 67], natural language generation [73], text-to-image synthesis [24, 32], audio generation [39, 35] and so on. A gradient view was proposed to further improve the diffusion sampling process with guidance [17]. Diffusion models were repurposed for some image understanding tasks in computer vision very recently, such as object detection [12] and image segmentation [5, 2]. However, only very few works targeted video-related tasks, including video forecasting and infilling [70, 28, 63]. A pre-trained diffusion model was fine-tuned to predict video memorability [60], and a frequency-aware diffusion module was proposed for video captioning [76]. In this paper, we recognize that diffusion models, given their iterative refinement properties, are especially suitable for temporal action segmentation. To the best of our knowledge, this work is the first one employing diffusion models for action analysis.

3. Preliminaries

We first familiarize the readers with the background of diffusion models. Diffusion models [26, 54] aim to approximate the data distribution $q(x_0)$ with a model distribution $p_\theta(x_0)$. The *forward process* or *diffusion process* corrupts the real data $x_0 \sim q(x_0)$ into a series of noisy data x_1, x_2, \dots, x_S . The *reverse process* or *denoising process* gradually removes the noise from $x_S \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to x_{S-1}, x_{S-2}, \dots until $x_0 \sim p_\theta(x_0)$ is achieved. S is the total number of steps.

Forward Process. Formally, the forward process adds Gaussian noise to the data at each step with a predefined variance schedule $\beta_1, \beta_2, \dots, \beta_S$: $q(x_s|x_{s-1}) = \mathcal{N}(x_s; \sqrt{1 - \beta_s}x_{s-1}, \beta_s \mathbf{I})$. By denoting $\alpha_s = 1 - \beta_s$ and $\bar{\alpha}_s = \prod_{i=1}^s \alpha_i$, we can directly obtain x_s from x_0 in a closed form without recursion: $q(x_s|x_0) = \mathcal{N}(x_s; \sqrt{\bar{\alpha}_s}x_0, (1 - \bar{\alpha}_s)\mathbf{I})$, which can be simplified using the reparameterization trick:

$$x_s = \sqrt{\bar{\alpha}_s}x_0 + \epsilon\sqrt{1 - \bar{\alpha}_s}. \quad (1)$$

The noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is sampled from a normal distribution at each step and its intensity is defined by $\sqrt{1 - \bar{\alpha}_s}$.

Reverse Process. The reverse process starts from x_S and progressively removes noise to recover x_0 , with one step in the process defined as $p_\theta(x_{s-1}|x_s)$:

$$p_\theta(x_{s-1}|x_s) = \mathcal{N}(x_{s-1}; \boldsymbol{\mu}_\theta(x_s, s), \sigma_s^2 \mathbf{I}), \quad (2)$$

where σ_s^2 is controlled by β_s , and $\boldsymbol{\mu}_\theta(x_s, s)$ is a predicted mean parameterized by a step-dependent neural network.

Several different ways [45] are possible to parameterize p_θ , including the prediction of mean as in Eq. 2, the prediction of the noise ϵ , and the prediction of x_0 . The x_0 prediction is used in our work. In such case, the model predicts x_0 by a neural network $f_\theta(x_s, s)$, instead of directly predicting the $\boldsymbol{\mu}_\theta(x_s, s)$ in Eq. 2. To optimize the model, a mean squared error loss can be utilized to match $f_\theta(x_s, s)$ and x_0 : $\mathcal{L} = \|f_\theta(x_s, s) - x_0\|^2, s \in_R \{1, 2, \dots, S\}$. The step s is randomly selected at each training iteration.

At the inference stage, starting from a pure noise $x_S \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the model can gradually reduce the noise according to the update rule [54] below using the trained f_θ :

$$x_{s-1} = \sqrt{\bar{\alpha}_{s-1}}f_\theta(x_s, s) + \frac{\sqrt{1 - \bar{\alpha}_{s-1} - \sigma_s^2}}{\sqrt{1 - \bar{\alpha}_s}}x_s - \frac{\sqrt{\bar{\alpha}_s}f_\theta(x_s, s)}{\sqrt{1 - \bar{\alpha}_s}} + \sigma_s \epsilon. \quad (3)$$

By iteratively applying Eq. 3, a new sample x_0 can be generated from p_θ via a trajectory x_S, x_{S-1}, \dots, x_0 . Some improved sampling strategy skips steps in the trajectory, *i.e.*, $x_S, x_{S-\Delta}, \dots, x_0$, for better efficiency [54].

Conditional information can be included in the diffusion model to control the generation process. The conditional model can be written as $f_\theta(x_s, s, \mathbb{C})$, with the conditional information \mathbb{C} as an extra input. In the literature, class labels [14], text prompts [24, 32], and image guidance [48] are the most common forms of conditional information.

4. Method

We formulate temporal action segmentation as a conditional generation problem of temporal action sequences (Fig. 2). Given the input features $F \in \mathbb{R}^{L \times D}$ with D dimensions for a video of L frames, we approximate the data distribution of the ground truth action sequence $Y_0 \in \{0, 1\}^{L \times C}$ in the one-hot form with C classes of actions.

Since the input features F are usually extracted per short clip by general-purpose pre-trained models, we employ an encoder h_ϕ to enrich the features with long-range temporal information and make them task-oriented, *i.e.*, $E = h_\phi(F)$. $E \in \mathbb{R}^{L \times D'}$ is the encoded feature with D' dimensions.

4.1. Diffusion Action Segmentation

The proposed method, *DiffAct*, constructs a diffusion process Y_0, Y_1, \dots, Y_S from the action ground truth Y_0 to a nearly pure noise Y_S at training, and a denoising process $\hat{Y}_S, \hat{Y}_{S-1}, \dots, \hat{Y}_0$ from a pure noise \hat{Y}_S to the action prediction \hat{Y}_0 at inference.

Training. To learn the underlying distribution of human actions, the model is trained to restore the action ground truth from its corrupted versions. Specifically, a random diffusion step $s \in_R \{1, 2, \dots, S\}$ is chosen at each training

iteration. Then we add noise to the ground truth sequence Y_0 according to the accumulative noise schedule as in Eq. 1 to attain a corrupted sequence $Y_s \in [0, 1]^{L \times C}$:

$$Y_s = \sqrt{\bar{\alpha}_s} Y_0 + \epsilon \sqrt{1 - \bar{\alpha}_s}, \quad (4)$$

where noise ϵ is from a Gaussian distribution.

Taking the corrupted sequence as input, a decoder g_ψ is designed to denoise the sequence:

$$P_s = g_\psi(Y_s, s, E \odot M), \quad (5)$$

where the resultant denoised sequence $P_s \in [0, 1]^{L \times C}$ indicates action probabilities for each frame. Apart from Y_s , the decoder also includes another two inputs, which are the step s and encoded video features E . Using step s as input makes the model step-aware so that we can share the same model to denoise at different noise intensities. The usage of encoded video features E as conditions ensures that the produced action sequences are not only broadly plausible but also consistent with the input video. More importantly, the conditional information is further diversified by element-wisely multiplying the feature E and a mask M to explicitly capture the three characteristics of human actions, which will be discussed later. The choices of the decoder g_ψ and encoder h_ϕ are flexible, which can be of common network architectures for action segmentation.

Loss Functions. With the denoised sequence P_s obtained, we impose three loss functions as below to match it with the original ground truth Y_0 .

Cross-Entropy Loss. The first loss is the standard cross-entropy for classification minimizing the negative log-likelihood of the ground truth action class for each frame:

$$\mathcal{L}_s^{\text{ce}} = \frac{1}{LC} \sum_{i=1}^L \sum_{c=1}^C -Y_{0,i,c} \log P_{s,i,c}, \quad (6)$$

where i is the frame index and c is the class index.

Temporal Smoothness Loss. To promote the local similarity along the temporal dimension, the second loss is computed as the mean squared error of the log-likelihoods between adjacent video frames [20, 41]:

$$\mathcal{L}_s^{\text{smo}} = \frac{1}{(L-1)C} \sum_{i=1}^{L-1} \sum_{c=1}^C (\log P_{s,i,c} - \log P_{s,i+1,c})^2. \quad (7)$$

Note that $\mathcal{L}_s^{\text{smo}}$ is clipped to avoid outlier values [20].

Boundary Alignment Loss. Accurate detection of boundaries is important for action segmentation. Therefore, the third loss is to align the action boundaries in the denoised sequence P_s and the ground truth sequence Y_0 . To this end, we need to derive the boundary probabilities from both P_s and Y_0 . First, a ground truth boundary sequence $B \in \{0, 1\}^{L-1}$ can be derived from the action ground truth

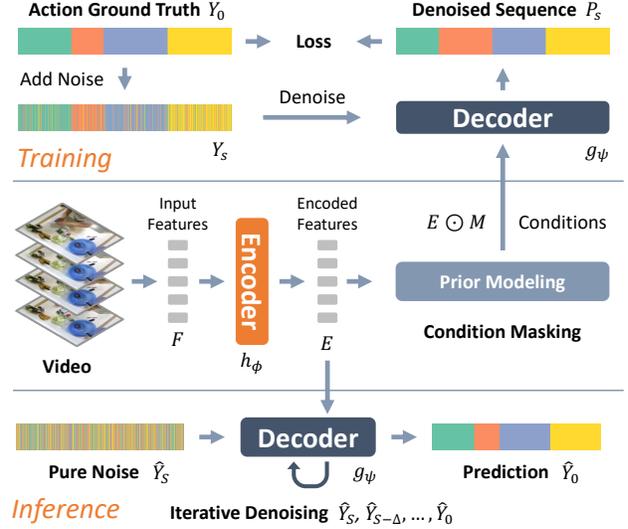


Figure 2. Method overview. During training, the proposed model is optimized to denoise corrupted action sequences. At inference, the model begins with a random noise sequence and obtains results via an iterative denoising process. The condition masking strategy strengthens the action prior modeling by blocking certain features during training. Input features are pre-extracted I3D features.

Y_0 , where $B_i = \mathbb{1}(Y_{0,i} \neq Y_{0,i+1})$. Since action transitions usually happen gradually, we smooth this sequence with a Gaussian filter for a soft version $\bar{B} = \lambda(B)$. As for the boundaries in the denoised sequence, their probabilities are computed with the dot product of the action probabilities from neighboring frames in P_s , i.e., $1 - P_{s,i} \cdot P_{s,i+1}$. The boundaries derived from the two sources are then aligned via a binary cross-entropy loss:

$$\mathcal{L}_s^{\text{bd}} = \frac{1}{L-1} \sum_{i=1}^{L-1} [-\bar{B}_i \log(1 - P_{s,i} \cdot P_{s,i+1}) - (1 - \bar{B}_i) \log(P_{s,i} \cdot P_{s,i+1})]. \quad (8)$$

The final training loss is a combination of the three losses at a randomly selected diffusion step: $\mathcal{L}^{\text{sum}} = \mathcal{L}_s^{\text{ce}} + \mathcal{L}_s^{\text{smo}} + \mathcal{L}_s^{\text{bd}}$, $s \in_R \{1, 2, \dots, S\}$.

Inference. Intuitively, the denoising decoder g_ψ is trained to adapt to sequences with arbitrary noise levels so that it can even handle a sequence made of purely random noise. Therefore, at inference, the denoising decoder starts from a pure noise sequence $\hat{Y}_S \in \mathcal{N}(\mathbf{0}, \mathbf{I})$ and gradually reduces the noise. Concretely, the update rule for each step is adapted from Eq. 3 as below:

$$\hat{Y}_{s-1} = \sqrt{\bar{\alpha}_{s-1}} P_s + \sqrt{1 - \bar{\alpha}_{s-1} - \sigma_s^2} \frac{\hat{Y}_s - \sqrt{\bar{\alpha}_s} P_s}{\sqrt{1 - \bar{\alpha}_s}} + \sigma_s \epsilon, \quad (9)$$

where the \hat{Y}_{s-1} is sent into the decoder to obtain the next



Figure 3. Illustration of the action prior modeling. The striped locations are masked. Different colors represent different actions.

P_{s-1} . This iterative refinement process yields a series of action sequences $\hat{Y}_S, \hat{Y}_{S-1}, \dots, \hat{Y}_0$, which leads to the \hat{Y}_0 at the end that can well approximate the underlying ground truth and is regarded as the final prediction. To speed up the inference, a sampling trajectory [54] with skipped steps $\hat{Y}_S, \hat{Y}_{S-\Delta}, \dots, \hat{Y}_0$ is utilized in our method. Note that the encoded features E are sent into the decoder without masking at inference. In addition, we fix the random seed for this denoising process to make it deterministic in practice.

4.2. Action Prior Modeling

Human behaviors follow regular patterns. We identify three valuable prior knowledge for action segmentation, which are the position prior, the boundary ambiguity, and the relation prior. One unique advantage of diffusion-based action segmentation is that it can capture the prior distribution of action sequences via its generative learning ability. This allows us to further exploit the three action priors by changing the conditional information for the model.

In detail, we devise a condition masking strategy to control the conditional information, which applies a mask $M \in \{0, 1\}^L$ to the features E in Eq. 5. At each training step, the mask is randomly sampled from a set, $M \in_R \{M^N, M^P, M^B, M^R\}$, with each element detailed below.

No Masking (N). The first type is a basic all-one mask $M^N = \mathbf{1}$ which lets all features pass into the decoder. This naive mask provides full conditional information for the model to map visual features to action classes.

Masking for Position Prior (P). The second type is an all-zero mask $M^P = \mathbf{0}$ that entirely blocks the features. Without any visual features, the only cues the decoder can rely on are the video length and the frame positions in the video¹. Therefore, the modeling of the positional prior of actions is promoted.

Masking for Boundary Prior (B). Due to the ambiguity of action transitions, the visual features around boundaries may not be reliable. For this reason, the third mask M^B removes the features close to boundaries based on the soft ground truth \bar{B} . This mask is defined as $M_i^B = \mathbb{1}(\bar{B}_i < 0.5)$, $i \in \{1, 2, \dots, L\}$. With mask M^B , the decoder is encouraged to further explore the context information about the action before and after the boundary and their durations, which can be more robust than the unreliable features solely.

¹Specifically, such positional information is available due to the temporal convolutions in the decoder.

Masking for Relation Prior (R). The ordinal relation is another fundamental characteristic of human actions. We thus mask the segments belonging to a random action class $\tilde{c} \in_R \{1, 2, \dots, C\}$ during training to enforce the model to infer the missing action based on its surrounding actions. The mask is denoted as M^R where $M_i^R = \mathbb{1}(Y_{0,i,\tilde{c}} \neq 1)$, $i \in \{1, 2, \dots, L\}$. Such masked segment modeling benefits the utilization of the relational dependencies among actions.

To summarize, by the condition masking strategy in Fig. 3, the three human action priors are simultaneously integrated into our diffusion-based framework. One interesting interpretation of this strategy is from the perspective of the classifier-free guidance of diffusion models [27]. Our model can be viewed to be fully conditional with the mask M^N , partially conditional with the mask M^B and M^R , and unconditional with the mask M^P . This brings a versatile model that captures the action distribution from different aspects. We also discuss more potential forms or the extending usage of this masking strategy in the supplementary.

5. Experiments

5.1. Setup

Datasets. Experiments are performed on three benchmark datasets. **GTEA** [21] contains 28 egocentric daily activity videos with 11 action classes. On average, each video is of one-minute duration and 19 action instances approximately. **50Salads** [59] includes 50 top-view videos regarding salad preparation with actions falling into 17 classes. The average video length is six minutes and the average number of instances per video is 20 roughly. **Breakfast** [34] comprises 1712 videos in the third-person view and 48 action classes related to making breakfasts. The videos are two-minute long on average but show a large variance in length. Seven action instances are contained in each video on average. Among the three datasets, Breakfast has the largest scale, and 50Salads consists of the longest videos and the most instances per video. Following previous works [40, 6, 69, 72, 41, 31, 20], five-fold cross-validation is performed on 50Salads and four-fold cross-validations are performed on GTEA and Breakfast. We use the same splits as in previous works.

Metrics. Following previous works, the frame-wise accuracy (Acc), the edit score (Edit), and the F1 scores at overlap thresholds 10%, 25%, 50% (F1@{10, 25, 50}) are reported. The accuracy assesses the results at the frame level, while the edit score and F1 scores measure the performance at the segment level. The same evaluation codes are used as in previous works [72, 20].

Implementation Details. For all the datasets, we leverage the I3D features [8] used in most prior works as the input features F with $D=2048$ dimensions. The encoder h_ϕ is a re-implemented ASFormer encoder [72]. The de-

Method	GTEA				50Salads				Breakfast			
	F1@{10, 25, 50}	Edit	Acc	Avg	F1@{10, 25, 50}	Edit	Acc	Avg	F1@{10, 25, 50}	Edit	Acc	Avg
[41]MS-TCN++, PAMI'20	88.8 / 85.7 / 76.0	83.5	80.1	82.8	80.7 / 78.5 / 70.1	74.3	83.7	77.5	64.1 / 58.6 / 45.9	65.6	67.6	60.4
[11]SSTDA, CVPR'20	90.0 / 89.1 / 78.0	86.2	79.8	84.6	83.0 / 81.5 / 73.8	75.8	83.2	79.5	75.0 / 69.1 / 55.2	73.7	70.2	68.6
[29]GTRM, CVPR'20	- / - / -	-	-	-	75.4 / 72.8 / 63.9	67.5	82.6	72.4	57.5 / 54.0 / 43.3	58.7	65.0	55.7
[68]BCN, ECCV'20	88.5 / 87.1 / 77.3	84.4	79.8	83.4	82.3 / 81.3 / 74.0	74.3	84.4	79.3	68.7 / 65.5 / 55.0	66.2	70.4	65.2
[10]MTDA, WACV'20	90.5 / 88.4 / 76.2	85.8	80.0	84.2	82.0 / 80.1 / 72.5	75.2	83.2	78.6	74.2 / 68.6 / 56.5	73.6	71.0	68.8
[23]G2L, CVPR'21	89.9 / 87.3 / 75.8	84.6	78.5	83.2	80.3 / 78.0 / 69.8	73.4	82.2	76.7	74.9 / 69.0 / 55.2	73.3	70.7	68.6
[1]HASR, ICCV'21	90.9 / 88.6 / 76.4	87.5	78.7	84.4	86.6 / 85.7 / 78.5	81.0	83.9	83.1	74.7 / 69.5 / 57.0	71.9	69.4	68.5
[31]ASRF, WACV'21	89.4 / 87.8 / 79.8	83.7	77.3	83.6	84.9 / 83.5 / 77.3	79.3	84.5	81.9	74.3 / 68.9 / 56.1	72.4	67.6	67.9
[72]ASFormer, BMVC'21	90.1 / 88.8 / 79.2	84.6	79.7	84.5	85.1 / 83.4 / 76.0	79.6	85.6	81.9	76.0 / 70.6 / 57.4	75.0	73.5	70.5
[9]UARL, IJCAI'22	92.7 / 91.5 / 82.8	88.1	79.6	86.9	85.3 / 83.5 / 77.8	78.2	84.1	81.8	65.2 / 59.4 / 47.4	66.2	67.8	61.2
[47]DPRN, PR'22	92.9 / 92.0 / 82.9	90.9	82.0	88.1	87.8 / 86.3 / 79.4	82.0	87.2	84.5	75.6 / 70.5 / 57.6	75.1	71.7	70.1
[33]SEDt, EL'22	93.7 / 92.4 / 84.0	91.3	81.3	88.5	89.9 / 88.7 / 81.1	84.7	86.5	86.2	- / - / -	-	-	-
[4]TCTr, IVC'22	91.3 / 90.1 / 80.0	87.9	81.1	86.1	87.5 / 86.1 / 80.2	83.4	86.6	84.8	76.6 / 71.1 / 58.5	76.1	77.5	72.0
[19]FAMMSDTN, NPL'22	91.6 / 90.9 / 80.9	88.3	80.7	86.5	86.2 / 84.4 / 77.9	79.9	86.4	83.0	78.5 / 72.9 / 60.2	77.5	74.8	72.8
[69]DTL, NeurIPS'22	- / - / -	-	-	-	87.1 / 85.7 / 78.5	80.5	86.9	83.7	<u>78.8 / 74.5 / 62.9</u>	<u>77.7</u>	75.8	<u>73.9</u>
[6]UVAST, ECCV'22	92.7 / 91.3 / 81.0	92.1	80.2	87.5	89.1 / 87.6 / <u>81.7</u>	83.9	<u>87.4</u>	85.9	76.9 / 71.5 / 58.0	77.1	69.7	70.6
[40]BrPrompt, CVPR'22	94.1 / 92.0 / 83.0	91.6	81.2	88.4	89.2 / 87.8 / 81.3	83.8	88.1	86.0	- / - / -	-	-	-
[30]MCFM, ICIP'22	91.8 / 91.2 / 80.8	88.0	80.5	86.5	90.6 / 89.5 / 84.2	84.6	90.3	87.8	- / - / -	-	-	-
<i>DiffAct, Ours</i>	92.5 / 91.5 / 84.7	89.6	82.2	88.1	90.1 / 89.2 / 83.7	85.0	88.9	87.4	80.3 / 75.9 / 64.6	78.4	76.4	75.1

Table 1. Comparison with state-of-the-art methods. Methods in gray are not suitable for direct comparison due to the extra usage of multi-modal features [40] or hand pose features [30]. We list them here for readers’ reference. Our method achieves superior results on 50Salads and Breakfast, and comparable performance on GTEA. The average number (Avg) of the five evaluation metrics is also presented.

coder g_ψ is a re-implemented ASFormer decoder modified to be step-dependent, which adds a step embedding to the input as in [26]. The encoder has 10, 10, 12 layers and 64, 64, 256 feature maps for GTEA, 50Salads, and Breakfast respectively. We adjust the decoder to be lightweight to reduce the computational cost of the iterative denoising process, which includes 8 layers and 24, 24, 128 feature maps for the three datasets respectively. The intermediate features from encoder layers with indices 5, 7, 9 are concatenated to be the conditioning features E with $D'=768$ for Breakfast and $D'=192$ for other datasets. The encoder and decoder are trained end-to-end using Adam with a batch size of 4. The learning rate is $1e-4$ for Breakfast and $5e-4$ for other datasets. In addition to the loss \mathcal{L}^{sum} for the decoder outputs, we append a prediction head to the encoder and apply \mathcal{L}^{ce} and \mathcal{L}^{smo} as auxiliary supervision. The total steps are set as $S=1000$ and 25 steps are utilized at inference based on the sampling strategy with skipped steps [54]. The action sequences are normalized to $[-1, 1]$ when adding and removing noise in Eq. 4 and Eq. 9. All frames are processed together and all actions are predicted together, without any auto-regressive method at training or inference.

5.2. Comparison to State-of-the-Art

Table 1 presents the experimental results of our method and other recent approaches on three datasets. Our proposed method advances the state-of-the-art by an evident margin on 50Salads and Breakfast, and achieves comparable performance on GTEA. Specifically, the average performance is improved from 86.2 to 87.4 on 50Salads and from 73.9 to 75.1 on Breakfast. On the smallest dataset

GTEA, our method obtains similar overall performance with higher accuracy and F1@50. The results show the effectiveness of our diffusion-based action segmentation as a new framework and its particular advantage on large or complex datasets. It is also promising to combine more recent backbones such as SEDT [33] and DPRN [47] into our framework to further improve the results.

5.3. Ablation Studies

Extensive ablation studies are performed to validate the design choices in our method. We select the 50Salads dataset for ablation studies because of its substantial complexity and proper data size.

Effect of Prior Modeling. To inspect the impact of prior modeling, experiments are conducted in Table 2 with different combinations of condition masking schemes. It is observed that our method reaches the best performance when all three priors are considered. Notably, the position prior is especially useful among the three priors.

Effect of Training Losses. In Table 3, we investigate the effect of loss functions, where each of the following losses is adopted, the full \mathcal{L}^{sum} loss, the \mathcal{L}^{sum} without \mathcal{L}^{bd} , the \mathcal{L}^{sum} without \mathcal{L}^{smo} , and the vanilla \mathcal{L}^{ce} loss. It is found that all the loss components are necessary for the best result. Our proposed boundary alignment loss \mathcal{L}^{bd} brings performance gain in terms of both frame-wise accuracy and temporal continuity on top of \mathcal{L}^{ce} and \mathcal{L}^{smo} .

Effect of Inference Steps. Experiment results using different number of inference steps are reported in Table 4, from which we can notice a steady increase in performance, with diminishing marginal benefits, as the step number gets

N	P	B	R	F1@{10, 25, 50}	Edit	Acc	Avg
✓				89.0 / 88.1 / 82.4	83.7	88.1	86.3
✓	✓			89.9 / 88.9 / 82.8	84.3	88.2	86.8
✓		✓		89.7 / 88.6 / 82.6	83.9	88.2	86.6
✓			✓	89.6 / 88.7 / 82.7	84.0	88.0	86.6
✓		✓	✓	89.4 / 88.7 / 83.0	84.4	88.2	86.7
✓	✓		✓	90.0 / 88.8 / 83.4	84.4	88.8	87.1
✓	✓	✓		90.2 / 89.3 / 83.6	84.6	88.5	87.2
✓	✓	✓	✓	90.1 / 89.2 / 83.7	85.0	88.9	87.4

Table 2. Ablation study on the prior modeling. N: Baseline. P: Position prior. B: Boundary prior. R: Relation prior. For each row, a scheme is randomly selected from the ticked ones at each training iteration.

\mathcal{L}^{ce}	\mathcal{L}^{smo}	\mathcal{L}^{bd}	F1@{10, 25, 50}	Edit	Acc	Avg
✓			86.7 / 85.3 / 79.2	80.8	87.0	83.8
✓	✓		89.8 / 88.9 / 83.1	84.0	88.8	86.9
✓		✓	86.9 / 86.1 / 78.7	81.0	85.4	83.6
✓	✓	✓	90.1 / 89.2 / 83.7	85.0	88.9	87.4

Table 3. Ablation study on the loss functions.

Steps	F1@{10, 25, 50}	Edit	Acc	Avg
1	64.9 / 63.8 / 59.3	56.5	88.6	66.6
2	81.7 / 80.5 / 75.5	74.5	88.9	80.2
4	87.6 / 86.6 / 81.2	82.1	89.1	85.3
8	89.3 / 88.3 / 83.1	83.5	89.0	86.6
16	90.0 / 88.8 / 83.3	84.5	89.0	87.1
25	90.1 / 89.2 / 83.7	85.0	88.9	87.4
50	90.4 / 89.5 / 84.0	85.3	89.0	87.6
100	90.4 / 89.7 / 84.3	85.3	88.9	87.7

Table 4. Ablation study on the number of inference steps.

Features	F1@{10, 25, 50}	Edit	Acc	Avg
Input Features F	82.5 / 80.6 / 72.3	75.7	82.5	78.7
h_ϕ Layer 5	90.3 / 89.2 / 83.9	85.1	89.1	87.5
h_ϕ Layer 7	90.4 / 89.4 / 83.4	85.0	88.8	87.4
h_ϕ Layer 9	90.0 / 89.0 / 83.4	84.4	88.8	87.1
h_ϕ Layer 5,7,9	90.1 / 89.2 / 83.7	85.0	88.9	87.4
h_ϕ Prediction	90.3 / 89.3 / 83.3	84.6	87.8	87.1

Table 5. Ablation study on the conditioning features.

larger. The computation grows linearly with the step number. We leverage 25 steps to keep a good balance between the performance and the computational cost.

Effect of Conditioning Features. For the condition of generation, the input video features F and the features from different layers of the encoder h_ϕ are explored as in Table 5. The performance drops remarkably when using the input feature F as the condition, suggesting the necessity of an encoder. On the other hand, the performance is not sensitive to which encoder layer the features are extracted from.

Effect of the Backbone. The choices of the encoder and decoder in *DiffAct* are flexible. Therefore, we change our backbone to MS-TCN [20] to show such flexibility. In detail, a single-stage TCN is directly used as the encoder and is modified with the step embedding to be the decoder.

Method	F1@{10, 25, 50}	Edit	Acc	Avg
[20]MS-TCN	76.3 / 74.0 / 64.5	67.9	80.7	72.7
[41]MS-TCN++	80.7 / 78.5 / 70.1	74.3	83.7	77.5
[1]HASR (MS-TCN)	83.4 / 81.8 / 71.9	77.4	81.7	79.2
[31]ASRF	84.9 / 83.5 / 77.3	79.3	84.5	81.9
[69]DTL (MS-TCN)	78.3 / 76.5 / 67.6	70.5	81.5	74.9
<i>DiffAct</i> (MS-TCN)	86.9 / 85.3 / 79.4	80.3	88.2	84.0

Table 6. Results on 50Salads using MS-TCN backbone.

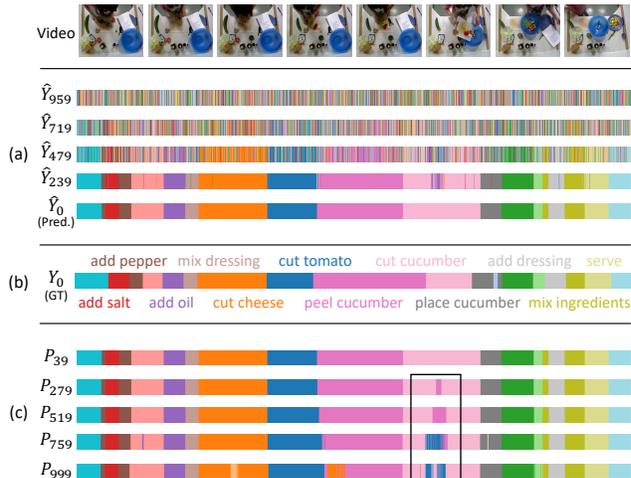


Figure 4. Visualization of the iterative denoising process. The ground truth is presented in (b), where some segments are marked with class labels. The (a) and (c) respectively plot the inference trajectory \hat{Y}_s and the denoised sequences P_s at different steps (Eq. 9). The video is ‘rgb-01-2’ from 50Salads.

Table 6 compares our results to recent methods with MS-TCN backbones, which show the superiority of our method.

5.4. Qualitative Result and Computational Cost

Qualitative Result. To illustrate the refinement process along the denoising steps, the step-wise results for a video from 50Salads are visualized in Fig. 4. The model refines an initial random noise sequence to generate the final action prediction in an iterative manner. For example, as in the black box in Fig. 4, the segment of ‘cut cucumber’ is broken up by ‘cut tomato’ and ‘peel cucumber’, which share similar visual representations. After a number of iterations, the relation between these actions is constructed and the error is gradually corrected. Finally, a continuous segment of ‘cut cucumber’ can be properly predicted.

Computational Cost. Table 7 compares the computational costs of our method and its backbone ASFormer [72]. Our method, which is equipped with a lightweight decoder, largely outperforms ASFormer with fewer FLOPs at inference when using 8 steps. Using 25 steps, our method further improves the result at an acceptable overhead.

Method	Avg	#params	FLOPs	Mem.	Time
ASFormer [72]	81.9	1.134M	6.66G	3.5G	2.38s
<i>DiffAct</i> (8 Steps)	86.6	0.975M	4.96G	1.9G	0.68s
<i>DiffAct</i> (16 Steps)	87.1	0.975M	7.73G	1.9G	1.30s
<i>DiffAct</i> (25 Steps)	87.4	0.975M	10.85G	1.9G	2.09s

Table 7. Computational cost comparison. The number of parameters, the average FLOPs at inference, the GPU memory cost during training, the average inference time, and the average performance (Avg) on 50Salads for our method and ASFormer.

Masking	F1@{10, 25, 50}	Edit	Acc	Avg
N	90.1 / 89.2 / 83.7	85.0	88.9	87.4
P	25.7 / 21.5 / 11.6	34.8	20.9	22.9
B	89.4 / 88.6 / 83.0	84.1	88.4	86.7
R	88.9 / 87.8 / 81.7	83.5	87.2	85.8

Table 8. Results on 50Salads using different condition masking types at the inference stage. Note that this is only for analysis purposes since the mask types B and R depend on the ground truth. The model performance is maintained at a reasonable level using different masks, suggesting the action priors are well captured.

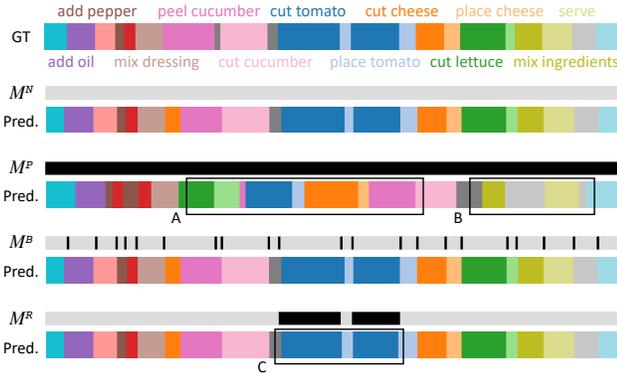


Figure 5. Visualization of the masks and the corresponding predictions using the masked conditions at inference. The video is ‘rgb-03-2’ from 50Salads. In M^N , M^P , M^B , M^R , masked locations are colored in black. More results using M^P at inference and further discussions are given in the supplementary material.

6. Discussion

Analysis of the Prior Modeling. In this section, an exploratory experiment is performed to analyze to what extent the position prior, boundary prior, and relation prior are captured in our model. Recall that the proposed method uses no masking (M^N) at inference by default, here in this experiment, we input the masked conditions with each masking type (M^P , M^B , M^R) for inference instead. As in Table 8, the model can still achieve reasonably good performance when the mask M^B or mask M^R is applied, indicating that the boundary prior and relation prior are well handled. It is also interesting to discover that the result using the completely masked condition (M^P), which has a 34.8 edit score, is much better than the random guess. This reveals that the

model has learned meaningful correlations between actions and time locations via our position prior modeling. We further visualize in Fig. 5 the condition masks and the corresponding action predictions for a video when each mask type is applied at inference. It is clear that the model produces a *generally plausible* action sequence when all the features are blocked by M^P . For example, the actions of cutting and placing ingredients are located in the middle of the video (Fig. 5 A), while the actions of mixing and serving occur at the end (Fig. 5 B). With mask M^B , the model is still able to find action boundaries. The missing action ‘cut tomato’ masked by M^R is successfully restored at Fig. 5 C. These analyses demonstrate the capability of our method in prior modeling.

Limitation and Future Work. One limitation of the proposed method is that its advantage on the small-scale dataset, GTEA, is not as significant as on large datasets. We speculate that it is more difficult to generatively learn the distribution of action sequences given only a few videos, which leads to a lower edit score. Note that this is not a problem on large datasets on which the model makes clear gains in terms of the edit score in Table 1. Potential remedies on small data include model reassembly [71] or replacing the Gaussian noise in the diffusion process with some perturbations based on the statistics of actions, *e.g.*, transforming the distribution towards the mean sequence obtained from the training set, to reduce the hypothesis space and thus the amount of data required. Future works can also combine the generation of frame-wise action sequences and segment-wise ordered action lists jointly in our diffusion-based action segmentation. It is also promising to extend the current framework for unified action segmentation and action anticipation in the future since our generative framework is intuitively appropriate for the anticipation task. We share other early attempts in the supplementary.

7. Conclusion

This paper proposes a new framework for temporal action segmentation which generates action sequences through an iterative denoising process. A flexible condition masking strategy is designed to jointly exploit the position prior, the boundary prior, and the relation prior of human actions. With its nature of iterative refinement, its ability of generative modeling, and its enhancement of the three action priors, the proposed framework achieves state-of-the-art results on benchmark datasets, unlocking new possibilities for action segmentation.

Acknowledgement. This work was supported in part by the Australian Research Council under Project DP210101859 and the University of Sydney Research Accelerator (SOAR) Prize. The training platforms supporting this work were provided by High-Flyer AI and National Computational Infrastructure Australia.

References

- [1] Hyemin Ahn and Dongheui Lee. Refining action segmentation with hierarchical video representations. In *ICCV*, 2021. 2, 6, 7
- [2] Tomer Amit, Eliya Nachmani, Tal Shaharbany, and Lior Wolf. SegDiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390*, 2021. 3
- [3] Evlampios Apostolidis, Eleni Adamantidou, Alexandros I Metsai, Vasileios Mezaris, and Ioannis Patras. Video summarization using deep neural networks: A survey. *Proceedings of the IEEE*, 2021. 1
- [4] Nicolas Aziere and Sinisa Todorovic. Multistage temporal convolution transformer for action segmentation. *Image and Vision Computing*, 2022. 2, 6
- [5] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khulkov, and Artem Babenko. Label-efficient semantic segmentation with diffusion models. *ICLR*, 2021. 3
- [6] Nadine Behrmann, S Alireza Golestaneh, Zico Kolter, Jürgen Gall, and Mehdi Noroozi. Unified fully and timestamp supervised temporal action segmentation via sequence to sequence translation. In *ECCV*, 2022. 2, 5, 6
- [7] Ankan Kumar Bhunia, Salman Khan, Hisham Cholakkal, Rao Muhammad Anwer, Jorma Laaksonen, Mubarak Shah, and Fahad Shahbaz Khan. Person image synthesis via denoising diffusion model. *arXiv preprint arXiv:2211.12500*, 2022. 3
- [8] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the Kinetics dataset. In *CVPR*, 2017. 5
- [9] Lei Chen, Muheng Li, Yueqi Duan, Jie Zhou, and Jiwen Lu. Uncertainty-aware representation learning for action segmentation. In *IJCAI*, 2022. 2, 6
- [10] Min-Hung Chen, Baopu Li, Yingze Bao, and Ghassan Al-Regib. Action segmentation with mixed temporal domain adaptation. In *WACV*, 2020. 2, 6
- [11] Min-Hung Chen, Baopu Li, Yingze Bao, Ghassan Al-Regib, and Zsolt Kira. Action segmentation with joint self-supervised temporal domain adaptation. In *CVPR*, 2020. 2, 6
- [12] Shoufa Chen, Peize Sun, Yibing Song, and Ping Luo. DiffusionDet: Diffusion model for object detection. *arXiv preprint arXiv:2211.09788*, 2022. 3
- [13] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *arXiv preprint arXiv:2209.04747*, 2022. 1, 3
- [14] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat GANs on image synthesis. *NeurIPS*, 2021. 1, 3
- [15] Guodong Ding, Fadime Sener, and Angela Yao. Temporal action segmentation: An analysis of modern technique. *arXiv preprint arXiv:2210.10352*, 2022. 1, 2
- [16] Li Ding and Chenliang Xu. Tricorner: A hybrid temporal convolutional and recurrent network for video action segmentation. *arXiv preprint arXiv:1705.07818*, 2017. 2
- [17] Anh-Dung Dinh, Daochang Liu, and Chang Xu. PixelAs-Param: A gradient view on diffusion sampling with guidance. In *ICML*, 2023. 3
- [18] Dazhao Du, Bing Su, Yu Li, Zhongang Qi, Lingyu Si, and Ying Shan. Do we really need temporal convolutions in action segmentation? *arXiv preprint arXiv:2205.13425*, 2022. 2
- [19] Zexing Du and Qing Wang. Dilated transformer with feature aggregation module for action segmentation. *Neural Processing Letters*, 2022. 2, 6
- [20] Yazan Abu Farha and Jurgen Gall. MS-TCN: Multi-stage temporal convolutional network for action segmentation. In *CVPR*, 2019. 1, 2, 4, 5, 7
- [21] Alireza Fathi, Xiaofeng Ren, and James M Rehg. Learning to recognize objects in egocentric activities. In *CVPR*, 2011. 2, 5
- [22] Harshala Gammulle, Tharindu Fernando, Simon Denman, Sridha Sridharan, and Clinton Fookes. Coupled generative adversarial network for continuous fine-grained action segmentation. In *WACV*, 2019. 2
- [23] Shang-Hua Gao, Qi Han, Zhong-Yu Li, Pai Peng, Liang Wang, and Ming-Ming Cheng. Global2local: Efficient structure search for video action segmentation. In *CVPR*, 2021. 2, 6
- [24] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *CVPR*, 2022. 3
- [25] Basavaraj Hampiholi, Christian Jarvers, Wolfgang Mader, and Heiko Neumann. Depthwise separable temporal convolutional network for action segmentation. In *3DV*, 2020. 2
- [26] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020. 1, 3, 6
- [27] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. 5
- [28] Tobias Höppe, Arash Mehrjou, Stefan Bauer, Didrik Nielsen, and Andrea Dittadi. Diffusion models for video prediction and infilling. *arXiv preprint arXiv:2206.07696*, 2022. 3
- [29] Yifei Huang, Yusuke Sugano, and Yoichi Sato. Improving action segmentation via graph-based temporal reasoning. In *CVPR*, 2020. 2, 6
- [30] Kenta Ishihara, Gaku Nakano, and Tetsuo Inoshita. MCFM: Mutual cross fusion module for intermediate fusion-based action segmentation. In *ICIP*, 2022. 6
- [31] Yuchi Ishikawa, Seito Kasai, Yoshimitsu Aoki, and Hirokatsu Kataoka. Alleviating over-segmentation errors by detecting action boundaries. In *WACV*, 2021. 2, 5, 6, 7
- [32] Gwanghyun Kim and Jong Chul Ye. DiffusionClip: Text-guided image manipulation using diffusion models. In *CVPR*, 2022. 3
- [33] Gyeong-hyeon Kim and Eunwoo Kim. Stacked encoder-decoder transformer with boundary smoothing for action segmentation. *Electronics Letters*, 2022. 2, 6
- [34] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *CVPR*, 2014. 2, 5
- [35] Max WY Lam, Jun Wang, Dan Su, and Dong Yu. BDDM: Bilateral denoising diffusion models for fast and high-quality speech synthesis. *arXiv preprint arXiv:2203.13508*, 2022. 3

- [36] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional networks for action segmentation and detection. In *CVPR*, 2017. 2
- [37] Colin Lea, Austin Reiter, René Vidal, and Gregory D Hager. Segmental spatiotemporal CNNs for fine-grained action segmentation. In *ECCV*, 2016. 2
- [38] Peng Lei and Sinisa Todorovic. Temporal deformable residual networks for action segmentation in videos. In *CVPR*, 2018. 2
- [39] Yichong Leng, Zehua Chen, Junliang Guo, Haohe Liu, Jiawei Chen, Xu Tan, Danilo Mandic, Lei He, Xiang-Yang Li, Tao Qin, et al. Binauralgrad: A two-stage conditional diffusion probabilistic model for binaural audio synthesis. *arXiv preprint arXiv:2205.14807*, 2022. 3
- [40] Muheng Li, Lei Chen, Yueqi Duan, Zhilan Hu, Jianjiang Feng, Jie Zhou, and Jiwen Lu. Bridge-prompt: Towards ordinal action understanding in instructional videos. In *CVPR*, 2022. 2, 5, 6
- [41] Shi-Jie Li, Yazan AbuFarha, Yun Liu, Ming-Ming Cheng, and Juergen Gall. MS-TCN++: Multi-stage temporal convolutional network for action segmentation. *IEEE TPAMI*, 2020. 1, 2, 4, 5, 6, 7
- [42] Yunheng Li, Zhuben Dong, Kaiyuan Liu, Lin Feng, Lianyu Hu, Jie Zhu, Li Xu, Shenglan Liu, et al. Efficient two-step networks for temporal action segmentation. *Neurocomputing*, 2021. 2
- [43] Daochang Liu, Qiyue Li, Tingting Jiang, Yizhou Wang, Rulin Miao, Fei Shan, and Ziyu Li. Towards unified surgical skill assessment. In *CVPR*, 2021. 1
- [44] Zhichao Liu, Leshan Wang, Desen Zhou, Jian Wang, Songyang Zhang, Yang Bai, Errui Ding, and Rui Fan. Temporal segment transformer for action segmentation. *arXiv preprint arXiv:2302.13074*, 2023. 2
- [45] Calvin Luo. Understanding diffusion models: A unified perspective. *arXiv preprint arXiv:2208.11970*, 2022. 3
- [46] Khoi-Nguyen C Mac, Dhiraj Joshi, Raymond A Yeh, Jinjun Xiong, Rogerio S Feris, and Minh N Do. Learning motion in feature space: Locally-consistent deformable convolution networks for fine-grained action detection. In *ICCV*, 2019. 2
- [47] Junyong Park, Daekyung Kim, Sejoon Huh, and Sungho Jo. Maximization and restoration: Action segmentation through dilation passing and temporal reconstruction. *Pattern Recognition*, 2022. 2, 6
- [48] Konpat Preechakul, Nattanat Chatthee, Suttisak Widadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *CVPR*, pages 10619–10629, 2022. 3
- [49] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 3
- [50] Fadime Sener, Dipika Singhania, and Angela Yao. Temporal aggregate representations for long-range video understanding. In *ECCV*, 2020. 2
- [51] Bharat Singh, Tim K Marks, Michael Jones, Oncel Tuzel, and Ming Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *CVPR*, 2016. 2
- [52] Dipika Singhania, Rahul Rahaman, and Angela Yao. Coarse to fine multi-resolution temporal convolutional network. *arXiv preprint arXiv:2105.10859*, 2021. 2
- [53] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. 3
- [54] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *ICLR*, 2021. 1, 3, 5, 6
- [55] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *NeurIPS*, 2019. 3
- [56] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *NeurIPS*, 2020. 3
- [57] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 3
- [58] Yaser Souri, Yazan Abu Farha, Fabien Despinoy, Gianpiero Francesca, and Juergen Gall. FIFA: Fast inference approximation for action segmentation. In *DAGM German Conference on Pattern Recognition*, 2021. 2
- [59] Sebastian Stein and Stephen J McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, 2013. 2, 5
- [60] Lorin Sweeney, Graham Healy, and Alan F Smeaton. Diffusing surrogate dreams of video scenes to predict video memorability. *arXiv preprint arXiv:2212.09308*, 2022. 3
- [61] Xiaoyan Tian, Ye Jin, and Xianglong Tang. Local-Global transformer neural network for temporal action segmentation. *Multimedia Systems*, 2022. 2
- [62] Sarvesh Vishwakarma and Anupam Agrawal. A survey on activity recognition and behavior understanding in video surveillance. *The Visual Computer*, 2013. 1
- [63] Vikram Voleti, Alexia Jolicoeur-Martineau, and Christopher Pal. Masked conditional video diffusion for prediction, generation, and interpolation. *arXiv preprint arXiv:2205.09853*, 2022. 3
- [64] Dong Wang, Yuan Yuan, and Qi Wang. Gated forward refinement network for action segmentation. *Neurocomputing*, 2020. 2
- [65] Jiahao Wang, Zhengyin Du, Annan Li, and Yunhong Wang. Atrous temporal convolutional network for video action segmentation. In *ICIP*, 2019. 2
- [66] Jiahui Wang, Zhenyou Wang, Shanna Zhuang, and Hui Wang. Cross-enhancement transformer for action segmentation. *arXiv preprint arXiv:2205.09445*, 2022. 2
- [67] Yunke Wang, Xiyu Wang, Anh-Dung Dinh, Bo Du, and Chang Xu. Learning to schedule in diffusion probabilistic models. In *KDD*, 2023. 3
- [68] Zhenzhi Wang, Ziteng Gao, Limin Wang, Zhifeng Li, and Gangshan Wu. Boundary-aware cascade networks for temporal action segmentation. In *ECCV*, 2020. 1, 2, 6
- [69] Ziwei Xu, Yogesh S Rawat, Yongkang Wong, Mohan Kankanhalli, and Mubarak Shah. Don't pour cereal into coffee: Differentiable temporal logic for temporal action segmentation. In *NeurIPS*, 2022. 2, 5, 6, 7

- [70] Ruihan Yang, Prakhar Srivastava, and Stephan Mandt. Diffusion probabilistic modeling for video generation. *arXiv preprint arXiv:2203.09481*, 2022. 3
- [71] Xingyi Yang, Daquan Zhou, Songhua Liu, Jingwen Ye, and Xinchao Wang. Deep model reassembly. In *NeurIPS*, 2022. 8
- [72] Fangqiu Yi, Hongyu Wen, and Tingting Jiang. ASFormer: Transformer for action segmentation. In *BMVC*, 2021. 1, 2, 5, 6, 7, 8
- [73] Peiyu Yu, Sirui Xie, Xiaojian Ma, Baoxiong Jia, Bo Pang, Ruigi Gao, Yixin Zhu, Song-Chun Zhu, and Ying Nian Wu. Latent diffusion energy-based model for interpretable text modeling. *arXiv preprint arXiv:2206.05895*, 2022. 3
- [74] Junbin Zhang, Pei-Hsuan Tsai, and Meng-Hsun Tsai. Semantic2graph: Graph-based multi-modal feature for action segmentation in videos. *arXiv preprint arXiv:2209.05653*, 2022. 2
- [75] Yunlu Zhang, Keyan Ren, Chun Zhang, and Tong Yan. SG-TCN: Semantic guidance temporal convolutional network for action segmentation. In *IJCNN*, 2022. 2
- [76] Xian Zhong, Zipeng Li, Shuqin Chen, Kui Jiang, Chen Chen, and Mang Ye. Refined semantic enhancement towards frequency diffusion for video captioning. *arXiv preprint arXiv:2211.15076*, 2022. 3