

# Self-Supervised Contrastive Learning for Videos using Differentiable Local Alignment

Keyne Oei<sup>1</sup>  
s8keoeii@uni-saarland.de

Amr Gomaa<sup>2</sup>  
amr.gomaa@dfki.de

Anna Maria Feit<sup>1</sup>  
feit@cs.uni-saarland.de

João Belo<sup>1</sup>  
jbelo@cs.uni-saarland.de

<sup>1</sup> Universität des Saarlandes  
Saarbrücken, Germany

<sup>2</sup> German Research Center for Artificial  
Intelligence (DFKI)  
Saarbrücken, Germany

---

## Abstract

Robust frame-wise embeddings are essential to perform video analysis and understanding tasks. We present a self-supervised method for representation learning based on aligning temporal video sequences. Our framework uses a transformer-based encoder to extract frame-level features and leverages them to find the optimal alignment path between video sequences. We introduce the novel Local-Alignment Contrastive (LAC) loss, which combines a differentiable local alignment loss to capture local temporal dependencies with a contrastive loss to enhance discriminative learning. Prior works on video alignment have focused on using global temporal ordering across sequence pairs, whereas our loss encourages identifying the best-scoring subsequence alignment. LAC uses the differentiable Smith-Waterman (SW) affine method, which features a flexible parameterization learned through the training phase, enabling the model to adjust the temporal gap penalty length dynamically. Evaluations show that our learned representations outperform existing state-of-the-art approaches on action recognition tasks.

## 1 Introduction

Video understanding and analysis have garnered significant attention in computer vision research in recent years [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]. These tasks require models to capture not only spatial information within individual frames but also the complex temporal dynamics that evolve across sequences of frames. Traditional supervised learning methods have achieved considerable success in advancing performance; models such as those presented in [1, 2, 3] are able to predict action categories effectively. However, these methods present several challenges: (i) Most approaches require fine-grained annotations, which are time-consuming and labor-intensive; (ii) Most approaches struggle to capture the complex temporal information and causal relationships that vary across diverse videos. To address these challenges, we explore how representation learning based on the pretext-task of video




Figure 1: We introduce a representation learning approach that aligns video sequences depicting the same processes. Our training objective is to use a novel LAC loss to optimize and learn an element-wise embedding function that supports the alignment process.

alignment can improve performance in downstream action recognition tasks such as phase classification and phase progression.

Prior self-supervised or weakly-supervised video alignment approaches [2, 5, 10, 24, 15, 27] train on pairs of videos that describe the same action using either cycle-consistency [10], soft dynamic time warping [24, 15], or sequential contrastive learning [5]. We notice that the pretext task of video alignment shares similarities with bioinformatics sequence alignment for studying DNA, RNA, and protein structures. Since some actions are composed of sub-actions or events that occur in a specific order, video alignment shares similarities with protein alignment, as both involve identifying regions of continuity and discontinuity within sequences. There are two main types of alignment algorithms commonly used in bioinformatics: *global* and *local*. The Needleman-Wunsch (NW) algorithm [25], a global alignment method, aligns entire sequences from start to finish. This approach is conceptually similar to Dynamic Time Warping (DTW) [24], which aligns sequences of varying lengths by matching their temporal patterns from start to finish. Differentiable versions of NW and DTW algorithms have been established by incorporating smooth maximum/minimum and argmax/argmin functions [2, 23]. In particular, Soft-DTW [2] has been used in prior video alignment approaches [2, 24, 15]. In contrast, the Smith-Waterman (SW) algorithm [29] implements local alignment, identifying regions of high similarity within sequences and allowing for gaps, insertions, and deletions. This ability to handle subsequences makes local alignment particularly suitable for modeling the complexity of actions in real-world video data, where actions may not always follow rigid, continuous sequences.

Specifically, in this work, we introduce a novel loss function termed *Local-Alignment Contrastive* (LAC) loss, integrated into our proposed end-to-end framework for the video alignment task, as visualized in Figure 1. Following prior works [2, 5], our framework uses a variation of convolutional and transformer encoders to extract spatio-temporal features from each frame. However, rather than relying on a differentiable global alignment loss like Soft-DTW [2], our novel LAC loss enables the comparison of video pairs through a consistent and differentiable local alignment loss. This loss specifically accommodates sequences with learned penalties for opening gaps and extending them, combined with a contrastive loss that effectively separates dissimilar frames. We demonstrate that the learned representations of our approach outperform previous methods on various downstream action recognition tasks, such as action phase classification and action phase progression on the Pouring [27] and PennAction [53] datasets. Code and models are available at: <https://github.com/keynekassapa13/LAC>.

## 2 Related Work

**Action Recognition.** Over the last decade, researchers have proposed various methods to address video recognition tasks [4, 12, 13, 30, 35], including action classification [4, 4, 11, 18, 28, 31], which involves mapping a video to an action category, and action segmentation [9, 12, 21, 32, 34], which aims to identify the spatial and temporal boundaries of specific actions within a video. However, these methods often rely heavily on extensive labeled data, which is time-consuming and labor-intensive to obtain. To address this limitation, self-supervised learning methods have been explored, enabling the extraction of video representations without requiring labeled data. Video alignment, as a self-supervised pretext task for action recognition, leverages temporal consistency to learn meaningful representations. Previous work has utilized traditional computer vision techniques to derive these representations; for instance, optical flow [8] computes motion between frames based on brightness constancy, the CONDENSATION algorithm [17] employs a probabilistic approach informed by prior motion data, and Space-Time Interest Points (STIPs) [20] detect salient points by analyzing both spatial and temporal patterns in video sequences.

**Self-Supervised Learning.** With the rise of deep learning, the field has shifted towards using Self-Supervised Learning (SSL) methods to learn video representations. Prior works [6, 10, 14, 15, 27] have addressed the problem of video alignment using self-supervised methods. Time Contrastive Network (TCN) [27] utilizes contrastive learning to distinguish frames from different segments while grouping those within the same segment. Temporal Cycle-Consistency (TCC) [10] introduces cycle-consistency loss to identify and match recurring action sequences within or across videos. Learning by Aligning Videos (LAV) [15] uses Soft-DTW [7] and Inverse Difference Moment (IDM) regularization to optimize alignment and ensure balanced frame distribution. Global Temporal Alignment (GTA) [14] implements a modified differentiable DTW with global consistency loss for consistent temporal alignment. Contrastive Action Representation Learning (CARL) [6] introduces Sequence Contrastive Loss (SCL), which minimizes the KL-divergence between augmented views based on timestamp distance. Learning Representation by position PROPagation (LRProp) [2] utilizes a Soft-DTW [7] alignment path and a novel pair-wise position propagation technique. We will use these state-of-the-art methods to evaluate our model’s performance on action recognition tasks.

**Global vs. Local Alignment.** Global alignment methods, such as the NW algorithm [25] and DTW [24], aim to align entire sequences from start to finish to achieve the best overall match; differentiable versions have been developed [2, 23] and utilized in prior video alignment approaches [2, 14, 15]. In contrast, the SW algorithm [29] focuses on identifying the most similar subsequences within larger sequences. Recent work by Petti et al. [26] introduced a differentiable SW algorithm in a supervised framework that jointly learns alignments and improves protein structure prediction. [27] shows the power of embedding sequences into high-dimensional spaces within a supervised framework for more precise local alignment. Our approach employs a differentiable version of SW in self-supervised framework to improve model performance on video understanding and alignment tasks compared to state-of-the-art temporal global alignment methods.




Figure 2: Our framework uses Local Alignment Loss and Contrastive Loss to optimize embeddings generated by an encoder that processes input videos after they have undergone spatio-temporal data augmentation.

## 3 Method

This section introduces our novel Local-Alignment Contrastive (LAC) loss function employed within our framework where video alignment serves as the pretext task, as visualized in Fig. 2. Specifically, we derive the forward and backward passes of the differentiable local alignment loss and explain how we ensure consistency within our method.

### 3.1 Notations

Let  $V^1$  and  $V^2$  represent two videos, each consisting of a sequence of frames. Specifically,  $V^1 = \{v_t^1\}_{t=1}^{T_1}$  and  $V^2 = \{v_t^2\}_{t=1}^{T_2}$ , where  $v_t^1$  and  $v_t^2$  denote individual frames at time  $t$ , and  $T_1$  and  $T_2$  are the total number of frames in each video. Following [9], we preprocess the videos using temporal random cropping to generate cropped versions  $V_c^1$  and  $V_c^2$ . The indices of the sampled frames during cropping are recorded in  $S^1$  and  $S^2$ , where  $S^1 \subseteq \{1, \dots, T_1\}$  and  $S^2 \subseteq \{1, \dots, T_2\}$ . Next, we apply data augmentations to the cropped videos, resulting in  $\tilde{V}^1$  and  $\tilde{V}^2$ . These augmented videos are then processed through an encoder to obtain embeddings  $Z^1$  and  $Z^2$ . Each embedding  $Z^i \in \mathbb{R}^{T_i \times E}$  represents a sequence of length  $T_i$  with an embedding dimension  $E$ . The objective is to train the encoder to minimize the distance between embeddings  $Z^1$  and  $Z^2$  corresponding to videos depicting the same action, as determined by our loss.

### 3.2 Background

Prior approaches [4, 5] use DTW as the main method for the temporal alignment loss. Given embeddings  $Z^1$  and  $Z^2$ , we compute a cost matrix  $C \in \mathbb{R}^{T_1 \times T_2}$ , where each element  $C_{i,j}$  measures the local distance between  $Z_i^1$  and  $Z_j^2$ . DTW seeks the optimal alignment path through this cost matrix that minimizes the total cumulative cost. The feasible paths must satisfy constraints of matching endpoints, monotonicity, and continuity. While the number of possible alignment paths grows exponentially with the lengths of the sequences, dynamic programming allows DTW to compute the optimal path efficiently in quadratic time  $O(T_1 T_2)$ . The DTW algorithm is formulated recursively as:

$$D_{i,j} = C_{i,j} + \min \begin{cases} D_{i-1,j-1} \\ D_{i-1,j} \\ D_{i,j-1} \end{cases} \quad (1)$$

where  $D_{i,j}$  represents the accumulated cost up to point  $(i, j)$ , and  $C_{i,j}$  is the local cost between elements  $Z_i^1$  and  $Z_j^2$ . The non-differentiable nature of the DTW formulation presents challenges for gradient-based optimization techniques. To address this, prior methods [14, 15] utilize Soft-DTW [7] which replaces the discrete min operator with a differentiable, smoothed version denoted as  $\text{softmin}^\gamma$ . As the smoothing parameter  $\gamma$  approaches zero, the behavior of  $\text{softmin}^\gamma$  approximates that of the discrete min operator. This substitution enables the use of gradient-based optimization in DTW while preserving its alignment capabilities.

### 3.3 Differentiable Local-Alignment

Despite the advantages offered by Soft-DTW [7], it inherently maintains the global alignment characteristic, ensuring continuity throughout the entire sequence. This attribute of global alignment may prove suboptimal in scenarios where localized alignment strategies, such as those implemented by the SW algorithm [19], are preferable. The SW algorithm [19] supports localized matching, providing enhanced flexibility and increased sensitivity to similarities within subsequences. Therefore, we propose a differentiable SW local-alignment approach as our method for temporal alignment loss.

We consider a pair of embeddings,  $(Z^1, S^1)$  and  $(Z^2, S^2)$ , with corresponding sampled indices, tasked with determining the optimal alignment between these embeddings. We begin by calculating a similarity matrix,  $S \in \mathbb{R}^{T \times T}$ , using the inverted Euclidean distance to yield higher values that indicate increased similarity between the points in  $Z^1$  and  $Z^2$ . Utilizing dynamic programming, the SW algorithm effectively determines the optimal alignment by maximizing a score derived from pairwise similarities within the matrix. The alignment algorithm is formally expressed as follows:

$$D_{i,j} = S_{i,j} + \max \begin{cases} 0 \\ D_{i-1,j-1} \\ I_{x_{i-1,j-1}} \\ I_{y_{i-1,j-1}} \end{cases} \quad (2)$$

$$I_{x_{i,j}} = \max \begin{cases} D_{i,j-1} - g_{o_{i,j}} \\ I_{x_{i,j-1}} - g_{e_{i,j}} \end{cases} \quad (3)$$

$$I_{y_{i,j}} = \max \begin{cases} D_{i-1,j} - g_{o_{i,j}} \\ I_{x_{i-1,j}} - g_{o_{i,j}} \\ I_{y_{i-1,j}} - g_{e_{i,j}} \end{cases} \quad (4)$$

where the initial values are set to  $D(i, 0) = D(0, j) = I_x(i, 0) = I_x(0, j) = I_y(i, 0) = I_y(0, j) = -\infty$  for all indices  $i, j = 1, \dots, T$ . The matrix  $D$  described in Eq. 2 acts as the primary scoring matrix, representing the cumulative maximum score at each matrix coordinate  $(i, j)$ . The matrices  $I_x$  and  $I_y$  detailed in Eq. 3 and Eq. 4, represent the scores associated with introducing gaps along the  $x$  and  $y$  axes. Additionally,  $g_o$  and  $g_e$  are used to specify the penalties for opening and extending gaps.

Similar to Soft-DTW [7], our differentiable temporal alignment loss function uses a smoothed version of the discrete max function, denoted as  $\text{max}^\gamma$ . This smooth approximation is formally defined as:

$$\max^\gamma(u_1, \dots, u_n) := \gamma \log \sum_{i=1}^n e^{u_i/\gamma} \quad (5)$$

where  $\gamma > 0$  is a smoothing parameter that controls the approximation’s fidelity to the original max function. As  $\gamma$  approaches zero, the smoothed  $\max^\gamma$  function increasingly approximates the behavior of the standard max function, enabling a differentiable formulation that can be integrated into gradient-based optimization frameworks. The gradient of the smoothed  $\max^\gamma$  with respect to each input  $u_i$  is calculated as follows:

$$\frac{\partial \max^\gamma(\{u_1, \dots, u_n\})}{\partial u_i} = \exp\left(\frac{u_i - \max^\gamma(\{u_1, \dots, u_n\})}{\gamma}\right) \quad (6)$$

This expression aligns with the softmax function. By differentiating the smoothed maximum score with respect to each input, our model directly utilizes derivative-sensitive parameters such as the similarity score ( $S$ ), gap open penalty ( $g^o$ ), and gap extend penalty ( $g^e$ ). As a result, we define our differentiable local alignment loss as  $\mathcal{L}_{swij} = \text{SW}(\text{sim}(Z^i, Z^j))$ . The (sim) function calculates the similarity between two embeddings and generates a similarity matrix  $S$ . Unlike prior differentiable SW implementations [22, 26] that depend on supervised learning, our approach integrates the differentiable SW algorithm within a self-supervised framework.

### 3.4 Final Loss

Following the approach in [9], our contrastive loss is modeled after the NT-Xent loss from SimCLR [9]. This loss function calculates absolute pairwise distances between embeddings ( $Z^1, S^1$ ) and ( $Z^2, S^2$ ), subsequently forming a Gaussian-weighted positive label distribution. This distribution is then contrasted against normalized similarity logits using the Kullback-Leibler divergence. The contrastive loss is mathematically expressed as:

$$\mathcal{L}_c = -\frac{1}{T} \sum_{i,j=1}^T \left( \frac{G(s_i^1 - s_j^2)}{\sum_{k=1}^T G(s_i^1 - s_k^2)} \log \frac{\exp(\text{sim}(z_i^1, z_j^2)/\tau)}{\sum_{k=1}^T \exp(\text{sim}(z_i^1, z_k^2)/\tau)} \right) \quad (7)$$

where  $\tau$  represents the temperature parameter,  $G$  denotes the Gaussian weighting function applied to the absolute pairwise distances, and (sim) measures the normalized similarity between embeddings. To ensure consistency between contrastive and local alignment losses, we synchronize our temporal local loss by utilizing the primary score matrix  $D_{12}$  and  $D_{21}$  from  $\mathcal{L}_{sw12}$  and  $\mathcal{L}_{sw21}$ . Following the application of the softmax function to  $D$ , the resulting matrix is referred to as  $\tilde{D}$ . The logits matrix  $L$  is then derived by performing an element-wise multiplication of  $\tilde{D}_{12}$  and  $\tilde{D}_{21}^T$ . The cross-entropy loss for this logits matrix against the Gaussian-weighted labels is calculated as follows:

$$\mathcal{L}_l = -\frac{1}{T} \sum_{i,j=1}^T \left( \frac{G(s_i^1 - s_j^2)}{\sum_{k=1}^T G(s_i^1 - s_k^2)} \log \left( \frac{\exp(L_{ij})}{\sum_{k=1}^T \exp(L_{ik})} \right) \right),$$

where  $\tilde{D}_{12,ij} = \frac{\exp(D_{12,ij}/\tau)}{\sum_{k=1}^T \exp(D_{12,ik}/\tau)}$ ,  $\tilde{D}_{21,ij} = \frac{\exp(D_{21,ij}/\tau)}{\sum_{k=1}^T \exp(D_{21,ik}/\tau)}$ ,  $L = \tilde{D}_{12} \cdot \tilde{D}_{21}^T$

Method	Class			AP@K			Progress	$\tau$
	10	50	100	K=5	K=10	K=15		
TCN [27]	80.32	81.44	83.56	76.26	76.71	77.26	82.30	83.51
TCC [10]	86.60	86.78	86.86	81.84	80.94	81.69	83.36	85.26
LAV [15]	89.77	90.35	91.77	87.48	88.36	88.40	85.20	88.75
GTA [14]	89.34	90.20	90.22	87.79	87.48	87.82	88.67	92.47
SCL [8]	<u>92.76</u>	92.80	93.05	88.75	88.51	88.97	91.26	<u>98.20</u>
LRPROP [9]	92.70	<u>94.44</u>	<u>94.36</u>	<u>92.41</u>	<u>90.33</u>	<u>90.86</u>	<u>94.09</u>	<b>99.46</b>
LAC	<b>95.87</b>	<b>95.78</b>	<b>95.16</b>	<b>92.76</b>	<b>91.07</b>	<b>91.37</b>	<b>94.24</b>	97.50

Table 1: Performance comparison of state-of-the-art methods on the Pouring Dataset [27]

Our proposed loss function, referred as Local-Alignment Contrastive (LAC), integrates contrastive and local alignment losses. It is formally defined as follows:

$$\mathcal{L} = \mathcal{L}_c + \alpha \cdot (\mathcal{L}_l + \beta \cdot (\mathcal{L}_{sw12} + \mathcal{L}_{sw21})) \quad (8)$$

where  $\alpha$  and  $\beta$  are weights used to balance the components of the loss, set to 0.01 and 1, respectively.

## 4 Experiments

**Datasets.** We evaluate the performance of the LAC model on two datasets using a variety of evaluation metrics. The Pouring dataset [27], which focuses on the action of pouring liquids, includes 70 training and 14 testing videos. The PennAction dataset [63], featuring 13 human actions, contains 1140 training and 966 testing videos. We utilize the key events and phases for the videos in both datasets as proposed by TCC [10].

**Implementation Details.** Our encoder,  $f: \mathbb{R}^{T \times C \times W \times H} \rightarrow Z$ , maps video inputs  $V$  into an embedding space  $Z$ . We use a ResNet50-v2 [16] as our backbone to extract features from the Conv4c layer with an output size of  $10 \times 10 \times 512$ . These features are then processed through adaptive max pooling, followed by two fully connected layers with ReLU activation. A subsequent linear layer projects the features into a 256-dimensional space. To enhance the model’s capacity to capture long-range dependencies, we integrate sine-cosine positional encoding and employ a two-layer Transformer encoder. To improve the model’s ability to capture long-range dependencies, we incorporate sine-cosine positional encoding and apply a two-layer Transformer encoder. The final embedding layer reduces the dimensionality to 128 for the frame-wise representations.

**Evaluation Metrics.** Following related work [2, 8, 10, 14, 15], we evaluate our model using the following metrics: (i) *Phase Classification*, which assesses the accuracy of action phase predictions by training an SVM classifier on our embeddings; (ii) *Phase Progression*, which evaluates how accurately our embeddings predict action progress using a linear regression model’s average R-squared value, based on normalized timestamp differences; (iii) *Average Precision@K (AP@K)*, which evaluates fine-grained frame retrieval accuracy by calculating the proportion of correctly matched phase labels within the  $K$  closest frames; (iv) *Kendall’s Tau*, which quantifies the temporal alignment between sequences by comparing the ratio of concordant to discordant frame pairs.



Method	Class			AP@K			Progress	$\tau$
	10	50	100	K=5	K=10	K=15		
TCN [27]	69.73	70.26	70.01	60.92	61.57	61.52	76.37	63.72
TCC [10]	86.60	86.78	86.86	81.84	80.94	81.69	83.36	85.26
LAV [15]	88.51	88.72	88.97	73.47	73.13	74.27	92.52	93.06
GTA [24]	84.21	84.68	85.28	71.72	72.17	71.52	90.51	83.35
SCL [9]	87.85	87.52	88.15	91.70	90.61	90.58	92.89	98.14
LRPROP [7]	91.90	92.96	93.25	92.46	92.2	92.03	93.03	99.09
LAC	<b>95.57</b>	<b>93.79</b>	<b>93.40</b>	<b>93.87</b>	<b>93.41</b>	<b>92.65</b>	<b>94.21</b>	94.10

Table 2: Performance comparison of state-of-the-art methods on the PennAction Dataset [63]




Figure 3: Similarity matrix (left) shows video alignment using an optimal path and the respective video aligned frame-by-frame (right).

## 4.1 Results

We apply the same four metrics to the Pouring dataset [27]. For the PennAction dataset [63], we evaluate each of the 13 action categories using the same four metrics and report the averaged results across all categories. To ensure a fair comparison (e.g., consistent data splits, identical preprocessing), we replicated the evaluations of previous approaches using the GitHub repositories<sup>1</sup> provided by the original authors. Each model was trained using its respective pre-trained backbone. An exception was made for LRPROP, as their GitHub repository is not available.

**Results on Pouring Dataset.** Table 1 presents a comparison of our method’s performance against state-of-the-art approaches on the Pouring dataset. Bold and underlined text denote the best and second-best results. Notably, it achieves a +3.11% improvement on Phase Classification using only 10% of the labels. Additionally, our model excels in AP@K and Progress metrics. However, we observe lower performance in Kendall’s Tau. We hypothesize this is due to how the SW’s algorithm encourages skipping unnecessary segments of the sequence. The introduction of gap open and extend penalties could disrupt the continuity needed for high Kendall’s Tau scores.

**Results on PennAction Dataset.** Table 2 compares the performance of our method with state-of-the-art approaches on the PennAction dataset. Bold and underlined text denote the best and second-best results. LAC consistently outperforms previous methods across most metrics, with the exception of Kendall’s Tau. Notably, the improvement in AP@K is more

<sup>1</sup>github.com/google-research/google-research/tcc, github.com/trquhuytin/LAV-CVPR21, github.com/hadjisma/VideoAlignment, github.com/mingchen/CARL\_code



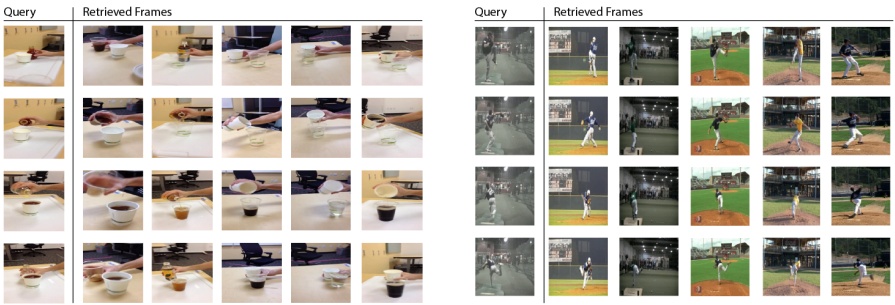


Figure 4: Fine-grained frame retrieval for Pouring (left) and PennAction (right) achieved using nearest neighbors within our embedding space.

pronounced on PennAction than on Pouring, likely due to the fewer number of frames in PennAction dataset, which may further impact alignment performance.

**Qualitative analysis of results.** Figure 3 illustrates the alignment process, with the optimal path depicted on the left and the frame-by-frame alignment on the right. This visualization demonstrates the synchronization of the two videos despite differences in their temporal progression and duration. Such alignment visualizations enable in-depth analysis of deviations or inefficiencies within specific actions in video understanding. Additionally, Figure 4 showcases our embedding-based retrieval system’s ability to accurately identify and retrieve frames corresponding to specific action sequences across different videos.

## 4.2 Ablation Study

This section presents multiple experiments on the *Pouring* dataset that analyze the different components of our framework.

Architecture	Class	Progress	$\tau$
ResNet-50 + Convolutional 3D	88.04	73.26	71.37
ResNet-50 + Transformer (w/o pretrained weights)	90.06	89.32	94.87
ResNet-50 + Transformer	<b>95.16</b>	<b>94.24</b>	<b>97.50</b>

Table 3: Different Encoder Architecture on Model Performance.

$\gamma$	Class	Progress	$\tau$	Loss $\mathcal{L}$	Class	Progress	$\tau$
0.6	93.44	92.12	94.93	$\mathcal{L}_c$	93.44	92.12	94.93
0.7	94.87	92.82	97.29	$\mathcal{L}_c + \alpha \cdot \mathcal{L}_l$	93.71	92.16	95.01
0.8	<b>95.16</b>	<b>94.24</b>	<b>97.50</b>	$\mathcal{L}_c + \alpha \cdot (\mathcal{L}_l + \beta \cdot (\mathcal{L}_{sw12} + \mathcal{L}_{sw21}))$	<b>95.16</b>	<b>94.24</b>	<b>97.50</b>
0.9	93.93	92.02	94.91				

Table 4: Different  $\gamma$  Values on Model Performance

Table 5: Different LAC Loss Formulations on Model Performance

**Network Architecture.** Table 3 demonstrates that our performance evaluation across various network architectures highlights the superiority of the ResNet-50 model with pretrained weights combined with a Transformer.

**Hyperparameter of LAC Loss.** Table 4 presents the optimal performance results obtained at the smoothing parameter  $\gamma = 0.8$ . Table 5 demonstrates that adding the local alignment loss ( $\mathcal{L}_l$ ) to the contrastive loss ( $\mathcal{L}_c$ ), the classification accuracy increases by 1.72, progression accuracy by 2.12, and the Tau score by 2.57.

## 5 Conclusion

The paper introduces a novel approach to representation learning through a Local-Alignment Contrastive (LAC) loss that integrates a differentiable local alignment loss with a contrastive loss, all within a self-supervised framework. The method employs a differentiable affine Smith-Waterman algorithm to enable temporal alignment that dynamically adjusts to variations in action sequences. This approach is distinct in its focus on capturing local temporal dependencies and enhancing the discriminative learning of video embeddings, accommodating differences in action lengths and sequences. Experimental results on the Pouring and PennAction datasets showcase the method’s superior performance over existing state-of-the-art approaches.

## References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6836–6846, 2021.
- [2] Guy Bar-Shalom, George Leifman, and Michael Elad. Weakly-supervised representation learning for video alignment and analysis. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6909–6919, 2024.
- [3] John L Barron, David J Fleet, and Steven S Beauchemin. Performance of optical flow techniques. *International journal of computer vision*, 12:43–77, 1994.
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [5] Minghao Chen, Fangyun Wei, Chong Li, and Deng Cai. Frame-wise action representations for long videos via sequence contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13801–13810, 2022.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [7] Marco Cuturi and Mathieu Blondel. Soft-dtw: a differentiable loss function for time-series. In *International conference on machine learning*, pages 894–903. PMLR, 2017.
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

- [9] Kevin Duarte, Yogesh Rawat, and Mubarak Shah. Videocapsulenet: A simplified network for action detection. *Advances in neural information processing systems*, 31, 2018.
- [10] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. Temporal cycle-consistency learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1801–1810, 2019.
- [11] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 203–213, 2020.
- [12] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019.
- [13] Rohit Girdhar, Joao Carreira, Carl Doersch, and Andrew Zisserman. Video action transformer network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 244–253, 2019.
- [14] Isma Hadji, Konstantinos G Derpanis, and Allan D Jepson. Representation learning via global temporal alignment and cycle-consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11068–11077, 2021.
- [15] Sanjay Haresh, Sateesh Kumar, Huseyin Coskun, Shahram N Syed, Andrey Konin, Zeeshan Zia, and Quoc-Huy Tran. Learning by aligning videos in time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5548–5558, 2021.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [17] Michael Isard and Andrew Blake. Condensation—conditional density propagation for visual tracking. *International journal of computer vision*, 29(1):5–28, 1998.
- [18] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014.
- [19] Myeongjun Kim, Federica Spinola, Philipp Benz, and Tae-Hoon Kim. A\*: Atrous spatial temporal action recognition for real time applications. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7014–7024, 2024.
- [20] Ivan Laptev. On space-time interest points. *International journal of computer vision*, 64:107–123, 2005.
- [21] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bsn: Boundary sensitive network for temporal action proposal generation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.

- [22] Felipe Llinares-López, Quentin Berthet, Mathieu Blondel, Olivier Teboul, and Jean-Philippe Vert. Deep embedding and alignment of protein sequences. *Nature methods*, 20(1):104–111, 2023.
- [23] James T Morton, Charlie EM Strauss, Robert Blackwell, Daniel Berenberg, Vladimir Gligorijevic, and Richard Bonneau. Protein structural alignments from sequence. *BioRxiv*, pages 2020–11, 2020.
- [24] Meinard Müller. Dynamic time warping. *Information retrieval for music and motion*, pages 69–84, 2007.
- [25] Saul B Needleman and Christian D Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of molecular biology*, 48(3):443–453, 1970.
- [26] Samantha Petti, Nicholas Bhattacharya, Roshan Rao, Justas Dauparas, Neil Thomas, Juannan Zhou, Alexander M Rush, Peter Koo, and Sergey Ovchinnikov. End-to-end learning of multiple sequence alignments with differentiable smith–waterman. *Bioinformatics*, 39(1):btac724, 2023.
- [27] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, Sergey Levine, and Google Brain. Time-contrastive networks: Self-supervised learning from video. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 1134–1141. IEEE, 2018.
- [28] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *Advances in neural information processing systems*, 27, 2014.
- [29] Temple F Smith, Michael S Waterman, et al. Identification of common molecular subsequences. *Journal of molecular biology*, 147(1):195–197, 1981.
- [30] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European conference on computer vision*, pages 20–36. Springer, 2016.
- [31] Mengmeng Wang, Jiazheng Xing, and Yong Liu. Actionclip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472*, 2021.
- [32] Chen-Lin Zhang, Jianxin Wu, and Yin Li. Actionformer: Localizing moments of actions with transformers. In *European Conference on Computer Vision*, pages 492–510. Springer, 2022.
- [33] Weiyu Zhang, Menglong Zhu, and Konstantinos G Derpanis. From actemes to action: A strongly-supervised representation for detailed action understanding. In *Proceedings of the IEEE international conference on computer vision*, pages 2248–2255, 2013.
- [34] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2914–2923, 2017.
- [35] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *Proceedings of the European conference on computer vision (ECCV)*, pages 803–818, 2018.