

Normalized Contrastive Learning for Text-Video Retrieval

Yookoon Park^{*1}, Mahmoud Azab², Bo Xiong², Seungwhan Moon²,
Florian Metze², Gourab Kundu², and Kirmani Ahmed²

¹Columbia University, New York, 10027, USA
yookoon.park@columbia.edu

²Meta, Menlo Park CA 94025, USA
{azab, bxiong, shanemoon, fmetze, gkfacebookny, kirmani}@fb.com

Abstract

Cross-modal contrastive learning has led the recent advances in multimodal retrieval with its simplicity and effectiveness. In this work, however, we reveal that cross-modal contrastive learning suffers from incorrect normalization of the sum retrieval probabilities of each text or video instance. Specifically, we show that many test instances are either over- or under-represented during retrieval, significantly hurting the retrieval performance. To address this problem, we propose Normalized Contrastive Learning (NCL) which utilizes the Sinkhorn-Knopp algorithm to compute the instance-wise biases that properly normalize the sum retrieval probabilities of each instance so that every text and video instance is fairly represented during cross-modal retrieval. Empirical study shows that NCL brings consistent and significant gains in text-video retrieval on different model architectures, with new state-of-the-art multimodal retrieval metrics on the ActivityNet, MSVD, and MSR-VTT datasets without any architecture engineering.

1 Introduction

With the advent of large-scale multimodal data and transformer-based architectures, cross-modal contrastive learning has contributed to the recent advances in multimodal representation learning (Luo et al., 2021; Radford et al., 2021). Cross-modal contrastive learning provides a simple, yet highly effective approach for learning representations from multimodal data without supervisions. In particular, CLIP (Contrastive Language-Image Pretraining) (Radford et al., 2021) learns image-text representations using vision and text transformer encoders on millions of image-text data from the Web and demonstrates that the learned vision and text encoders can perform zero-shot transfer on various vision tasks. More recently, CLIP4Clip (Luo et al.,

2021) extends the pretrained CLIP model and fine-tunes it on text-video datasets for embedding-based text-video retrieval, achieving state-of-the-art performance.

In embedding-based retrieval, the retrieval probabilities of each text or video instance are defined by their embedding similarity to their cross-modal queries. When there is one-to-one correspondence between text and video instances as in many standard benchmarks, one would expect that the retrieval probabilities of a video summed over the text queries should be normalized to 1 and vice versa, so that in overall, all text and video instances are equally represented during retrieval. However, we show that in practice, cross-modal contrastive learning suffers from significant normalization errors of the sum retrieval probabilities of each instance (Fig. 2). This suggests that many test instances are either over- or under-represented during retrieval, leading to high false positive and false negative rates (Fig. 3) and consequently harming the text-video retrieval performance.

To address this problem, we propose *Normalized Contrastive Learning (NCL)* which computes instance-wise biases using the Sinkhorn-Knopp algorithm (Cuturi, 2013) and adjusts the cross-modal embedding similarity scores so that the sum retrieval probabilities of each instance are properly normalized to 1 (Fig. 1). At test time where the test queries are not known a priori, we show that we can approximate the test query distribution by storing a subset of train queries during training in a queue. We show that this approach consistently reduces the normalization errors (Fig. 2) and thereby improves the text-video retrieval performance significantly.

We evaluate NCL on text-video retrieval on the ActivityNet, MSVD, and MSR-VTT datasets. Empirical results show that NCL consistently improves both text-to-video and video-to-text retrieval across all datasets on different base architectures, including those of CLIP4Clip (Luo et al., 2021) and SSB

* Work done during internship at Meta

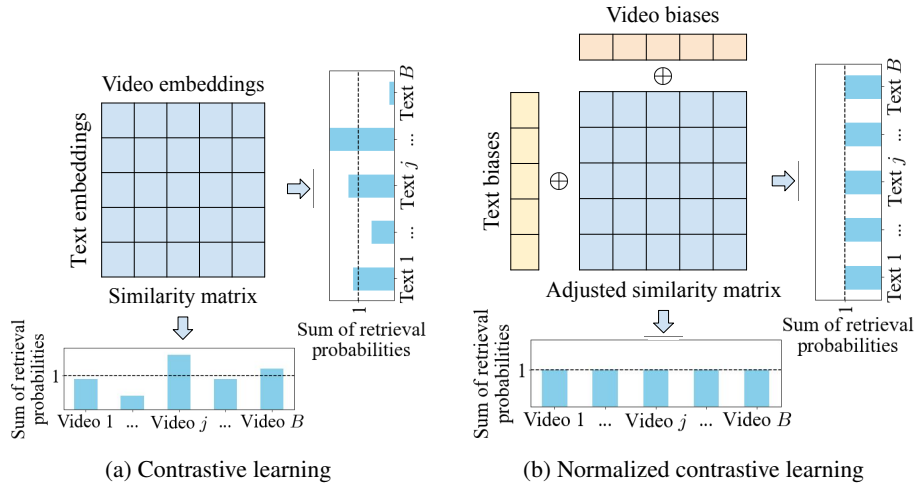


Figure 1: Overview of the proposed approach. (a) Cross-modal contrastive learning suffers from incorrect normalization of the sum retrieval probabilities of each text/video instance (b) Normalized contrastive learning computes the instance-wise biases that normalize the sum retrieval probabilities so that all instances are fairly represented

(Patrick et al., 2021) and further advances state-of-the-art text-video retrieval performance without any architecture engineering.

In summary, our contributions are:

1. Revealing that cross-modal contrastive learning suffers from incorrect normalization of the sum retrieval probabilities of each instance and how its text-video retrieval performance is impaired by this problem (Sec. 3.2).
2. Proposing the novel approach of Normalized Contrastive Learning (NCL) to address the normalization errors in cross-modal contrastive learning. NCL computes instance-wise biases using the Sinkhorn-Knopp algorithm (Cuturi, 2013) to adjust the cross-modal similarity scores so that every instance is fairly represented during retrieval (Secs. 3.3 to 3.4).
3. Establishing new state-of-the-art results on text-video retrieval tasks on multiple benchmark datasets (ActivityNet, MSVD and MSR-VTT) without any architecture engineering. Moreover, NCL brings consistent and significant gains across different base model architectures including CLIP4Clip (Luo et al., 2021) and SSB (Patrick et al., 2021) (Sec. 4).

2 Related Work

Contrastive representation learning InfoNCE (Noise Contrastive Estimation of mutual Information) has been developed for learning unsupervised representations of natural images (Wu et al., 2018;

Oord et al., 2018; Tian et al., 2019; Chen et al., 2020). It has quickly gained popularity due to its simplicity and effectiveness, leading the state-of-the-art advances in visual representation learning. Specifically, it first samples two different views of an image using random data augmentations. The views of a same image constitute a positive pair, while negative pairs consist of views of different images in a mini-batch. The encoder is then trained to minimize the cross-entropy loss to classify the positive pairs from the set of negative pairs based on the embedding similarity between the representations.

Recently, the scope of contrastive learning has been extended to multimodal data (Zhang et al., 2020; Miech et al., 2020; Radford et al., 2021; Amrani et al., 2021; Luo et al., 2021). In cross-modal contrastive learning, different modalities of the data are mapped to a shared embedding space using separate encoders for each modality, where their cross-modal similarity is defined. In particular, CLIP (Radford et al., 2021) trains its image and text encoders on millions of image-text pairs using cross-modal contrastive learning and demonstrates that the learned representations can perform zero-shot transfer on various vision tasks. Li et al. (2021, 2022) further improves the vision-language contrastive learning by incorporating additional language modeling and image-text matching losses.

Aside from InfoNCE where each example is assigned its own class label, a recent line of works (Asano et al., 2020b,a; Caron et al., 2020) assume that the data can be clustered into K latent classes.

They generate pseudo labels for the data by solving an optimal transport problem with entropy regularization, using the Sinkhorn-Knopp algorithm (Cucuri, 2013). The encoder is then trained to predict the pseudo-labels of the data. In contrast to previous work that used the Sinkhorn-Knopp algorithm for self-labeling, in this work we use the algorithm to compute the instance-wise biases that properly normalizes the sum retrieval probabilities of each text or video instance so that all instances are fairly represented during retrieval.

Embedding-based text-video retrieval maps video and text data into a shared multimodal embedding space where the similarity between a text-video pair is defined as the cosine similarity of their embeddings. Many recent works build upon pre-trained vision and text encoders (Liu et al., 2019; Miech et al., 2019, 2020; Liu et al., 2021; Amrani et al., 2021; Patrick et al., 2021; Croitoru et al., 2021; Luo et al., 2021) and finetune their models on downstream text-video datasets. For example, Liu et al. (2019) aggregates multiple expert features including objects, motion, appearance, and audio using a collaborative gating mechanism and minimizes ranking loss to align video and text embeddings for retrieval. Similarly, Gabeur et al. (2020) applies self-attention to video expert features to get video-level representations. On the other hand, Croitoru et al. (2021) trains multiple teacher models using different pretrained text encoders and distills their knowledge to a student model. Patrick et al. (2021) introduces an auxiliary task of reconstructing the caption of a video from other similar videos in a mini-batch to improve the learning of multimodal representation. Bain et al. (2021) tailors vision transformer architecture to train the model on both image and video together and applies curriculum learning by gradually increasing the number of frames the vision encoder takes. Liu et al. (2021) builds hierarchical transformers and performs hierarchical cross-modal contrastive matching at both the feature and semantic levels. In particular, CLIP4Clip (Luo et al., 2021) achieves state-of-the-art performance on text-video retrieval by loading the pretrained CLIP model (Radford et al., 2021) and finetuning it on text-video datasets using cross-modal contrastive learning. However, we show that CLIP4Clip significantly suffers from incorrect normalization of the sum retrieval probabilities of each instance and propose Normalized Contrastive Learning to address this problem.

3 Approach

3.1 Cross-modal Contrastive Learning

We start with a brief introduction of cross-modal contrastive learning. Given a batch of B ground-truth text-video pairs, video and text encoders map the input to embedding vector pairs $\{(t_1, v_1), \dots, (t_B, v_B)\}$ where each embedding lies on the unit hypersphere S^D . Cross-modal similarity between a text-video pair is defined as the cosine similarity of their embeddings. As the embeddings have the unit norm, the cosine similarity is simply equivalent to their inner product:

$$\text{sim}(t_i, v_j) = \langle t_i, v_j \rangle \quad (1)$$

Text-to-video ($t2v$) and video-to-text ($v2t$) retrieval distributions are defined based on the cross-modal similarity as

$$P_{t2v}(v_j|t_i) = \frac{\exp(\text{sim}(t_i, v_j))^{1/\gamma}}{\sum_{k=1}^B \exp(\text{sim}(t_i, v_k))^{1/\gamma}} \quad (2)$$

$$P_{v2t}(t_j|v_i) = \frac{\exp(\text{sim}(t_j, v_i))^{1/\gamma}}{\sum_{k=1}^B \exp(\text{sim}(t_k, v_i))^{1/\gamma}}, \quad (3)$$

where γ is a temperature parameter that controls the concentration of the distributions. In text-to-video retrieval, text embedding t_i serves as a query for video embedding v_i and vice versa.

Cross-modal contrastive learning (Radford et al., 2021; Luo et al., 2021) minimizes the sum of text-to-video and video-to-text cross-entropy losses:

$$\begin{aligned} \mathcal{L}_{t2v} &= -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(\text{sim}(t_i, v_i))^{1/\gamma}}{\sum_{k=1}^B \exp(\text{sim}(t_i, v_k))^{1/\gamma}} \\ \mathcal{L}_{v2t} &= -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(\text{sim}(t_i, v_i))^{1/\gamma}}{\sum_{k=1}^B \exp(\text{sim}(t_k, v_i))^{1/\gamma}} \\ \mathcal{L} &= \frac{1}{2} (\mathcal{L}_{t2v} + \mathcal{L}_{v2t}) \end{aligned} \quad (4)$$

In particular, CLIP (Radford et al., 2021) is trained on millions of image-text pairs using cross-modal contrastive learning. CLIP4Clip (Luo et al., 2021) finetunes the pretrained CLIP model for text-video retrieval, achieving state-of-the-art performance.

3.2 Normalization Errors in Retrieval

When there is one-to-one correspondence between the video and text instances as in many standard benchmarks, one would expect that the retrieval probabilities for a video v_j summed over the text queries to be 1, i.e., $\sum_i P_{t2v}(v_j|t_i) = 1, \forall j$ and

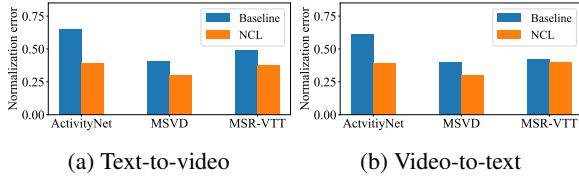


Figure 2: Normalization errors (e.g. Eq. (5)) of the sum retrieval probabilities at test time. The baseline is CLIP4Clip (Luo et al., 2021) using cross-modal contrastive learning. Normalized Contrastive Learning (NCL) consistently reduces the normalization errors across the datasets. The remaining errors of NCL comes from approximating the unknown test query distribution with a subset of train queries.

vice versa: $\sum_i P_{v2t}(t_j|v_i) = 1, \forall j$, so that all instances are equally represented during retrieval.

However, Fig. 2 shows that in practice, CLIP4Clip (Luo et al., 2021) trained using cross-modal contrastive learning suffers from significant normalization errors of the sum retrieval probabilities, where we define the text-to-video normalization error as average the absolute deviation of the sum of video retrieval probabilities from 1:

$$\mathbb{E}[|1 - \sum_{i=1}^N P_{t2v}(v_j|t_i)|], \quad (5)$$

where N is the number of test text queries. The video-to-text normalization error is defined in a symmetrical manner.

Incorrect normalization of the sum retrieval probabilities compromises the retrieval performance:

1. If $\sum_{i=1}^N P_{t2v}(v_j|t_i) < 1$, v_j is *under-represented* than the average. It will have higher chance of not being retrieved by the true query t_j (false negative)
2. If $\sum_{i=1}^N P_{t2v}(v_j|t_i) > 1$, v_j is *over-represented* than the average. It will have higher chance of being wrongly retrieved by irrelevant queries t_k for $k \neq j$ (false positive).

Figure 3 illustrates these phenomena on ActivityNet, demonstrating how normalization error is correlated with false negative and false positive rates in retrieval. Given that cross-modal contrastive learning suffers from significant normalization errors (Fig. 2), this suggests that its retrieval performance is substantially impaired by incorrect normalization of the sum retrieval probabilities.

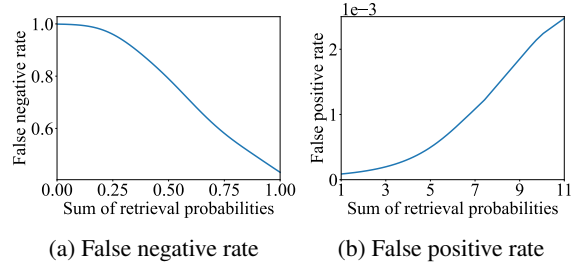


Figure 3: False negative/positive rates vs. sum of retrieval probabilities $\sum_i P_{t2v}(v_j|t_i)$ for a video in text-to-video retrieval on ActivityNet. The false negative/positive rates rapidly increase as the sum of retrieval probabilities of a video deviates from 1

3.3 Normalized Contrastive Learning

To address the problem, we propose Normalized Contrastive Learning (NCL) that normalizes the sum retrieval probabilities of each instance so that all instances are equally represented during retrieval. First, we introduce instance-wise biases and define the adjusted cross-modal similarity score as

$$\text{sim}(t_i, v_j) = a_i + b_j + \langle t_i, v_j \rangle, \quad (6)$$

where a_i is a text bias for text i and b_j is a video bias for video j . These instance-wise biases adjust the overall weights of the instances during retrieval. For example, a positive video bias b_j will increase the overall retrieval probabilities for v_j , while a negative video bias will decrease the overall retrieval probabilities for v_j . Therefore, by setting instance-wise biases a_i, b_j to appropriate values, we can properly normalize the retrieval probabilities of text and video instances.

NCL utilizes the Sinkhorn-Knopp algorithm (Cuturi, 2013) to compute the optimal values of biases for text: $\{a_1^*, \dots, a_B^*\}$, and videos: $\{b_1^*, \dots, b_B^*\}$. Specifically, given a non-negative matrix $M \in \mathbb{R}_+^{m \times n}$, the Sinkhorn-Knopp algorithm computes the normalization vectors $\alpha \in \mathbb{R}_+^m, \beta \in \mathbb{R}_+^n$ using fixed-point iterations such that

$$P = \text{diag}(\alpha) M \text{diag}(\beta), \quad (7)$$

is a valid transportation polytope, i.e.,

$$P \mathbf{1}_n = \frac{1}{n} \mathbf{1}_n, \quad (8)$$

$$P^T \mathbf{1}_m = \frac{1}{m} \mathbf{1}_m. \quad (9)$$

In other words, P represents a joint probability distribution for two random variables X, Y such that $P(X = i, Y = j) = p_{ij}$ with uniform marginal

constraints $P(X = i) = 1/m$, $P(Y = j) = 1/n$. For retrieval, this means that all instances are equally represented given the set of queries. In this work, we focus on the standard setting where there is one-to-one correspondence between the video and text captions; hence we assume uniform marginal priors for the instances. However, note that the Sinkhorn-Knopp algorithm generalizes to arbitrary prior distributions $P(X)$, $P(Y)$ (e.g. see [Asano et al. \(2020a\)](#)) and if a video is matched to multiple text captions or vice versa, we can easily modify the priors accordingly so that its sum of retrieval probabilities scales proportionally to the number of matching queries. Specifically, the RHS of Eqs. (8) to (9) will be placed with simplex vectors \mathbf{r} and \mathbf{c} that represent the marginal distributions of text and video instances. The Sinkhorn-Knopp algorithm will normalize the retrieval distributions so that the sum retrieval probabilities of each instance normalizes to its marginal weight.

Given the text-video similarity matrix $S = \{s_{ij}\} \in \mathbb{R}^{B \times B}$, $s_{ij} = \langle t_i, v_j \rangle / \gamma$, we define the non-negative matrix as $M = \exp(S)$. After computing the normalization vectors α, β for M using the Sinkhorn-Knopp algorithm, the optimal values of the text and video biases are derived as

$$a_i^* = \gamma \log \frac{\alpha_i}{\sum_k \alpha_k}, \text{ for } i = 1, \dots, B, \quad (10)$$

$$b_j^* = \gamma \log \frac{\beta_j}{\sum_k \beta_k}, \text{ for } j = 1, \dots, B. \quad (11)$$

Figure 4 gives a PyTorch-style implementation of the algorithm for computing the instance-wise biases using fixed-point iterations. We set the number of fixed point iterations to 4 and find that the residuals are sufficiently small for our experiments.

NCL uses the computed instance-wise biases to adjust the cross-modal similarity score for retrieval:

$$\text{sim}^*(t_i, v_j) = a_i^* + b_j^* + \langle t_i, v_j \rangle, \quad (12)$$

We can easily verify that the adjusted similarity score of Eq. (12) defines a properly normalized retrieval distributions such that

$$\sum_{i=1}^N P_{t2v}^*(v_j | t_i) = \sum_{i=1}^N P_{v2t}^*(t_j | v_i) = 1, \forall ij \quad (13)$$

The proof is in Appendix A.

During training, NCL computes the optimal biases for the current batch of examples and use the adjusted similarity score (Eq. (12)) for learning.

```
def sinkhorn_knopp(S, gamma, n_iters):
    M = S.exp()
    beta = 1 / M.sum(0)
    for i in range(n_iters):
        alpha = 1 / (M @ beta)
        beta = 1 / (alpha @ M)

    alpha /= alpha.sum(keepdims=True)
    beta /= beta.sum(keepdims=True)
    a = gamma * alpha.log()
    b = gamma * beta.log()
    return a, b
```

Figure 4: PyTorch-style implementation of Sinkhorn-Knopp algorithm for computing the optimal biases a, b for similarity matrix S . The computed bias vectors are used to adjust the cross-modal similarity scores (Eq. (12))

3.4 Normalization at Test Time

If the test query distribution is known a priori, we can readily construct the test similarity matrix $S = \{s_{ij}\} \in \mathbb{R}^{N \times N}$, $s_{ij} = \langle t_i, v_j \rangle / \gamma$ where N is the number of test instances. In this case, we can directly compute the optimal biases a^*, b^* that normalizes the retrieval weights of the test instances using the algorithm of Fig. 4. In this case, the normalization error will be exactly zero.

In general, however, the test queries may not be known in advance. In such cases, we propose to approximate the unseen test query distribution with a subset of train queries. For this purpose, we introduce two query queues that store the last K train queries during training, one for text and video queries each. These query queues can be easily integrated with the existing training loops with negligible computational overhead. The stored queries are only used at test time, as plug-in approximations to the unknown test query distributions. At test time, we normalize the retrieval probabilities of test instances using the subset of train queries stored in the query queues. For text-to-video retrieval, for example, we first construct a pseudo similarity matrix $\tilde{S} \in \mathbb{R}^{K \times N}$ with K train text queries in the queue and N test video instances. We then apply our normalization algorithm (Fig. 4) to compute the video biases $\tilde{b}_1, \dots, \tilde{b}_N$ for the test videos. Video-to-text retrieval is handled in a symmetrical manner. We use the computed biases $\tilde{a}_1, \dots, \tilde{a}_N$ and $\tilde{b}_1, \dots, \tilde{b}_N$ to adjust the cross-modal similarity scores (Eq. (6)) between the test queries and test instances during retrieval. Due to the approximation error from using a subset of train queries as a substitute of actual test queries, the normalization

error will be nonzero. However, Fig. 2 shows that the proposed approach still consistently reduces the normalization errors on all datasets.

The computational complexity of the test time normalization is $O(KN)$ where K is the size of the train query queue and N is the number of test instances, scaling linearly with the test set size. Given that embedding-based retrieval already has $O(N^2)$ complexity, NCL does not increase the overall test time complexity as long as $K = O(N)$. We study the effects of our approximation and the query queue size on retrieval performance in Sec. 4.2.

4 Experiments

We evaluate Normalized Contrastive Learning (NCL) on multimodal retrieval on popular text-video datasets including ActivityNet, MSVD, and MSR-VTT and report recall-at-K (R@K) metrics (higher is better), median and mean rank (lower is better). The main goals of the experiments are:

1. Compare NCL to state-of-the-art models in text-video retrieval using different model architectures (Sec. 4.1)
2. Analyze the effects of the proposed test time normalization method and the size of the query queue on NCL’s text-video retrieval performance (Sec. 4.2)

For a fair comparison, we assume that the test query distribution is not known in advance and use train query queues for test time normalization.

Datasets **ActivityNet** (Krishna et al., 2017; Fabian Caba Heilbron and Niebles, 2015) is a collection of 20,000 YouTube videos. Following (Zhang et al., 2018; Gabeur et al., 2020), we concatenate the text descriptions of a video into one paragraph and perform video-paragraph retrieval on ActivityNet. We use the ‘vall’ split for evaluation which contains 5K videos and paragraphs. The train split has 10K video-paragraphs. **MSVD** (Chen and Dolan, 2011) has 1,970 videos with each video having approximately 40 captions. The train split has 1200 videos, with 100 videos in the validation and 760 videos in the test split. In MSVD, each video in the test set has multiple captions associated with it. **MSR-VTT** (Xu et al., 2016) consists of 10,000 videos with 20 text captions per each video. We use the 9K train split with 180K captions and the 1K test split following (Yu et al.,

2018). The test split only contains one caption per video.

Architecture We use the state-of-the-art architecture of CLIP4Clip (Luo et al., 2021) with the code released by the authors. CLIP4Clip is based on the pretrained CLIP model (Radford et al., 2021) which consists of transformer-based vision and text encoders trained on large-scale image-text data. For the experiments, we adopt the mean pooling (meanP) architecture for aggregating the frame features as it was shown to deliver the most consistent performance on both text-to-video and video-to-text retrieval across different datasets (Luo et al., 2021). We do not add any trainable model components and use the CLIP4Clip architecture as-is.

Implementations details are described in Appendix B.

4.1 Comparison to State of the Art

We compare NCL with the state-of-the-art multimodal retrieval models (Croitoru et al., 2021; Patrick et al., 2021; Liu et al., 2021; Luo et al., 2021) on ActivityNet, MSVD, and MSR-VTT. NCL uses the same architecture as CLIP4Clip (Luo et al., 2021). Tables 1 to 3 summarize the results. On all datasets, NCL brings significant improvements on state-of-the-art recall metrics in both text-to-video and video-to-text retrieval across different datasets. On ActivityNet, NCL gives 13% and 10% relative gains on text-to-video and video-to-text R@1, respectively, compared to the previous state-of-the-art of CLIP4Clip (Luo et al., 2021). Especially, the gain on MSVD video-to-text retrieval is substantial with more than 23% relative boost in R@1 and the mean rank being reduced by more than half. This may be attributed to the significant imbalance between the number of video queries and text captions in the MSVD test set with 670 videos and 28K captions, making the proper normalization of the caption retrieval probabilities more crucial. In addition, NCL improves most of the recall metrics on MSR-VTT. We emphasize that these results are achieved without any architecture engineering and the additional computational overhead introduced by NCL is negligible.

In addition, Tab. 4 presents the results for NCL implemented on the Support-Set Bottleneck (SSB) (Patrick et al., 2021) architecture, which demonstrate that NCL consistently improves the retrieval performance regardless of the base architecture. The implementation details for the SSB architec-

Method	Text \rightarrow Video					Video \rightarrow Text				
	R@1 \uparrow	R@5 \uparrow	R@50 \uparrow	MdR \downarrow	MnR \downarrow	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	MdR \downarrow	MnR \downarrow
FSE (Zhang et al., 2018)	18.2	44.8	89.1	7	-	16.7	43.1	-	7	-
CE (Liu et al., 2019)	18.2	47.7	91.4	6	23.1	17.7	46.6	-	6	24.4
HSE (Zhang et al., 2018)	20.5	49.3	-	-	-	18.7	48.1	-	-	-
TT-CE+ (Croitoru et al., 2021)	23.5	57.2	96.1	4	-	23.0	56.1	-	4	-
MMT (Gabeur et al., 2020)	28.7	61.4	94.5	3.3	16.0	28.9	61.1	-	4	17.1
SSB (Patrick et al., 2021)	29.2	61.6	94.7	3	-	28.7	60.8	-	2	-
CLIP4Clip (Luo et al., 2021)	40.5	72.4	98.1	2	7.4	42.5	74.1	80.6	2	6.6
NCL (Ours)	45.9	76.8	98.3	2	6.7	46.8	76.5	86.8	2	6.2

Table 1: Multimodal retrieval results on ActivityNet evaluated on the val1 split (Fabian Caba Heilbron and Niebles, 2015). NCL brings significant 13% and 10% relative gains in state-of-the-art R@1 on text-to-video and video-to-text retrieval, respectively

Method	Text \rightarrow Video					Video \rightarrow Text				
	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	MdR \downarrow	MnR \downarrow	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	MdR \downarrow	MnR \downarrow
VSE (Kiros et al., 2014)	12.3	30.1	42.3	14	-	-	-	-	-	-
VSE++ (Faghri et al., 2018)	15.4	39.6	53.0	9.	-	-	-	-	-	-
CE (Liu et al., 2019)	19.8	49.0	63.8	6	-	-	-	-	-	-
NoiseE (Amrani et al., 2021)	20.3	49.0	63.3	6	-	-	-	-	-	-
TT-CE+ (Croitoru et al., 2021)	25.4	56.9	71.3	4	-	27.1	55.3	67.1	4	-
SSB (Patrick et al., 2021)	28.4	60.0	72.9	4	-	34.7	59.9	70.0	3	-
CLIP4Clip (Luo et al., 2021)	46.2	76.1	84.6	2	10.0	56.6	79.7	84.3	1	7.6
NCL (Ours)	47.8	77.5	85.9	2	9.9	69.6	89.9	95.4	1	3.3

Table 2: Multimodal retrieval results on MSVD. NCL consistently improves all retrieval metrics with substantial 23% relative gain in video-to-text R@1

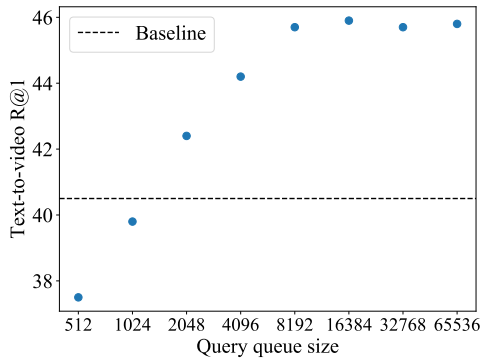


Figure 5: Query queue size vs. text-to-video R@1 on ActivityNet. The performance improves as the size of query queue increases and plateaus at about 16K

ture are described in Appendix C.

4.2 Analysis of NCL

Tab. 5 studies the effect of the approximation proposed in Sec. 3.4. Using the test queries directly

for normalization gives the best performance (last row). This is the oracle setting where the normalization errors become zero. However, in general, the test queries may not be known in advance and we approximate the unknown test query distribution using a subset of train queries stored in the query queues. We observe that this approximation reduces the test normalization errors (Fig. 2) and brings significant gains compared to the CLIP4Clip baseline (Luo et al., 2021) even without a prior knowledge of test queries. Still, there is a considerable gap when compared to the oracle.

Figure 5 presents the effect of the train query queue size on text-video retrieval performance on ActivityNet. The performance improves with growing query queue size and plateaus at about 16K queries. Based on this result, we set the query queue size to 16384 in our experiments.

We visualize example retrieval results for text-to-video retrieval on MSR-VTT in Fig. 6. It il-

Models	Text → Video					Video → Text				
	R@1↑	R@5↑	R@10↑	MdR↓	MnR↓	R@1↑	R@5↑	R@10↑	MdR↓	MnR↓
JSFusion (Yu et al., 2018)	10.2	31.2	43.2	13	-	-	-	-	-	-
HT100M (Miech et al., 2019)	12.1	35.0	48.0	12	-	16.8	41.7	55.1	8	-
JPoSE (Wray et al., 2019)	14.3	38.1	53.0	9	-	16.4	41.3	54.4	8.7	-
CE (Liu et al., 2019)	20.9	48.8	62.4	6	-	20.6	50.3	64.0	5.3	-
MMT (Gabeur et al., 2020)	26.6	57.1	69.6	4	24.0	27.0	57.5	69.7	3.7	21.3
TT-CE+ (Croitoru et al., 2021)	29.6	61.6	74.2	3	-	32.1	62.7	75.0	3	-
SSB (Patrick et al., 2021)	30.1	58.5	69.3	3	-	28.5	58.6	71.6	3	-
HiT (Liu et al., 2021)	30.7	60.9	73.2	2.6	-	32.1	62.7	75.0	3	-
CLIP4Clip(Luo et al., 2021)	43.1	70.4	80.8	2	16.2	43.1	70.5	81.2	2	12.4
NCL (Ours)	44.9	71.2	81.5	2	15.5	44.9	71.8	80.7	2	12.8

Table 3: Multimodal retrieval results on MSR-VTT evaluated on the 1K test split (Yu et al., 2018). NCL again improves most of the retrieval metrics

Dataset/Model	Text → Video					Video → Text				
	R@1↑	R@5↑	R@50↑	MdR↓	MnR↓	R@1↑	R@5↑	R@10↑	MdR↓	MnR↓
<i>ActivityNet</i>										
SSB (Patrick et al., 2021)	26.8	58.1	93.5	3	-	25.5	57.3	-	3	-
SSB + NCL	32.8	64.8	96.8	2	11.2	33.0	65.0	76.2	2	11.4
<i>MSVD</i>										
SSB (Patrick et al., 2021)	23.0	52.8	65.8	5	-	27.3	50.7	60.8	5	-
SSB + NCL	24.7	55.3	68.1	3	21.6	29.3	54.0	64.0	4	45.2
<i>MSR-VTT</i>										
SSB (Patrick et al., 2021)	27.4	56.3	67.7	3	-	26.6	55.1	67.5	3	-
SSB + NCL	28.5	57.8	70.5	3	22.3	28.5	57.3	69.4	3	22.8

Table 4: Text-video retrieval using the Support Set Bottleneck (SSB) (Patrick et al., 2021) architecture on ActivityNet, MSVD and MSR-VTT. We report the results without additional pretraining on the HowTo100M dataset. The results show that NCL brings consistent gains across all datasets regardless of the base architecture

illustrates how the sum retrieval probabilities are not well-normalized in contrastive learning (CL) with the test videos being severely over- or under-represented during retrieval. NCL consistently alleviates the normalization errors.

5 Conclusion

We’ve presented Normalized Contrastive Learning (NCL) to improve cross-modal contrastive learning for retrieval. NCL applies the Sinkhorn-Knopp algorithm to normalize the retrieval probabilities of text and video instances so that each instance is fairly represented during retrieval. When the test queries are not known a priori, NCL approximates the test query distribution with a subset of train queries stored during training. Empirical studies

show that NCL consistently reduces the normalization errors and brings significant gains in state-of-the-art text-video retrieval performance without any architecture engineering. Moreover, the gains are consistent over different model architectures. For future work, it will be worthwhile to explore if NCL can help retrieval tasks in other domains such as image and text, and whether it can improve general representation learning for downstream tasks such as classification.

Limitations

The proposed approach of Normalized Contrastive Learning (NCL) is broadly applicable to general embedding-based retrieval tasks on unimodal or multimodal domains. However, the scope of the

Normalization	Text \rightarrow Video					Video \rightarrow Text				
	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	MdR \downarrow	MnR \downarrow	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	MdR \downarrow	MnR \downarrow
None (Luo et al., 2021)	40.5	72.4	83.8	2	7.4	42.5	74.1	80.6	2	6.6
Train queries	45.9	76.8	86.5	2	6.7	46.8	76.5	86.8	2	6.2
Test queries (oracle)	54.1	80.7	89.2	1	5.7	54.2	76.7	89.4	1	5.1

Table 5: The effect of the proposed approximation (Sec. 3.4) on ActivityNet text-video retrieval. The baseline is CLIP4Clip (Luo et al., 2021) without any normalization (first row). While using the test queries for normalization gives the best performance (the oracle, last row), using a subset of train queries for normalization still brings significant gains compared to the baseline (second row) even without any knowledge of the test queries

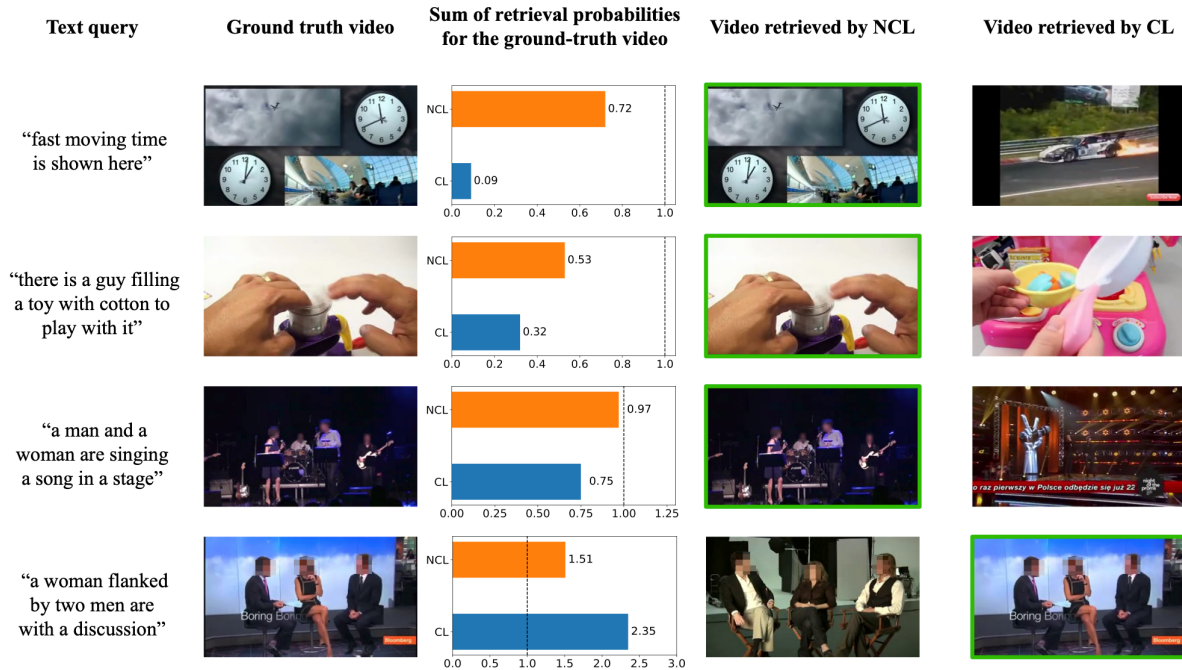


Figure 6: Example retrieval results for text-to-video retrieval on MSR-VTT comparing Normalized Contrastive Learning (NCL) to Contrastive Learning (CL). NCL consistently reduces the normalization errors in all examples. The first three rows show the examples where the ground-truth videos are under-represented and CL fails to correctly retrieve the videos, while NCL correctly retrieves the ground truth videos

empirical studies in this work was limited to the video-text domain only. Other limitations include that NCL may not generalize well if the test query distribution is significantly different from the training query distribution or if the test query distribution changes dynamically over time. Such discrepancies may be alleviated by storing the previous test queries on the fly in the query queue to dynamically adapt to the unseen test query distribution. We leave related investigations to future work.

Ethics Statement

The proposed approach retrieves relevant videos given a text query and vice versa. The encoders for video and text were pretrained on large collec-

tion of image-text data and finetuned on video-text datasets and may reflect biases in those data, including those with potentially negative societal impacts. In addition, there is a risk that the model will be used for tasks of societal sensitivity, such as surveillance. These concerns call for caution and consideration when deploying or building upon the proposed model.

References

Elad Amrani, Rami Ben-Ari, Daniel Rotman, and Alex Bronstein. 2021. Noise estimation using density estimation for self-supervised multimodal learning. In *AAAI*.

- Yuki M Asano, Mandela Patrick, Christian Rupprecht, and Andrea Vedaldi. 2020a. Labelling unlabelled videos from scratch with multi-modal self-supervision. In *NeurIPS*.
- Yuki M. Asano, Christian Rupprecht, and Andrea Vedaldi. 2020b. Self-labelling via simultaneous clustering and representation learning. In *International Conference on Learning Representations (ICLR)*.
- Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. 2021. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*.
- Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. 2020. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*.
- David Chen and William B Dolan. 2011. Collecting highly parallel data for paraphrase evaluation. In *ACL-HLT*.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *ICML*.
- Ioana Croitoru, Simion-Vlad Bogolin, Yang Liu, Samuel Albanie, Marius Leordeanu, Hailin Jin, and Andrew Zisserman. 2021. Teachtext: Cross-modal generalized distillation for text-video retrieval. *arXiv preprint arXiv:2104.08271*.
- Marco Cuturi. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26:2292–2300.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.
- Bernard Ghanem Fabian Caba Heilbron, Victor Escorcia and Juan Carlos Niebles. 2015. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*.
- Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. 2018. Vse++: Improving visual-semantic embeddings with hard negatives. In *BMVC*.
- Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. 2020. Multi-modal transformer for video retrieval. In *ECCV*.
- Deepti Ghadiyaram, Du Tran, and Dhruv Mahajan. 2019. Large-scale weakly-supervised pre-training for video action recognition. In *CVPR*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. 2014. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*.
- Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. Dense-captioning events in videos. In *Proceedings of the IEEE international conference on computer vision*, pages 706–715.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint arXiv:2201.12086*.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation.
- Song Liu, Haoqi Fan, Shengsheng Qian, Yiru Chen, Wenkui Ding, and Zhongyuan Wang. 2021. Hit: Hierarchical transformer with momentum contrast for video-text retrieval. *arXiv preprint arXiv:2103.15049*.
- Yang Liu, Samuel Albanie, Arsha Nagrani, and Andrew Zisserman. 2019. Use what you have: Video retrieval using representations from collaborative experts. *arXiv preprint arXiv:1907.13487*.
- Ilya Loshchilov and Frank Hutter. 2017. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*.
- Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. 2021. Clip4clip: An empirical study of clip for end to end video clip retrieval. *arXiv preprint arXiv:2104.08860*.
- Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. 2020. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*.
- Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. 2019. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *ICCV*.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Mandela Patrick, Po-Yao Huang, Yuki Asano, Florian Metze, Alexander Hauptmann, João Henriques, and Andrea Vedaldi. 2021. Support-set bottlenecks for video-text representation learning. In *ICLR*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*.

Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2019. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*.

Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. 2018. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*.

Tongzhou Wang and Phillip Isola. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *ICML*. PMLR.

Michael Wray, Diane Larlus, Gabriela Csurka, and Dima Damen. 2019. Fine-grained action retrieval through multiple parts-of-speech embeddings. In *ICCV*.

Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*.

Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, and Florian Metze Luke Zettlemoyer Christoph Feichtenhofer. 2021. Videoclip: Contrastive pre-training for zero-shot video-text understanding. In *EMNLP*.

Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*.

Youngjae Yu, Jongseok Kim, and Gunhee Kim. 2018. A joint sequence fusion model for video question answering and retrieval. In *ECCV*.

Bowen Zhang, Hexiang Hu, and Fei Sha. 2018. Cross-modal and hierarchical modeling of video and text. In *ECCV*.

Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D Manning, and Curtis P Langlotz. 2020. Contrastive learning of medical visual representations from paired images and text. *arXiv preprint arXiv:2010.00747*.

A Proof for Eq. (13)

We prove that the adjusted cross-modal similarity score (Eq. (12)) using the instance-wise biases computed by the algorithm of Fig. 4 properly nor-

malizes the retrieval distributions:

$$\sum_{i=1}^B P_{t_2v}^*(v_j|t_i) \quad (14)$$

$$= \sum_{i=1}^B \frac{\exp(\text{sim}^*(t_i, v_j))^{1/\gamma}}{\sum_{k=1}^B \exp(\text{sim}^*(t_i, v_k))^{1/\gamma}} \quad (15)$$

$$= \sum_{i=1}^B \frac{\exp(a_i^* + b_j^* + \langle t_i, v_j \rangle)^{1/\gamma}}{\sum_{k=1}^B \exp(a_i^* + b_k^* + \langle t_i, v_k \rangle)^{1/\gamma}} \quad (16)$$

$$= \sum_{i=1}^B \frac{\alpha_i \beta_j M_{ij}}{\sum_{k=1}^B \alpha_i \beta_k M_{ik}} \quad (17)$$

$$= \sum_{i=1}^B \frac{p_{ij}}{\sum_k p_{ik}} = B \sum_{i=1}^B p_{ij} = B \frac{1}{B} = 1, \quad (18)$$

where the last line follows from the normalization properties (Eqs. (8) to (9)).

B Implementation Details for CLIP4Clip

We mostly follow the experimental settings of CLIP4Clip (Luo et al., 2021). We initialize the text and vision encoders with CLIP (ViT-B/32) (Radford et al., 2021) and finetune the model using Adam optimizer (Kingma and Ba, 2015) and a cosine learning rate schedule without restart (Loshchilov and Hutter, 2017) with initial learning of 1e-7. The maximum caption token length is set to 32 for MSVD and MSR-VTT, and 64 for ActivityNet. The number of frames sampled is 12 for MSVD and MSR-VTT, and 64 for ActivityNet whose videos are generally much longer in length. The only hyperparameters we tweak are: (1) the number of training epochs for each dataset and (2) the size of the query queue. These hyperparameters are selected based on their validation performance. We train the model for 3, 5, and 10 epochs for MSVD, MSR-VTT, and ActivityNet, respectively. We apply linear learning rate warm-up over the first 10% of the training iterations for all datasets except for MSVD where we find that using warm-up slightly degrades the performance. For test time normalization, NCL introduces two query queues to store a subset of train queries during training, one for text and video queries each. The query queue size is set to 16384.

C Implementation Details for SSB

We describe the experimental details for the experiments using the Support-Set Bottleneck (SSB) architecture. The code for SSB was provided by the authors.

Architecture. SSB uses the pretrained T5-base model (Raffel et al., 2020) for its text encoder. For vision, it first extracts motion and appearance features using the 34-layer R(2+1)-D model (Tran et al., 2018) pretrained on IG65M (Ghadiyaram et al., 2019) and ResNet152 (He et al., 2016) pretrained on ImageNet (Deng et al., 2009), respectively. It then concatenates the motion and appearance features for its visual input. SSB applies CNN and RNN networks on its text and visual features, respectively, followed by transformer pooling layers. For more details, we refer to the original paper (Patrick et al., 2021).

Implementation details. In our experiments, we use the SSB model (Patrick et al., 2021) without additional pretraining on HowTo100M (Miech et al., 2019). SSB employs max-margin triplet ranking loss with hard negative mining (Faghri et al., 2018) and we replace the margin loss with the NCL loss (Eq. 31 to 33 in the paper). The temperature parameter for NCL is set to 0.07 following (Chen et al., 2020). NCL uses query queues of size 16384 to store the queries during training. The scale of the NCL loss is multiplied by 15 in order to approximately match the scale of the original margin loss. We apply dropout with ratio 0.1 for MSVD and MSR-VTT, and 0.0 for ActivityNet. For the rest of the hyperparameters, we follow the setting used in (Patrick et al., 2021).

D Multimodal Embedding Space of Cross-modal Contrastive Learning

Many recent works (Radford et al., 2021; Luo et al., 2021; Xu et al., 2021; Miech et al., 2020; Bain et al., 2021) have demonstrated the promise of cross-modal contrastive learning for multimodal data. However, its behavior and properties in multimodal environments have not been well-understood until now. In this section, we study the multimodal embedding space of the CLIP4Clip (Luo et al., 2021) model trained using cross-modal contrastive learning.

Figure 7 visualizes the video-text embedding space of the CLIP4CLIP model at initialization on the MSR-VTT dataset. The video and text embeddings do not overlap with each other, being highly clustered to their modalities in the embedding space. This is surprising, as the success of contrastive learning in the unimodal setting of natural images has previously been attributed to the alignment and uniformity of the embeddings (Wang

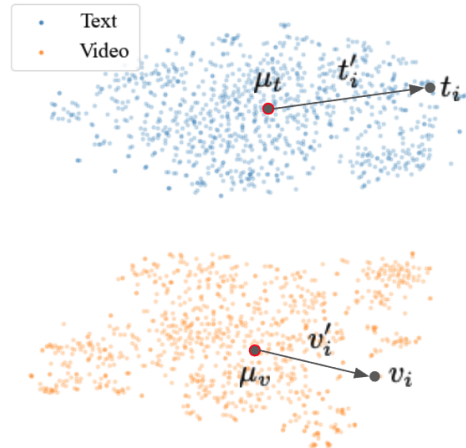


Figure 7: TSNE visualization of the multimodal embedding space of CLIP4Clip (Luo et al., 2021) on MSR-VTT. The figure perceptually illustrates the modal mean decomposition (Eqs. (19) to (20)). μ_v, μ_t denote the modal means of video and text embeddings and v'_i, t'_i represents the displacements from their respective modal means. The text and video embeddings are highly clustered to their modalities and do not overlap with each other.

and Isola, 2020). The figure shows that the video-text embeddings are neither well-aligned with each other nor evenly distributed on the unit hypersphere. This can also be confirmed from the average embedding similarities between the modalities in Fig. 8, which shows that within-modal embedding similarities are still high even after finetuning on MSR-VTT. In addition, the relatively low average similarity between text and video suggests that the embeddings of different modalities do not overlap as depicted in Fig. 7. This problem may be due to the distribution shift caused by finetuning the pretrained CLIP model on MSR-VTT. While finetuning the model longer might alleviate this issue, we find that longer finetuning harms the retrieval performance due to overfitting.

To analyze the implications of this phenomenon on cross-modal contrastive learning, we decompose the embeddings as

$$t_i = \mu_t + t'_i \quad (19)$$

$$v_i = \mu_v + v'_i, \quad (20)$$

where μ_v, μ_t denote the modal means of video and

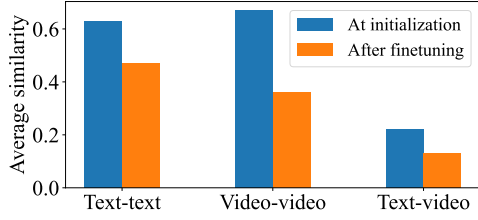


Figure 8: Average embedding similarities between modalities on MSR-VTT at initialization and after finetuning. Even after finetuning, with-in modal similarities are high, showing that the embeddings are clustered to their modalities. On the other hand, relatively low text-video similarity suggests that the embeddings of different modalities do not overlap with each other as depicted in Fig. 7

text embeddings:

$$\mu_t = \frac{1}{B} \sum_{i=1}^B t_i, \quad (21)$$

$$\mu_v = \frac{1}{B} \sum_{i=1}^B v_i, \quad (22)$$

and v'_i, t'_i represent the displacements from their respective modal means such that $\mathbb{E}[t'_i] = \mathbb{E}[v'_i] = 0$. Figure 7 perceptually illustrates the decomposition.

Using this modal mean decomposition, the cross-modal similarity between a text-video pair can be written as

$$\langle t_i, v_j \rangle = \langle \mu_t, \mu_v \rangle + \langle \mu_v, t'_i \rangle + \langle \mu_t, v'_j \rangle + \langle t'_i, v'_j \rangle. \quad (23)$$

The first term is constant irrespective of i or j and can be discarded without loss of generality. The second and third terms are implicit video- and text-wise *biases* that depend either on i or j . The last term is the text-video alignment between the displacements from their respective modal means. This decomposition shows that when there are clustering structures within different modalities, there are implicit instance-wise biases that control the overall retrieval weights of each instance.

For example, we can rewrite the text-to-video

retrieval distribution using the decomposition:

$$P_{t2v}(v_j|t_i) \quad (24)$$

$$= \frac{\exp(\langle t_i, v_j \rangle)^{1/\gamma}}{\sum_{k=1}^B \exp(\langle t_i, v_k \rangle)^{1/\gamma}} \quad (25)$$

$$= \frac{\exp(\langle \mu_t, v'_j \rangle + \langle t'_i, v'_j \rangle)^{1/\gamma}}{\sum_{k=1}^B \exp(\langle \mu_t, v'_k \rangle + \langle t'_i, v'_k \rangle)^{1/\gamma}} \quad (26)$$

$$= \frac{\beta_j \exp(\langle t'_i, v'_j \rangle)^{1/\gamma}}{\sum_{k=1}^B \beta_k \exp(\langle t'_i, v'_k \rangle)^{1/\gamma}}, \quad (27)$$

where $\beta_j = \exp(\langle \mu_t, v'_j \rangle)^{1/\gamma}$ and other common terms have been canceled out. Equation (27) can be seen as a weighted Softmax and reveals how the video bias $\langle \mu_t, v'_j \rangle$ controls the weights of the videos in text-to-video retrieval.

The role of these biases in cross-modal contrastive learning can be better understood through gradient analysis of the cross entropy loss. Denote the video bias as $b_j = \langle \mu_t, v'_j \rangle$. The gradient with respect to the bias is

$$\frac{\partial \mathcal{L}}{\partial b_j} = -\frac{1}{2B} \left(1 - \sum_{i=1}^B P_{t2v}(v_j|t_i) \right), \quad (28)$$

where $\sum_{i=1}^B P_{t2v}(v_j|t_i)$ is the retrieval probabilities of video j summed over the set of text queries. Equation (28) shows that the video bias will be learned to minimize the normalization error of the sum retrieval probability. Therefore cross-modal contrastive learning actually learns to normalize the sum retrieval probabilities for each instance during training. However, we empirically find that cross-modal contrastive learning suffers from significant normalization errors of retrieval probabilities and that the explicit normalization proposed in the paper substantially improves its retrieval performance. For the reason why cross-modal contrastive learning suffers from such significant normalization errors in practice, we leave it to future work.