

KDMCSE: Knowledge Distillation Multimodal Sentence Embeddings with Adaptive Angular margin Contrastive Learning

Cong-Duy Nguyen¹, Thong Nguyen², Xiaobao Wu¹, Luu Anh Tuan¹

¹Nanyang Technological University, Singapore

²National University of Singapore, Singapore

Abstract

Previous work on multimodal sentence embedding has proposed multimodal contrastive learning and achieved promising results. However, by taking the rest of the batch as negative samples without reviewing when forming contrastive pairs, those studies encountered many suspicious and noisy negative examples, significantly affecting the methods' overall performance. In this work, we propose KDMCSE (Knowledge Distillation Multimodal contrastive learning of Sentence Embeddings), a novel approach that enhances the discrimination and generalizability of multimodal representation and inherits the knowledge from the teacher model to learn the difference between positive and negative instances and via that, can detect noisy and wrong negative samples effectively before they are calculated in the contrastive objective. Furthermore, to overcome the limitation of modeling the variation within negative pairs, we introduce a new contrastive objective, AdapACSE (Adaptive Angular Margin Supervised Contrastive Learning for Multimodal sentence embeddings), that enhances the discriminative representation by strengthening the margin within the angular space while capturing varying semantics within the negative. Experimental results on widely used Semantic Textual Similarity (STS) benchmarks demonstrate the effectiveness of our approach. The source code is available at <https://github.com/duyngtr16061999/KDMCSE>.

1 Introduction

Learning sentence embeddings, which involves translating sentences into consistent-length vectors that capture their semantic connections, remains a pivotal task in the NLP field. While pre-trained language models like BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) have witnessed immense success, studies indicate that the default sentence embeddings from these PLMs,

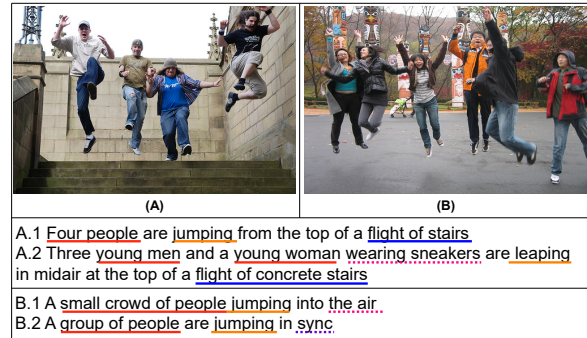


Figure 1: Example image-caption pairs in Flickr. Solid lines of the same color talk about the same instance, and a dot line means the additional information that does not occur in the other caption.

without specific fine-tuning, may not perform as well as simply averaging GloVe vectors (Pennington et al., 2014) in capturing semantic similarity (Reimers and Gurevych, 2019). Consequently, recent works (Li et al., 2020; Zhang et al., 2020; Su et al., 2021) are geared towards refining sentence embeddings from PLMs without supervision.

While text-centric models have made significant strides, their depth of understanding of sentence semantics remains a point of contention. True semantic understanding often stems from associations in the real world rather than mere textual statistics (Wang et al., 2019; Bender and Koller, 2020; Bisk et al., 2020; Wei et al., 2022). Numerous recent studies have explored the enhancement of language representations through visual information (Tan and Bansal, 2020b; Tang et al., 2021; Nguyen et al., 2023b). The recent multimodal contrast in sentence embedding learning (Zhang et al., 2022a) offers an encouraging direction. This method augments SimCSE with a multimodal contrastive objective, which seeks congruence between sentences and their associated images within a unified space. However, samples drawn randomly from grounded datasets during training can present semantic differences, subtle similarities, or over-



T	Two chefs in white are working in the kitchen, while a woman sits at a nearby counter	CLIP score
a	There is a woman sitting at a table by the kitchen	0.92
b	People are in a restaurant eating	0.83
c	A woman is preparing food at the store	0.49
d	The small dog stands inside of a suitcase on a bed	0.03

Figure 2: Example image-caption pairs in Flickr. The green caption is a true annotation of the image while 4 red captions is randomly picked from the dataset. The scores on the right are the cosine similarity between its caption representation and image representation extracted from CLIP model.

lapping attributes. For example, in Figure 1, two captions B.1 and B.2 describe the second image B, but also correctly represent the first image A. Due to their inherent noise, such samples can hinder the learning process. Mitigating this issue requires preemptively filtering out these semantically similar samples before loss calculation, ensuring a noise-reduced training environment. In addition, comparing two captions, A.1 and A.2, for the left image reveals that the latter offers a more comprehensive description, detailing aspects such as the individual’s gender, attire, sneakers, and even the material of the stairs (concrete). Thus, using only visual images cannot capture the depth of language representations.

Furthermore, MCSE did not consider the distinctions between negative sample pairs while focusing on constructing pairs of positive and negative representations. As shown in Figure 2, although none of the captions perfectly describes the image, their levels of discrimination vary between pairs. On the contrary, the caption *c* registers a score of 0.49, hinting at its partial relevance to the image. Caption *d*, with a very weak score of 0.03, is completely unrelated. Consequently, within the feature space, both captions *c* and *d* should be substantially distanced from the accurate depiction. Furthermore, the gap between the true caption and caption *d* should considerably exceed that between the true caption and the caption *c*.

To address the aforementioned limitation of noisy negative sampling and complexity of linguis-

tic representation, we present KDMCSE: a Knowledge Distillation Multimodal Contrastive Learning framework for Sentence Embeddings that utilizes CLIP (Radford et al., 2021) as the teacher model for both image and text modalities. Our approach harnesses images for multimodal objectives and textual representations derived from the CLIP text encoder. Relying solely on images falls short of capturing the nuances in sentence representation. Emulating the representation of the CLIP’s textual encoder - especially when paired with image matching - can augment the richness of linguistic representations, transcending a mere alignment with visual features.

Moreover, in response to the identified challenge of diversity among negative pairs, we introduce AdapACSE: Adaptive Angular Margin Contrastive Learning. This method bolsters discriminative representation by amplifying the margin within the angular space, specifically accounting for the various semantics present in negative samples. Leveraging the teacher model, CLIP, we produce soft labels that signify the similarity between samples. Instead of merely utilizing feature representations as positive or negative samples during the multimodal objective, we employ CLIP’s soft labels to intensely penalize samples exhibiting low similarity and provide leniency to pairs exhibiting certain resemblances.

Our contributions can be summarized as follows.

- We propose a novel multimodal framework for sentence embedding KDMCSE: Knowledge Distillation Multimodal Contrastive learning for Sentence Embedding. The knowledge of the teacher model is transferred to the student model through multimodal contrastive learning, spanning both text and visual modalities.
- We introduce self-supervised contrastive learning: AdapACSE that improves the discriminative representation of samples with varying degrees of similarity.
- We evaluate our approach using standard semantic textual similarity (STS) benchmarks and SentEval transfer tasks. Our method outperforms the performance of earlier state-of-the-art approaches.

2 Related Work

Sentence Representation Learning In previous research, sentence representations were typ-

ically learned by extending the principles of word2vec (Mikolov et al., 2013), predicting neighboring sentences (Kiros et al., 2015a; Hill et al., 2016; Logeswaran and Lee, 2018) or aggregating n-gram embeddings (Pagliardini et al., 2018), the pre-trained transformer-based model BERT (Li et al., 2020; Reimers and Gurevych, 2019) and its enhancements BERT-Flow (Li et al., 2020) and BERT-Whitening (Su et al., 2021). More recently, various studies have embraced the contrastive learning framework for learning sentence representation. These studies have proposed different methods to form contrastive pairs, either through various data transformation techniques (Yan et al., 2021; Zhang et al., 2020; Giorgi et al., 2021) or through encoders with distinct structures or parameters (Carlsson et al., 2021; Kim et al., 2021). A notable example is SimCSE (Gao et al., 2021), which employs dropout as a data augmentation strategy. Many studies in the field of sentence contrastive learning have made a significant contribution by focusing on the mining of challenging negative samples (Wang et al., 2022b; Zhou et al., 2022; Zhang et al., 2022b; Wang et al., 2022a; Wei et al., 2023; He et al., 2023).

Deep Metric Learning Objectives Contrastive learning (Chopra et al., 2005) has gained traction across multiple domains (Nguyen and Luu, 2021; Nguyen et al., 2022, 2024, 2023a; Wu et al., 2020, 2022, 2023a,b, 2024; Wei et al., 2024). Renowned training objectives such as N-Pair Loss (Sohn, 2016), Triplet Margin Loss (Balntas et al., 2016), and ArcCon (Zhang et al., 2022c) are deeply rooted in metric learning principles. In supervised tasks, objectives that leverage softmax have shown efficacy, especially when they integrate class centers and impose penalties on the distances between deep features and their corresponding centers. Among these, Center loss (Wen et al., 2016), SphereFace (Liu et al., 2017), CosFace (Wang et al., 2018), and ArcFace (Deng et al., 2018) stand out and are frequently used in computer vision and natural language processing applications. Nonetheless, these loss functions primarily aim at classification tasks, rendering them unsuitable for regression labels. ArcCSE (Zhang et al., 2022c) presents training objectives using additional margin tailored to enhance the discriminative prowess in pairwise relations and capture the entailment dynamics within triplet sentence structures.

Visually Grounded Representation Learning

Numerous studies have underscored the advantages of integrating NLP models with visual insights to improve textual representation learning (Lazaridou et al. (2015); Kiela et al. (2018)). Based on the Skip-Thought model of Kiros et al. (2015b), Borde et al. (2019) created a unified space that harmoniously accommodates both visual and textual dimensions. Recent pioneering efforts by Tan and Bansal (2020a), Tang et al. (2021) and (Nguyen et al., 2023b) have laid the foundation for vast language models using multimodal guidance, aiming to improve language understanding. Similarly, MCSE by Zhang et al. (2022a) suggested sentence embedding learning using a multimodal contrastive objective to align sentences with their respective images coherently. Distinguishing our work from Zhang et al. (2022a), we present an innovative angular margin contrastive learning framework, building upon the cutting-edge multimodal contrastive learning proposed by Zhang et al. (2022a). By weaving in this new contrastive technique, our goal is to push the envelope in STS by leveraging multimodal semantic data.

3 Method

3.1 Background: Unsupervised SimCSE and Multimodal Contrastive Learning MCSE

The idea of unsupervised SimCSE is to use dropout noise as a simple, yet effective, data augmentation strategy. Given a collection of sentences $\{x_i\}_{i=1}^m$, we construct a positive pair for each input x_i by encoding it twice using different dropout masks: $\mathbf{h}_i^z = g_\phi(f_\theta(x_i, z))$ and $\mathbf{h}_i^{z'} = g_\phi(f_\theta(x_i, z'))$, where z and z' denote different dropout masks¹, $f_\theta(\cdot)$ is a pre-trained language encoder such as BERT, and $g_\phi(\cdot)$ is a projection head² on top of the [CLS] token. The training objective is:

$$\ell_i^S = -\log \frac{e^{\text{sim}(\mathbf{h}_i^{z_i}, \mathbf{h}_i^{z'_i})/\tau}}{\sum_{j=1}^N e^{\text{sim}(\mathbf{h}_i^{z_i}, \mathbf{h}_j^{z'_j})/\tau}}, \quad (1)$$

where N is the size of the mini-batch, τ is a temperature parameter, and $\text{sim}(\mathbf{h}_1, \mathbf{h}_2)$ is the cosine similarity $\frac{\mathbf{h}_1^T \mathbf{h}_2}{\|\mathbf{h}_1\| \cdot \|\mathbf{h}_2\|}$. After training, the [CLS] token outputs of the language encoder are taken as the sentence embeddings.

¹The standard dropout masks in Transformers are used.

²There is a MLP pooler layer over [CLS] in BERT’s implementation. Gao et al. (2021) use it with reinitialization.

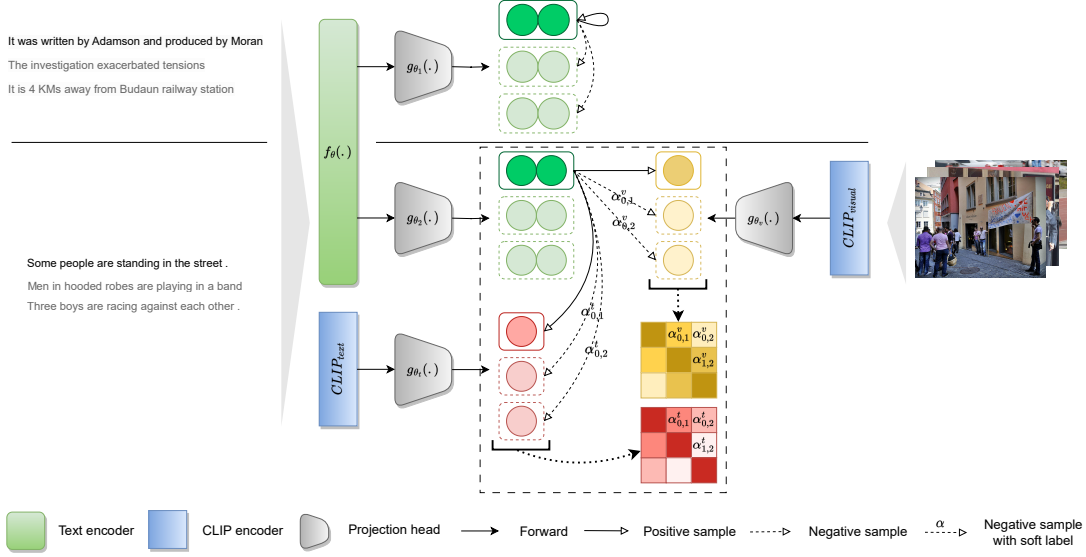


Figure 3: The overall architecture of KDMCSE. The upper part is the original SimCSE, the below part is the multimodal contrastive learning approach with knowledge distillation from CLIP model.

To take advantage of visual and textual information, MCSE (Zhang et al., 2022a) adopt SimCSE as the textual baseline and extend it with a multimodal contrastive learning objective:

$$\mathbf{s}_i^z = g_{\phi_1}(f_\theta(x_i, z)), \mathbf{v}_i = g_{\phi_2}(f^v(y_i)), \quad (2)$$

where $f^v(\cdot)$ is a pre-trained image encoder. $g_{\phi_1}(\cdot)$ and $g_{\phi_2}(\cdot)$ are two projection heads for text and image modalities. They define the multimodal contrastive learning objective as:

$$\ell_i^M = - \sum_{z \in \{z_i, z_i^z\}} \log \frac{e^{\text{sim}(\mathbf{s}_i^z, \mathbf{v}_i) / \tau'}}{\sum_{j=1}^N e^{\text{sim}(\mathbf{s}_i^z, \mathbf{v}_j) / \tau'}}, \quad (3)$$

3.2 Knowledge Distillation Multimodal Contrastive learning for Sentence Embedding

In addition to the multimodal objective, we present a Knowledge Distillation framework designed to harness both visual and textual insights from the vision-language model CLIP. An overview of our KDMCSE model can be found in Figure 3. Given a set of sentence-image pairs represented as $D = \{x_i, y_i\}_{i=1}^m$, we initially project the sentence x_i and the image y_i into a shared or grounded space.

$$\mathbf{s}_i^z = g_{\phi_g}(f_\theta(x_i, z)), \quad (4)$$

$$\mathbf{t}_i = g_{\phi_t}(\hat{\mathbf{t}}_i), \hat{\mathbf{t}}_i = CLIP_{text}(x_i), \quad (5)$$

$$\mathbf{v}_i = g_{\phi_v}(\hat{\mathbf{v}}_i), \hat{\mathbf{v}}_i = CLIP_{visual}(y_i), \quad (6)$$

where $CLIP_{text}$, $CLIP_{visual}$ are a pre-trained text and image CLIP encoder, $\hat{\mathbf{v}}_i$ and $\hat{\mathbf{t}}_i$ are the visual

and text representation extracted from CLIP, $g_{\phi_g}(\cdot)$ is the projection heads for the language student model to project the sentence representation into grounded space, $g_{\phi_t}(\cdot)$ and $g_{\phi_v}(\cdot)$ are distinct projection heads of the teacher model for the text and image modality, respectively. During training, we freeze the pre-trained CLIP model.

Threshold Filtering During training, negative samples were randomly taken from data in Eq.3. We not only transfer CLIP’s knowledge of the text and visual representation of each sample, but also leverage the similarity within samples in a batch. We calculate the cosine similarity:

$$\alpha_{i,j}^{m,n} = \frac{m_i \cdot n_j}{\|m_i\|_2 \|n_j\|_2}, m, n \in \{\hat{\mathbf{v}}, \hat{\mathbf{t}}\} \quad (7)$$

where m, n are the defined modality (text or visual). Thus, as shown in Figure 4, we can generate four similarity mappings: text-text, visual-visual, text-visual, and visual-text. As we strive to align the sentence representation generated by the language encoder (text) with the information gleaned from the teacher model encompassing both sentence (text) and image (visual) data, we exclusively employ ‘text-text’ and ‘text-visual’ pairings for our soft labels. We define a threshold filter function to remove noisy negative samples:

$$\varphi_{i,j}^{m,n} = \begin{cases} 0 & \text{if } \alpha_{i,j}^{m,n} < \text{threshold} \\ 1 & \text{else} \end{cases} \quad (8)$$

where *threshold* is a fixed hyperparameter during

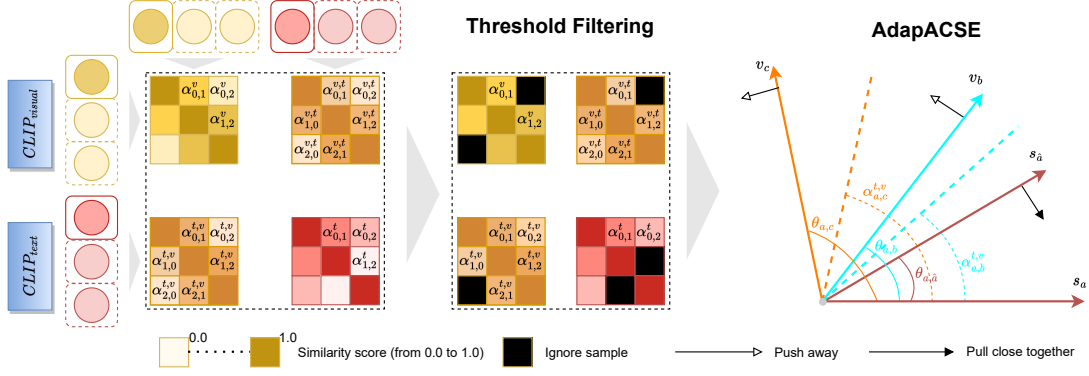


Figure 4: The overall framework of knowledge distillation with Adaptive Angular margin contrastive learning. The pipeline first calculates the soft-label similarity scores between text and visual representation, then we apply threshold filtering to remove the noisy negative pairs, and finally, we transfer the soft-label matrices into our proposed AdapACSE to flexibly find the margin. $s_{\hat{a}}$ is the positive sample for s_a , v_b and v_c are its negative counterparts. In particular, the difference in pairs $s_a - v_c$ is more pronounced than in $s_a - v_b$. As a result, the margin (depicted as a dashed line) for c (in orange) is greater than for b (in cyan).

training. The proposed contrastive objective with Threshold Filtering is defined as follows:

$$\ell_i^m = - \sum_{z \in \{z_i, z'_i\}} \log \frac{e^{\text{sim}(s_i^z, \mathbf{m}_i)/\tau'}}{\sum_{j=1}^N \varphi_{i,j}^{t,m} e^{\text{sim}(s_i^z, \mathbf{m}_j)/\tau'}}, \quad (9)$$

where m is the modality (text or visual), $m \in \{t, v\}$, \mathbf{m}_i is the projected representation of CLIP model, $\mathbf{m}_i \in \{v_i, t_i\}$.

3.3 Adaptive Angular margin Contrastive learning

Previous studies (Zhang et al., 2022c) introduced the ArcCSE objective. This was designed to increase the discriminative capacity between pairs and encapsulate the entailment relationship among triplet sentences. ArcCSE adds an additive angular margin m_c between positive pairs h_i^z and $h_i^{z'}$. Compared to Eq.1, it further pushed h_i towards the area where $\theta_{i,i'}$ becomes smaller and $\theta_{j,i}$ becomes larger, increasing the compactness of sentence representations with the same semantics and enlarging the discrepancy of different semantic representations. ArcCSE is defined as:

$$\ell_i^{\text{ArcCSE}} = -\log \frac{e^{\phi(\theta_{i,i'} + m_c)/\tau}}{e^{\phi(\theta_{i,i'} + m)/\tau} + \sum_{j \neq i} e^{\phi(\theta_{j,i})/\tau}} \quad (10)$$

where ϕ is the \cos function, angular $\theta_{i,j}$ is denoted as follows:

$$\theta_{i,j} = \arccos \left(\frac{h_i^T h_j}{\|h_i\| * \|h_j\|} \right) \quad (11)$$

where h_i is the vector representation, in this work, $h_i \in \{v_i, t_i, s_i^z\}$.

There are distinct differences among the negative pairs, especially when we observe that the examples marked in red in Figure 2 serve as negative samples for the green caption. In this grounded space, the distance between the samples becomes crucial. We have refined the contrastive objective to harness the knowledge from the teacher model and to capture the magnitude of the differences between negative pairs. To address this challenge, we introduce a novel training objective, Adaptive Angular margin Contrastive Loss, AdapACSE, which is an adaptation of Eq.10. Instead of fixing margin m_c for all samples in the data, we adapt it to be flexible depending on its negative samples. This new objective improves sentence representation learning by an adaptive margin, $m_c \Delta_{i,j}$, between negative pairs h_i and h_j . Thus, if the $\Delta_{i,j}$ is large and the two samples i, j are much different, the margin is also large to push the representation away, and vice versa. Our approach is illustrated in the right part of Figure 4, we formalize this objective function as AdapACSE:

$$\ell_i^{\text{AdapACSE}} = -\log \frac{e^{\phi(\theta_{i,i^*})/\tau}}{e^{\phi(\theta_{i,i^*})/\tau} + \sum_{j \neq i} e^{\phi(\theta_{i,j} - m_c \Delta_{i,j})/\tau}} \quad (12)$$

where $\Delta_{i,j} = |1 - \alpha_{i,j}|$ is the cosine distance between two samples i and j . Together with the threshold filtering contrastive learning, we rewrite the AdapACSE objective to be:

$$\ell_i^{\text{AdapACSE}'} = \ell_i^{m'} = -\log \frac{e^{\phi(\theta_{i,i^*})/\tau}}{e^{\phi(\theta_{i,i^*})/\tau} + \sum_{j \neq i}^n \varphi_{i,j}^{t,m} e^{\phi(\theta_{i,j} - m_c \Delta_{i,j})/\tau}} \quad (13)$$

Finally, we define our KDMCSE objective as follows:

$$\ell_i^{\text{KDMCSE}} = \frac{\ell_i^{v'} + \ell_i^{t'}}{2} \quad (14)$$

where $\ell_i^{v'}$ and $\ell_i^{t'}$ are AdapACSE Eq. 13 with t, v are text and visual modalities, respectively.

4 Experiments Setup

4.1 Dataset

We employ Flickr30k (Young et al., 2014) and MSCOCO (Lin et al., 2014) as our multimodal datasets. Flickr30k contains 29,783 training images and MSCOCO has 82,783 training images. Every image in these collections comes with multiple captions, usually five captions. As in Gao et al. (2021), we utilize Wiki1M for our text-based dataset, consisting of 10^6 sentences extracted randomly from English Wikipedia.

4.2 Implementation

Language Encoder - Student model We have implemented our work in the Hugging Face Transformers library³ as described by (Wolf et al., 2020). We load our language encoder from the checkpoints of bert-base-uncased and roberta-base. Then, we fine-tune these foundation models with our introduced contrastive objective. To evaluate sentence embeddings, we use the 768-dimensional outputs of the [CLS] token preceding the MLP pooler layer from the transformer-based models.

Multimodal encoder - Teacher model Our teacher model is CLIP (Radford et al., 2021), the model’s weights are initialized with pre-trained clip-vit-base-patch32, where ViT-B/32 Transformer architecture is built as an image encoder and the patch size P is 32.

³<https://github.com/huggingface/transformers>

Because during training, the teacher model will not be fine-tuned, we will pre-extract image and text features \hat{v}, \hat{t} to reduce the computation cost during the training phase.

MLP Projection Heads We use 4 different MLP modules $g_{\{\theta_1, \theta_2, \theta_t, \theta_v\}}$ to project different modalities and objectives. For the pure textual objective using the Wiki1M dataset, sentence embeddings are projected into a 768-dimensional space. For CLIP projection, both the sentence embeddings (student and teacher models) and the image feature vectors are projected into a shared 256-dimensional space.

Parameter Settings In line with the experimental setup described by Zhang et al. (2022a), we explored two primary training scenarios: *wiki+flickr* and *wiki+coco*. Mini-batches were drawn either from the Wiki1M corpus or from the respective caption datasets, with proportions reflecting their relative sizes. Our parameter choices were largely inspired by the guidelines provided by Gao et al. (2021). Specifically, we set the temperature parameters τ and τ' , to 0.05. For model assessment, we performed evaluations on the STS-B dev set after every 125 training iteration, preserving the best performing checkpoint for our final evaluation. Regarding the BERT encoder, we adopted a learning rate of $3e - 5$ and a batch size of 64. For RoBERTa, these parameters were adjusted to $1e - 5$ and 128, respectively. Our training was executed on an A6000 GPU, with each experimental run spanning roughly 5-6 hours and equivalent to MCSE training time.

4.3 Evaluation

We evaluate our trained models on seven STS (Semantic Textual Similarity) tasks, including STS 2012-2016 (Agirre et al., 2012, 2013, 2014, 2015, 2016), the STS Benchmark (Cer et al., 2017), and SICK-Relatedness (Marelli et al., 2014). Each dataset is made up of pairs of sentences with the aim of assigning a similarity score to each pair. In alignment with Gao et al. (2021), we present the Spearman correlation (multiplied by 100) between the official annotations and our predicted scores in an "all" context. This means that we merge all subsets for each task and then provide a comprehensive Spearman correlation.

	Model	STS12	STS13	STS14	STS15	STS16	STS-B	SICK-R	Avg.↑
wiki	SimCSE-BERT \diamond	67.8 \pm 1.6	80.0 \pm 2.1	72.5 \pm 1.7	80.1 \pm 0.8	77.6 \pm 0.8	76.5 \pm 0.8	70.1 \pm 0.9	74.9 \pm 1.1
	SimCSE-RoBERTa \diamond	68.7 \pm 1.0	82.0 \pm 0.5	74.0 \pm 1.0	82.1 \pm 0.4	81.1 \pm 0.4	80.6 \pm 0.3	69.2 \pm 0.2	76.8 \pm 0.5
wiki+flickr	SimCSE-BERT \diamond	69.9 \pm 1.7	79.8 \pm 1.5	72.9 \pm 0.9	81.9 \pm 0.8	77.8 \pm 0.9	76.6 \pm 1.1	68.4 \pm 0.8	75.3 \pm 0.9
	MCSE-BERT \diamond	71.4 \pm 0.9	81.8 \pm 1.3	74.8 \pm 0.9	83.6 \pm 0.9	77.5 \pm 0.8	79.5 \pm 0.5	72.6 \pm 1.4	77.3 \pm 0.5
	KDMCSE-BERT	74.4* \pm 1.4	83.1* \pm 0.9	76.3* \pm 1.1	83.7* \pm 0.8	78.8* \pm 0.9	81.3* \pm 0.9	73.0* \pm 0.9	78.6* \pm 0.8
wiki+coco	SimCSE-RoBERTa \diamond	69.5 \pm 0.9	81.6 \pm 0.5	74.1 \pm 0.6	82.4 \pm 0.3	80.9 \pm 0.5	79.9 \pm 0.3	67.3 \pm 0.5	76.5 \pm 0.4
	MCSE-RoBERTa \diamond	71.7 \pm 0.2	82.7 \pm 0.4	75.9 \pm 0.3	84.0 \pm 0.4	81.3 \pm 0.3	82.3 \pm 0.5	70.3 \pm 1.3	78.3 \pm 0.1
	KDMCSE-RoBERTa	73.6* \pm 0.7	83.8* \pm 0.6	77.4* \pm 0.4	84.0 \pm 0.3	81.5 \pm 0.7	82.3 \pm 0.6	71.2* \pm 0.4	79.1* \pm 0.3
wiki+coco	SimCSE-BERT \diamond	69.1 \pm 1.0	80.4 \pm 0.9	72.7 \pm 0.7	81.1 \pm 0.3	78.2 \pm 0.9	73.9 \pm 0.6	66.6 \pm 1.2	74.6 \pm 0.2
	MCSE-BERT \diamond	71.2 \pm 1.3	79.7 \pm 0.9	73.8 \pm 0.9	83.0 \pm 0.4	77.8 \pm 0.9	78.5 \pm 0.4	72.1 \pm 1.4	76.6 \pm 0.5
	KDMCSE-BERT	73.2* \pm 1.2	80.5* \pm 1.0	75.4* \pm 0.9	83.2 \pm 0.3	79.7* \pm 0.8	79.7* \pm 0.7	73.7* \pm 1.4	77.9* \pm 1.2
	SimCSE-RoBERTa \diamond	66.4 \pm 0.9	80.7 \pm 0.7	72.7 \pm 1.1	81.3 \pm 0.9	80.2 \pm 0.8	76.8 \pm 0.6	65.7 \pm 0.7	74.8 \pm 0.5
wiki+coco	MCSE-RoBERTa \diamond	70.2 \pm 1.7	82.0 \pm 0.7	75.5 \pm 1.2	83.0 \pm 0.6	81.5 \pm 0.7	80.8 \pm 1.0	69.9 \pm 0.6	77.6 \pm 0.8
	KDMCSE-RoBERTa	72.8* \pm 1.5	81.7 \pm 0.9	76.1 \pm 1.1	83.4 \pm 1.0	81.5 \pm 0.6	80.7 \pm 0.8	69.9 \pm 0.6	78.0* \pm 0.7

Table 1: Performance comparison on STS tasks. STS-B: STS Benchmark, SICK-R: SICK-Relatedness, Avg.: average across 7 tasks. \diamond : results from (Zhang et al., 2022a). We train the models using a random seed of 5, presenting the average and standard deviations of our findings. *: difference between MCSE and KDMCSE is significant at $\alpha = 0.05$ according to an independent t-test.

Model	alignment \downarrow		uniformity \downarrow	
	flickr	coco	flickr	coco
MCSE-BERT	0.293	0.267	-2.491	-2.350
KDMCSE-BERT	0.245	0.261	-2.387	-2.383
MCSE-RoBERTa	0.209	0.195	-1.721	-1.418
KDMCSE-RoBERTa	0.174	0.149	-1.952	-1.748

Table 2: The alignment uniformity results of the models when using the BERT and RoBERTa encoder. All models are trained in the *wiki-flickr* setting.

5 Experiments Results

5.1 Main Results

Following the methodology of previous studies, we conducted our model five times using the same hyperparameter settings. Averaged results are delineated in Table 1. A collective overview of our results reveals that our model consistently matches or exceeds the benchmarks set by baseline approaches. MCSE advances beyond SimCSE in the majority of STS evaluations. By integrating auxiliary visual and textual insights derived from the CLIP teacher model, our KDMCSE model showcases pronounced improvements in various downstream tasks. Specifically, when trained on Wiki1M and Flickr30k datasets, KDMCSE increases the performance metrics for BERT (from 77.3 to 78.6) and RoBERTa (from 78.3 to 79.1). Compared to MCSE-RoBERTa, our model exhibits comparable metrics on STS15 and STS-B, with a slight advancement on STS16. In the context of the *wiki+coco* dataset, KDMCSE-BERT notably sur-

passes other methods across most tasks, with scores rising from 76.6 to 77.9 (STS15 being the exception with a modest increase). Our RoBERTa encoder, when trained with the KDMCSE framework, showed higher average results (an increase from 77.6 to 78.0), where tasks such as STS12, STS14, and STS15 registered significant improvements, while others retained or slightly underperformed in comparison.

5.2 Alignment and Uniformity

The concepts of alignment and uniformity are closely linked to contrastive learning, serving as potential metrics to evaluate the quality of the representations derived (Wang and Isola, 2020). Alignment is a property that encourages encoders to produce analogous representations for instances that are alike. This can be quantified using the expected distance between the embeddings of positively paired instances.

$$\mathcal{L}_{align} \triangleq \mathbb{E}_{(x, x^+) \sim p_{pos}} \|f(x) - f(x^+)\|_2^2. \quad (15)$$

And the *uniformity loss* prefers a uniform distribution in the hypersphere:

$$\mathcal{L}_{uniform} \triangleq \log \mathbb{E}_{x, y \stackrel{i.i.d.}{\sim} p_{data}} e^{-2\|f(x) - f(y)\|_2^2}. \quad (16)$$

To delve deeper into the mechanics of our methodology, we assess alignment and uniformity metrics using the STS-B development set. In our

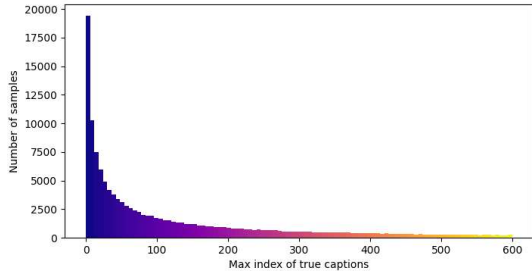


Figure 5: Statistic of the maximum index of true captions when sorting the similarity score between image and text of Flickr dataset.

comparison with MCSE, the results are presented in Table 2. Evidently, when compared to MCSE for both BERT and RoBERTa-based encoders, KDMCSE shows superior performance in both alignment and uniformity metrics. This not only reaffirms the foundational rationale of our method, but also suggests that AdapACSE can effectively enhance the caliber of sentence embeddings.

6 Analysis

6.1 Exploring the Impact of Angular Margin

In our AdapACSE loss function, the angular margin, represented as m , is critical to determining the discriminative capability. We conducted an experiment to better understand the impact of m by varying it in steps of 0.025 radians from 0.025 radians to 0.225 radians. The performance metrics averaged on all tasks are depicted in Figure 7. From the data, we discern an optimal performance when m is set to 0.125. Any deviation from this optimal point, be it an increase or decrease, results in a performance drop. This aligns with our preliminary hypotheses: A smaller value of m could have a negligible effect, whereas an overly large m could negatively skew the representations within the fusion space.

6.2 Threshold Filtering selection

In this section, we will show how to define the threshold hyperparameter during training. We make statistics based on text-text and text-visual similarity scores in the Flickr set, and select based on that statistic. We perform two calculations: (1) the histogram for the similarity threshold and (2) the position of the largest true captions (the annotation captions of that image) that align with the image when sorted with all captions in the dataset with that image. Looking at Figure 5, we see that

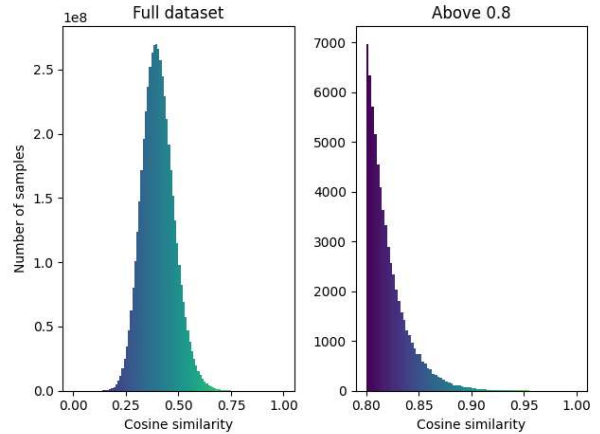


Figure 6: Statistic of the similarity score between image and text of Flickr dataset.

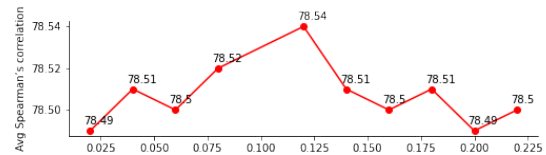


Figure 7: Impact of angular margin m in AdapACSE. The results are reported on average across 7 tasks, and our experiments are conducted on BERT encoder and in *wiki-flickr* setting.

most of the true captions of an image are in the top 100 (out of about 150,000 captions). However, there are still quite a few captions that, even though they accurately describe that image when sorted, can still be pushed to very far indexes (>200). This may be because there are still many similar images in the dataset, leading to captions that describe another image, but are still correct for this image. The image 6 shows that the scores are in a normal distribution. We focus more clearly on samples with high similarity and see that there are only a few pairs with similarity higher than 0.8. and decreases rapidly as it approaches 1.0. Based on the above statistics, we have chosen a threshold of about 0.85-0.9 to eliminate pairs in the contrastive learning objective.

7 Conclusion

In this paper, we introduce KDMCSE, a novel approach to sentence embedding learning that transfers knowledge from a vision and language model with a multimodal contrastive objective to not only align sentences and the corresponding image and text representations of the teacher model but also avoid noisy negative samples. Additionally, our innovative contrastive objective, AdapACSE, ad-

dresses the challenges associated with capturing nuanced differences within negative pairs, strengthening discriminative representation. Our framework lays the groundwork for continued exploration in this domain, fostering advancements in vision and language applications.

8 Limitations

In this study, we shed light on the promising enhancements of KDMCSE in STS benchmarks. However, it is imperative to acknowledge its limitations. In particular, there is a pronounced disparity in both the distribution and volume of word tokens when visual datasets are compared to traditional language corpora. Taking the Book Corpus and Wikipedia as instances, these standard language repositories house billions of words, spanning millions of unique tokens. In contrast, MS COCO, a frequently cited visual-grounded dataset, encompasses merely a million words and about a thousand unique tokens. Another challenge encountered in our methodology is the intricate correlation of the hyperparameter in our contrastive objective, which remains elusive. Its optimal configuration was determined not through systematic understanding, but rather through exhaustive fine-tuning across many experiments. Furthermore, the characterization of "semantic similarity" is highly dependent on the specific task at hand. Although we have seen significant advancements in STS benchmarks, there remains a need to evaluate the disparity in performance between text-only models and their multimodal counterparts across diverse benchmarks. This would provide a more holistic understanding of the efficacy of sentence representations in varied contexts.

Acknowledgement

We thank all anonymous reviewers for their helpful comments. This research/project is supported by the National Research Foundation, Singapore under its AI Singapore Programme, AISG Award No: AISG2-TC-2022-005.

References

Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Inigo Lopez-Gazpio, Montse Maritxalar, Rada Mihalcea, et al. 2015. [Semeval-2015 task 2: Semantic textual similarity, english, spanish and pilot on interpretability](#). In *Proceedings of the 9th international*

workshop on semantic evaluation (SemEval 2015), pages 252–263.

Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2014. [Semeval-2014 task 10: Multilingual semantic textual similarity](#). In *Proceedings of the 8th international workshop on semantic evaluation (SemEval 2014)*, pages 81–91.

Eneko Agirre, Carmen Banea, Daniel Cer, Mona Diab, Aitor Gonzalez Agirre, Rada Mihalcea, German Rigau Claramunt, and Janyce Wiebe. 2016. [Semeval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval 2016)*, pages 497–511.

Eneko Agirre, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. 2012. [Semeval-2012 task 6: A pilot on semantic textual similarity](#). In ** SEM 2012: The First Joint Conference on Lexical and Computational Semantics—Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 385–393.

Eneko Agirre, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and Weiwei Guo. 2013. [* SEM 2013 shared task: Semantic textual similarity](#). In *Second joint conference on lexical and computational semantics (* SEM), volume 1: proceedings of the Main conference and the shared task: semantic textual similarity*, pages 32–43.

Vassileios Balntas, Edgar Riba, Daniel Ponsa, and Krystian Mikolajczyk. 2016. [Learning local feature descriptors with triplets and shallow convolutional neural networks](#). In *British Machine Vision Conference*.

Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 5185–5198.

Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. [Experience grounds language](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8718–8735.

Patrick Bordes, Eloi Zablocki, Laure Soulier, Benjamin Piwowarski, and Patrick Gallinari. 2019. [Incorporating visual semantics into sentence representations within a grounded space](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 696–707.

- Fredrik Carlsson, Amaru Cuba Gyllensten, Evangelia Gogoulou, Erik Ylipää Hellqvist, and Magnus Sahlgren. 2021. [Semantic re-tuning with contrastive tension](#). In *International Conference on Learning Representations*.
- Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. [Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval 2017)*, pages 1–14.
- S. Chopra, R. Hadsell, and Y. LeCun. 2005. [Learning a similarity metric discriminatively, with application to face verification](#). In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 539–546 vol. 1.
- Jiankang Deng, Jia Guo, and Stefanos Zafeiriou. 2018. [Arcface: Additive angular margin loss for deep face recognition](#). *CoRR*, abs/1801.07698.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HIT)*, pages 4171–4186.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6894–6910.
- John Giorgi, Osvald Nitski, Bo Wang, and Gary Bader. 2021. [DeCLUTR: Deep contrastive learning for unsupervised textual representations](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 879–895, Online. Association for Computational Linguistics.
- Hongliang He, Junlei Zhang, Zhenzhong Lan, and Yue Zhang. 2023. [Instance smoothed contrastive learning for unsupervised sentence embedding](#).
- Felix Hill, Kyunghyun Cho, and Anna Korhonen. 2016. [Learning distributed representations of sentences from unlabelled data](#).
- Douwe Kiela, Alexis Conneau, Allan Jabri, and Maximilian Nickel. 2018. [Learning visually grounded sentence representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HIT)*, pages 408–418.
- Taeuk Kim, Kang Min Yoo, and Sang-goo Lee. 2021. [Self-guided contrastive learning for BERT sentence representations](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, pages 2528–2540.
- Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2015a. [Skip-thought vectors](#). *CoRR*, abs/1506.06726.
- Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015b. [Skip-thought vectors](#). In *Advances in neural information processing systems (NeurIPS)*, pages 3294–3302.
- Angeliki Lazaridou, Marco Baroni, et al. 2015. [Combining language and vision with a multimodal skip-gram model](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 153–163.
- Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020. [On the sentence embeddings from pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9119–9130.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. [Microsoft coco: Common objects in context](#). In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 740–755.
- Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. 2017. [Sphereface: Deep hypersphere embedding for face recognition](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6738–6746.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Lajanugen Logeswaran and Honglak Lee. 2018. [An efficient framework for learning sentence representations](#).
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, Roberto Zamparelli, et al. 2014. [A sick cure for the evaluation of compositional distributional semantic models](#). In *International Conference on Language Resources and Evaluation (LREC)*, pages 216–223.
- Tomás Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). *CoRR*, abs/1310.4546.

- Cong-Duy Nguyen, Thong Nguyen, Duc Anh Vu, and Luu Anh Tuan. 2023a. [Improving multimodal sentiment analysis: Supervised angular margin-based contrastive learning for enhanced fusion representation](#).
- Cong-Duy Nguyen, The-Anh Vu-Le, Thong Nguyen, Tho Quan, and Anh-Tuan Luu. 2023b. [Expand bert representation with visual information via grounded language learning with multimodal partial alignment](#). In *Proceedings of the 31st ACM International Conference on Multimedia*, MM '23, page 5665–5673, New York, NY, USA. Association for Computing Machinery.
- Thong Nguyen and Anh Tuan Luu. 2021. [Contrastive learning for neural topic model](#). *Advances in Neural Information Processing Systems*, 34.
- Thong Nguyen, Xiaobao Wu, Anh Tuan Luu, Zhen Hai, and Lidong Bing. 2022. [Adaptive contrastive learning on multimodal transformer for review helpfulness prediction](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10085–10096.
- Thong Thanh Nguyen, Xiaobao Wu, Xinshuai Dong, Cong-Duy T Nguyen, See-Kiong Ng, and Anh Tuan Luu. 2024. [Topic modeling as multi-objective optimization with setwise contrastive learning](#). In *The Twelfth International Conference on Learning Representations*.
- Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. 2018. [Unsupervised learning of sentence embeddings using compositional n-gram features](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pages 8748–8763.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992.
- Kihyuk Sohn. 2016. [Improved deep metric learning with multi-class n-pair loss objective](#). In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Jianlin Su, Jiarun Cao, Weijie Liu, and Yangyiwen Ou. 2021. [Whitening sentence representations for better semantics and faster retrieval](#). *arXiv preprint arXiv:2103.15316*.
- Hao Tan and Mohit Bansal. 2020a. [Vokenization: Improving language understanding via contextualized, visually-grounded supervision](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2066–2080.
- Hao Tan and Mohit Bansal. 2020b. [Vokenization: Improving language understanding with contextualized, visual-grounded supervision](#). *CoRR*, abs/2010.06775.
- Zineng Tang, Jaemin Cho, Hao Tan, and Mohit Bansal. 2021. [Vidlankd: Improving language understanding via video-distilled knowledge transfer](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 24468–24481.
- Anran Wang, Anh Tuan Luu, Chuan-Sheng Foo, Hongyuan Zhu, Yi Tay, and Vijay Chandrasekhar. 2019. [Holistic multi-modal memory network for movie question answering](#). *IEEE Transactions on Image Processing*, 29:489–499.
- Hao Wang, Yangguang Li, Zhen Huang, Yong Dou, Lingpeng Kong, and Jing Shao. 2022a. [Snscse: Contrastive learning for unsupervised sentence embedding with soft negative samples](#).
- Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Zhifeng Li, Dihong Gong, Jingchao Zhou, and Wei Liu. 2018. [Cosface: Large margin cosine loss for deep face recognition](#). *CoRR*, abs/1801.09414.
- Tongzhou Wang and Phillip Isola. 2020. [Understanding contrastive representation learning through alignment and uniformity on the hypersphere](#). In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 9929–9939.
- Wei Wang, Liangzhu Ge, Jingqiao Zhang, and Cheng Yang. 2022b. [Improving contrastive learning of sentence embeddings with case-augmented positives and retrieved negatives](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*. ACM.
- Jie Wei, Guanyu Hu, Luu Anh Tuan, Xinyu Yang, and Wenjing Zhu. 2023. [Multi-scale receptive field graph model for emotion recognition in conversations](#). In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.

- Jie Wei, Guanyu Hu, Xinyu Yang, Anh Tuan Luu, and Yizhuo Dong. 2022. Audio-visual domain adaptation feature fusion for speech emotion recognition. In *INTERSPEECH*, pages 1988–1992.
- Jie Wei, Guanyu Hu, Xinyu Yang, Anh Tuan Luu, and Yizhuo Dong. 2024. Learning facial expression and body gesture visual information for video emotion recognition. *Expert Systems with Applications*, 237:121419.
- Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. 2016. A discriminative feature learning approach for deep face recognition. In *Computer Vision – ECCV 2016*, pages 499–515, Cham. Springer International Publishing.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 38–45.
- Xiaobao Wu, Xinshuai Dong, Thong Nguyen, Chaoqun Liu, Liang-Ming Pan, and Anh Tuan Luu. 2023a. **Infoctm: A mutual information maximization perspective of cross-lingual topic modeling**. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 13763–13771.
- Xiaobao Wu, Xinshuai Dong, Thong Nguyen, and Anh Tuan Luu. 2023b. **Effective neural topic modeling with embedding clustering regularization**. In *International Conference on Machine Learning*. PMLR.
- Xiaobao Wu, Chunping Li, Yan Zhu, and Yishu Miao. 2020. **Short text topic modeling with topic distribution quantization and negative sampling decoder**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1772–1782, Online.
- Xiaobao Wu, Anh Tuan Luu, and Xinshuai Dong. 2022. **Mitigating data sparsity for short text topic modeling by topic-semantic contrastive learning**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2748–2760, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xiaobao Wu, Fengjun Pan, Thong Nguyen, Yichao Feng, Chaoqun Liu, Cong-Duy Nguyen, and Anh Tuan Luu. 2024. **On the affinity, rationality, and diversity of hierarchical topic modeling**. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Yuanmeng Yan, Rumei Li, Sirui Wang, Fuzheng Zhang, Wei Wu, and Weiran Xu. 2021. **ConSERT: A contrastive framework for self-supervised sentence representation transfer**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5065–5075.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. **From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions**. *Transactions of the Association for Computational Linguistics (TACL)*, pages 67–78.
- Miaoran Zhang, Marius Mosbach, David Ifeoluwa Adelani, Michael A. Hedderich, and Dietrich Klakow. 2022a. **Mcse: Multimodal contrastive learning of sentence embeddings**.
- Yan Zhang, Ruidan He, Zuozhu Liu, Kwan Hui Lim, and Lidong Bing. 2020. **An unsupervised sentence embedding method by mutual information maximization**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1601–1610.
- Yanzhao Zhang, Richong Zhang, Samuel Mensah, Xudong Liu, and Yongyi Mao. 2022b. Unsupervised sentence representation via contrastive learning with mixing negatives.
- Yuhao Zhang, Hongji Zhu, Yongliang Wang, Nan Xu, Xiaobo Li, and Binqiang Zhao. 2022c. **A contrastive framework for learning sentence representations from pairwise and triple-wise perspective in angular space**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4892–4903, Dublin, Ireland. Association for Computational Linguistics.
- Kun Zhou, Beichen Zhang, Wayne Xin Zhao, and Jirong Wen. 2022. **Debiased contrastive learning of unsupervised sentence representations**.

Model	STS12	STS13	STS14	STS15	STS16	STS-B	SICK-R	Avg.↑
KDMCSE-BERT	74.4	83.1	76.3	83.7	78.8	81.3	73.0	78.6
KDMCSE without AdapACSE	73.8	81.3	75.4	83.2	79.0	80.5	74.1	78.2
KDMCSE without Threshold Filtering	73.4	82.7	76.6	83.9	78.8	80.8	72.9	78.4

Table 3: Ablation results on STS tasks. STS-B: STS Benchmark, SICK-R: SICK-Relatedness, Avg.: average across 7 tasks. The models are trained with the BERT encoder and in *wiki-flickr* setting.

Model	STS12	STS13	STS14	STS15	STS16	STS-B	SICK-R	Avg.↑
Subset 1								
MCSE	71.54	77.9	73.6	80.83	78.38	79.73	74.66	76.66
KDMCSE	73.00	78.53	74.93	82.14	79.88	81.37	74.08	77.7
Subset 2								
MCSE	71.67	78.56	74.02	81.19	78.54	80.29	73.87	76.88
KDMCSE	73.35	79.51	75.58	83.43	80.07	81.89	73.78	78.23

Table 4: Ablation results on STS tasks. STS-B: STS Benchmark, SICK-R: SICK-Relatedness, Avg.: average across 7 tasks. The models are trained with the BERT encoder and in 1M image-text pairs subsets of the CC12 dataset.

Model	STS12	STS13	STS14	STS15	STS16	STS-B	SICK-R	Avg.↑
MCSE-BERT _{CLIPText}	72.60	81.86	75.28	83.28	78.74	80.77	72.78	77.90
KDMCSE-BERT	74.40	83.10	76.26	83.68	78.79	81.25	73.03	78.64
MCSE-BERT _{CLIPText}	71.39	79.25	73.84	82.65	77.96	78.85	71.74	76.53
KDMCSE-BERT	73.38	80.71	75.55	83.36	79.91	79.85	73.92	78.10

Table 5: Ablation study on STS tasks. STS-B: STS Benchmark, SICK-R: SICK-Relatedness, Avg.: average across 7 tasks.. We train the MCSE models using additional CLIP text embeddings.

A Regularization

To validate the significance of our objectives, we select the best-performing models for each dataset and gradually run with three settings: our KDMCSE, KDMCSE without AdapACSE, and KDMCSE without threshold filtering. The results are presented in Table 3, indicating that the incorporation of all the proposed approaches leads to the best performance in all tasks. The second line shows that not applying threshold filtering can marginally harm performance. This result can be attributed to the fact that, during training, noisy negative samples do not often appear in the batch. In the last row of the table, we show that the removal of AdapACSE significantly affects the performance of both BERT and RoBERTa.

B Larger experiments on image-text paired data

To mitigate scalability and generalization concerns, we conducted extensive experiments using a significantly larger dataset containing one million image-text pairs sourced from the CC12 dataset, equivalent in sample size to Wiki1M. In Table 4, we present the results of these experiments. Our findings consistently demonstrate that our KDMCSE model surpasses MCSE in terms of average performance (Avg) across STS, across various subtasks within CC12M subsets.

C Impact of Adaptive Angular Margin Contrastive objective

To substantiate the significance of our objectives and ensure fairness in our comparisons, we carried out experiments involving MCSE with the inclusion of additional text embeddings from the CLIP text encoder. The results of these experiments are presented in Table 5. Our findings unequivocally indicate that our proposed contrastive objectives have a positive impact on the model’s performance, thereby strengthening our core objectives.

D KDMCSE algorithm

To enhance the presentation of our method, we provide a detailed description of our implementation, adhering to the pseudo code outlined in Algorithm 1.

Algorithm 1 Our KDMCSE algorithm

Input:

Collection of sentences (text-only dataset) $D = \{x_i\}_{i=1}^n$

Collection of image-text pairs (multimodal dataset) $D^M = \{x_i, y_i\}_{i=1}^{n_m}$

Paired sample step $p = \left\lceil \frac{|D|}{|D^M|} \right\rceil$

for $t = 1, 2, 3, \dots, iter$ **do**

if not $t|p$ **then**

 Sample batch $\{x_i\}_{b_s}$ from D

 // Forward through language model and projection

$\mathbf{h}_i^z, \mathbf{h}_i^{z'} = g_\phi(f_\theta(x_i))$

 // Calculate loss

$\frac{1}{b_s} \sum_i^{b_s} \ell_i^S$

else

 Sample batch $\{x_i, y_i\}_{b_s}$ from D^M

$\mathbf{s}_i^z = g_{\phi_g}(f_\theta(x_i, z))$

$\hat{\mathbf{t}}_i = g_{\phi_t}(\hat{\mathbf{t}}_i), \hat{\mathbf{t}}_i = CLIP_{text}(x_i)$

$\mathbf{v}_i = g_{\phi_v}(\hat{\mathbf{v}}_i), \hat{\mathbf{v}}_i = CLIP_{visual}(y_i)$

 // Calculate cosine similarity

$\alpha_{i,j}^{t,t} = \frac{\mathbf{s}_i \cdot \hat{\mathbf{t}}_j}{\|\mathbf{s}_i\|_2 \|\hat{\mathbf{t}}_j\|_2}$

$\alpha_{i,j}^{t,v} = \frac{\mathbf{s}_i \cdot \mathbf{v}_j}{\|\mathbf{s}_i\|_2 \|\mathbf{v}_j\|_2}$

 // Calculate soft-label (text-text, text-visual)

$tt_{score} = \frac{\hat{\mathbf{t}}_i \cdot \hat{\mathbf{t}}_j}{\|\hat{\mathbf{t}}_i\|_2 \|\hat{\mathbf{t}}_j\|_2}$

$tv_{score} = \frac{\hat{\mathbf{t}}_i \cdot \mathbf{v}_j}{\|\hat{\mathbf{t}}_i\|_2 \|\mathbf{v}_j\|_2}$

 // Apply threshold filtering TF based on soft-label

$TF(\alpha_{i,j}^{t,t})$

$TF(\alpha_{i,j}^{t,v})$

 // Calculate AdapACSE loss

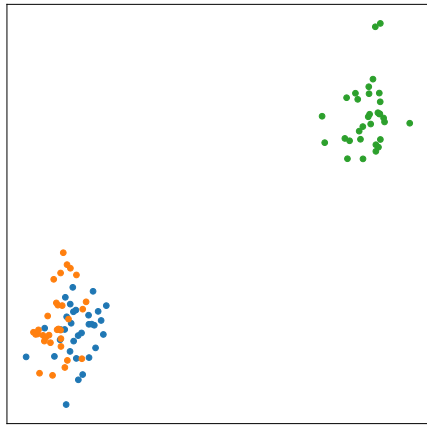
$\frac{1}{b_s} \sum_i^{b_s} \ell_i^{KDMCSE}$

end if

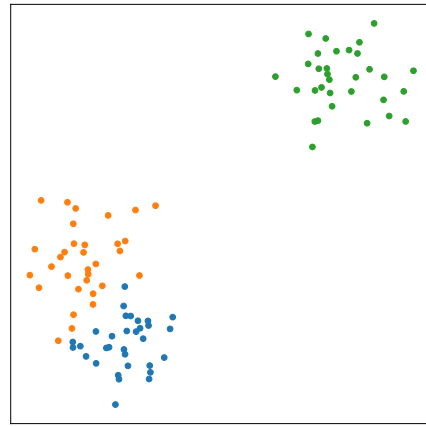
end for

E Sentence embedding visualization

In this appendix, we provide visualization of sentence embedding of our KDMCSE and MCSE models, illustrated in Figures 8, 9, 10, 11 and 12. We visualize sentence representations by repeatedly processing three closely related sentences through the MCSE and KDMCSE models, employing various dropout masks to produce diverse representations for each sentence. Subsequently, we normalize the embeddings and apply t-SNE for dimensionality reduction, facilitating a clearer visualization of the data. For each example, we take three captions: two captions are almost similar in semantics (labeled with orange and blue), and the remaining caption (labeled with green) is different from the other two captions.

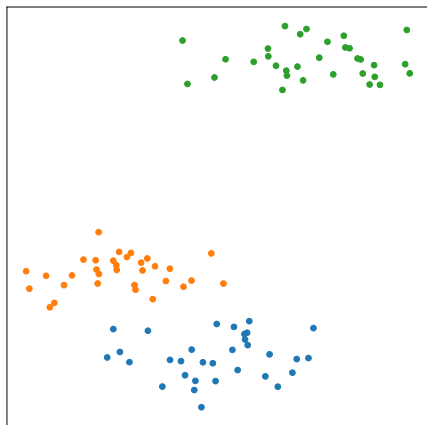


(a) KDMCSE

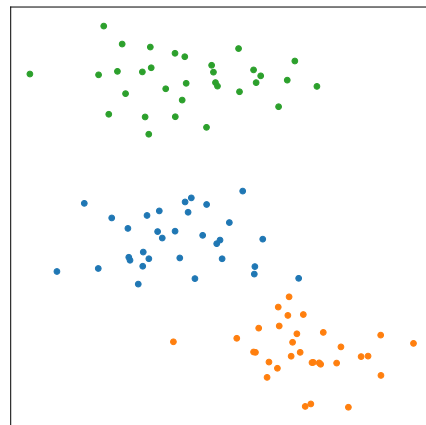


(b) MCSE

Figure 8: **Sentence 1**: A white and black dog and a brown dog in sandy terrain. **Sentence 2**: Brown , black and white dog standing on a sandy slope. **Sentence 3**: Two girls stand up against a red wall.

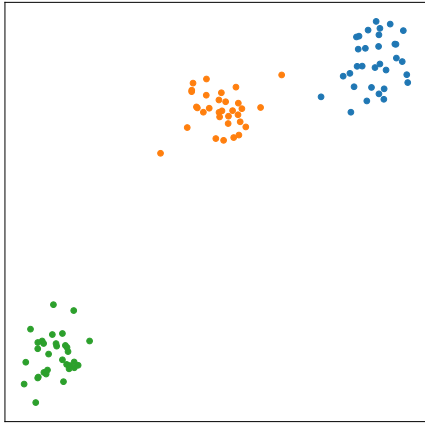


(a) KDMCSE




(b) MCSE

Figure 9: **Sentence 1**: An elderly man holding a pitchfork and doing some yard work. **Sentence 2**: An old bearded man is tilling the soil with a simple wooden plow , with a wooden fence and an old barn in the background. **Sentence 3**: A man in a red shirt is performing an aerial trick with a skateboard on a sidewalk.

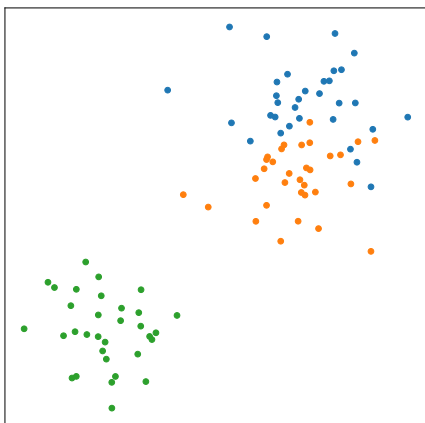


(a) KDMCSE




(b) MCSE

Figure 10: **Sentence 1**: A person wearing a blue shirt , rides a white horse along a dusty country road. **Sentence 2**: A man in a blue tee-shirt , wearing a cap , is riding a white horse down a dirt road , in a rural setting of grass and hills. **Sentence 3**: A man climbing a mountain.

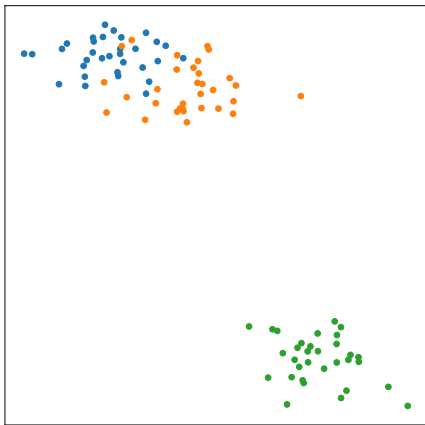


(a) KDMCSE

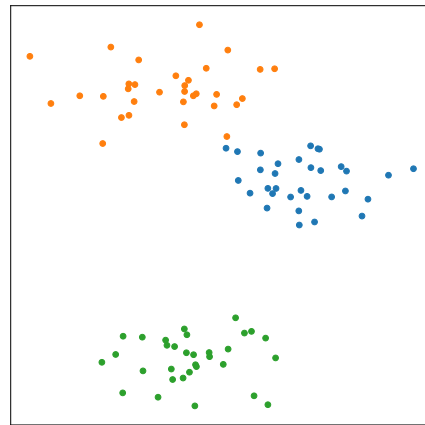


(b) MCSE

Figure 11: **Sentence 1**: A dog is jumping into a swimming pool after a duck. **Sentence 2**: A dog is jumping into a pool to get a duck floating there. **Sentence 3**: A white curly-haired dog runs with a stick in its mouth.



(a) KDMCSE



(b) MCSE

Figure 12: **Sentence 1**: A tattooed woman in a black dress holds a drink while sitting at a table in a dimly-lit room. **Sentence 2**: Scantly clad woman in black waits in a restaurant with other patrons. **Sentence 3**: Woman bending down to pick up a tennis ball outside in front of a wall with graffiti on it.