

Exploring the Impact of Negative Samples of Contrastive Learning: A Case Study of Sentence Embedding

Rui Cao^{1,*} and Yihao Wang^{1,*} and Yuxin Liang¹

Ling Gao^{1,†} and Jie Zheng¹ and Jie Ren² and Zheng Wang³

¹Northwest University, ²Shaanxi Normal University, and ³University of Leeds
{caorui, wangyihao, liangyuxin}@stumail.nwu.edu.cn,
{gl, jzheng}@nwu.edu.cn, renjie@snnu.edu.cn, z.wang5@leeds.ac.uk

Abstract

Contrastive learning is emerging as a powerful technique for extracting knowledge from unlabeled data. This technique requires a balanced mixture of two ingredients: positive (similar) and negative (dissimilar) samples. This is typically achieved by maintaining a queue of negative samples during training. Prior works in the area typically uses a fixed-length negative sample queue, but how the negative sample size affects the model performance remains unclear. The opaque impact of the number of negative samples on performance when employing contrastive learning aroused our in-depth exploration. This paper presents a momentum contrastive learning model with negative sample queue for sentence embedding, namely MoCoSE. We add the prediction layer to the online branch to make the model asymmetric and together with EMA update mechanism of the target branch to prevent the model from collapsing. We define a maximum traceable distance metric, through which we learn to what extent the text contrastive learning benefits from the historical information of negative samples. Our experiments find that the best results are obtained when the maximum traceable distance is at a certain range, demonstrating that there is an optimal range of historical information for a negative sample queue. We evaluate the proposed unsupervised MoCoSE on the semantic text similarity (STS) task and obtain an average Spearman’s correlation of 77.27%. Source code is available [here](#).

1 Introduction

In recent years, unsupervised learning has been brought to the fore in deep learning due to its ability to leverage large-scale unlabeled data. Various unsupervised contrastive models is emerging, continuously narrowing down the gap between supervised and unsupervised learning. Contrastive learning

suffers from the problem of model collapse, where the model converges to a constant value and the samples all mapped to a single point in the feature space. Negative samples are an effective way to solve this problem.

In computer vision, SimCLR from Chen (Chen et al., 2020) and MoCo from He (He et al., 2020) is known for using negative samples and get the leading performance in the contrastive learning. SimCLR uses different data augmentation (e.g., rotation, masking, etc.) on the same image to construct positive samples, and negative samples are from the rest of images in the same batch. MoCo goes a step further by randomly select the data in entire unlabeled training set to stack up a first-in-first-out negative sample queue.

Recently in natural language processing, contrastive learning has been widely used in the task of learning sentence embedding. One of current state-of-the-art unsupervised method is SimCSE (Gao et al., 2021). Its core idea is to make similar sentences in the embedding space closer while keeping dissimilar away from each other. SimCSE uses dropout mask as augmentation to construct positive text sample pairs, and negative samples are picked from the rest of sentences in the same batch. The mask adopted from the standard Transformer makes good use of the minimal form of data augmentation brought by the dropout. Dropout results in a minimal difference without changing the semantics, reducing the negative noise introduced by augmentation. However, the negative samples in SimCSE are selected from the same training batch with a limited batch size. Our further experiments show that SimCSE does not obtain improvement as the batch size increases, which arouses our interest in using the negative sample queue.

To better digging in the performance of contrastive learning on textual tasks, we build a contrastive model consisting of a two-branch structure and a negative sample queue, namely MoCoSE

*Authors contributed equally to this manuscript.

†Corresponding author.

(Momentum Contrastive Sentence Embedding with negative sample queue). We also introduce the idea of asymmetric structure from BYOL (Grill et al., 2020) by adding a prediction layer to the upper branch (i.e., the online branch). The lower branch (i.e., the target branch) is updated with exponential moving average (EMA) method during training. We set a negative sample queue and update it using the output of target branch. Unlike directly using negative queue as in MoCo, for research purpose, we set an initialization process with a much smaller negative queue, and then filling the entire queue through training process, and update normally. We test both character-level (e.g., typo, back translation, paraphrase) and vector-level (e.g., dropout, shuffle, etc.) data augmentations and found that for text contrastive learning, the best results are obtained by using FGSM and dropout as augmentations.

Using the proposed MoCoSE model, we design a series of experiments to explore the contrastive learning for sentence embedding. We found that using different parts of samples from the negative queue leads to different performance. In order to test how much text contrastive learning benefit from historical information of the model, we proposed a maximum traceable distance metric. The metric calculates how many update steps before the negative samples in the queue are pushed in, and thus measures the historical information contained in the negative sample queue. We find that the best results can be achieved when the maximum traceable distance is within a certain range, reflected in the performance of uniformity and alignment of the learned text embedding. Which means there is an optimal interval for the length of negative sample queue in text contrastive learning model.

Our main contributions are as follows:

1. We combine several advantages of frameworks from image contrastive learning to build a more generic text unsupervised contrastive model. We carried out a detailed study of this model to achieve better results on textual data.

2. We evaluate the role of negative queue length and the historical information that the queue contains in text contrastive learning. By slicing the negative sample queue and using different positions of negative samples, we found those near the middle of the queue provides a better performance.

3. We define a metric called 'maximum traceable distance' to help analyze the impact of negative

sample queue by combining the queue length, EMA parameter, and batch size. We found that changes in MTD reflects in the performance of uniformity and alignment of the learned text embedding.

2 Related Work

Contrastive Learning in CV

Contrast learning is a trending and effective unsupervised learning framework that was first applied to the computer vision (Hadsell et al., 2006). The core idea is to make the features of images within the same category closer and the features in different categories farther apart. Most of the current work are using two-branch structure (Chen et al., 2021). While influential works like SimCLR and MoCo using positive and negative sample pairs, BYOL (Grill et al., 2020) and SimSiam (Chen and He, 2021) can achieve the same great results with only positive samples. BYOL finds that by adding a prediction layer to the online branch to form an asymmetric structure and using momentum moving average to update the target branch, can train the model using only positive samples and avoid model collapsing. SimSiam explores the possibility of asymmetric structures likewise. Therefore, our work introduces this asymmetric idea to the text contrastive learning to prevent model collapse. In addition to the asymmetric structure and the EMA mechanism to avoid model collapse, some works consider merging the constraint into the loss function, like Barlow Twins (Zbontar et al., 2021), W-MSE (Ermolov et al., 2021), and ProtoNCE (Li et al., 2021).

Contrastive Learning in NLP

Since BERT (Devlin et al., 2018) redefined state-of-the-art in NLP, leveraging the BERT model to obtain better sentence representation has become a common task in NLP. A straightforward way to get sentence embedding is by the $[CLS]$ token due to the Next Sentence Prediction task of BERT. But the $[CLS]$ embedding is non-smooth anisotropic in semantic space, which is not conducive to STS tasks, this is known as the representation degradation problem (Gao et al., 2019). BERT-Flow (Li et al., 2020) and BERT-whitening (Su et al., 2021) solve the degradation problem by post-processing the output of BERT. SimCSE found that utilizing contrasting mechanism can also alleviate this problem.

Data augmentation is crucial for contrastive learning. In CLEAR (Wu et al., 2020), word and

phrase deletion, phrase order switching, synonym substitution is served as augmentation. CERT (Fang and Xie, 2020) mainly using back-and-forth translation, and CLINE (Wang et al., 2021) proposed synonym substitution as positive samples and antonym substitution as negative samples, and then minimize the triplet loss between positive, negative cases as well as the original text. ConSERT (Yan et al., 2021) uses adversarial attack, token shuffling, cutoff, and dropout as data augmentation. CLAE (Ho and Nvasconcelos, 2020) also introduces Fast Gradient Sign Method, an adversarial attack method, as text data augmentation. Several of these augmentations are also introduced in our work. The purpose of data augmentation is to create enough distinguishable positive and negative samples to allow contrastive loss to learn the nature of same data after different changes. Works like (Mitrovic et al., 2020) points out that longer negative sample queues do not always give the best performance. This also interests us how the negative queue length affects the text contrastive learning.

3 Method

Figure 1 depicts the architecture of proposed MoCoSE. In the embedding layer, two versions of the sentence embedding are generated through data augmentation ($dropout = 0.1$ and $fgsm = 5e - 9$). The resulting two slightly different embeddings then go through the online and target branch to obtain the query and key vectors respectively. The structure of encoder, pooler and projection of online and target branch is identical. We add a prediction layer to the online branch to make asymmetry between online and target branch. The pooler, projection and prediction layers are all composed of several fully connected layers.

Finally, the model calculates contrasting loss between query, key and negative queue to update the online branch. In the process, key vector serves as positive sample with respect to the query vector, while the sample from queue serves as negative sample to the query. The target branch truncates the gradient and updated with the EMA mechanism. The queue is a first-in-first-out collection of negative samples with size K which means it sequentially stores the key vectors generated from the last few training steps.

The PyTorch style pseudo-code for training MoCoSE with the negative sample queue is shown in

Algorithm 1 in Appendix A.3.

Data Augmentation Comparing with SimCSE, we tried popular methods in NLP such as paraphrasing, back translation, adding typos etc., but experiments show that only adversarial attacks and dropout have improved the results. We use FGSM (Goodfellow et al., 2015) (Fast Gradient Sign Method) as adversarial attack. In a white-box environment, FGSM first calculates the derivative of model with respect to the input, and use a sign function to obtain its specific gradient direction. Then, after multiplying it by a step size, the resulting 'perturbation' is added to the original input to obtain the sample under the FGSM attack.

$$x' = x + \varepsilon \cdot \text{sign}(\nabla_x \mathcal{L}(x, \theta)) \quad (1)$$

Where x is the input to the embedding layer, θ is the online branch of the model, and $\mathcal{L}(\cdot)$ is the contrastive loss computed by the query, key and negative sample queue. ∇_x is the gradient computed through the network for input x , $\text{sign}(\cdot)$ is the sign function, and ε is the perturbation parameter which it controls how much noise it added.

EMA and Asymmetric Branches Our model uses EMA mechanism to update the target branch. Formally, denoting the parameters of online and target branch as θ_o and θ_t , EMA decay weight as η , we update θ_t by:

$$\theta_t \leftarrow \eta\theta_t + (1 - \eta)\theta_o \quad (2)$$

Experiments demonstrate that not using EMA leads to model collapsing, which means the model did not converge during training. The prediction layer we added on the online branch makes two branches asymmetric to further prevent the model from collapsing. For more experiment details about symmetric model structure without EMA mechanism, please refer to Appendix A.2.

Negative Sample Queue The negative sample queue has been theoretically proven to be an effective means of preventing model from collapsing. Specifically, both the queue and the prediction layer of the upper branch serves to disperse the output feature of the upper and lower branches, thus ensuring that the contrastive loss obtains features with sufficient uniformity. We also set a buffer for the initialization of the queue, i.e., only a small portion of the queue is randomly initialized at the beginning, and then enqueue and dequeue normally until the end.

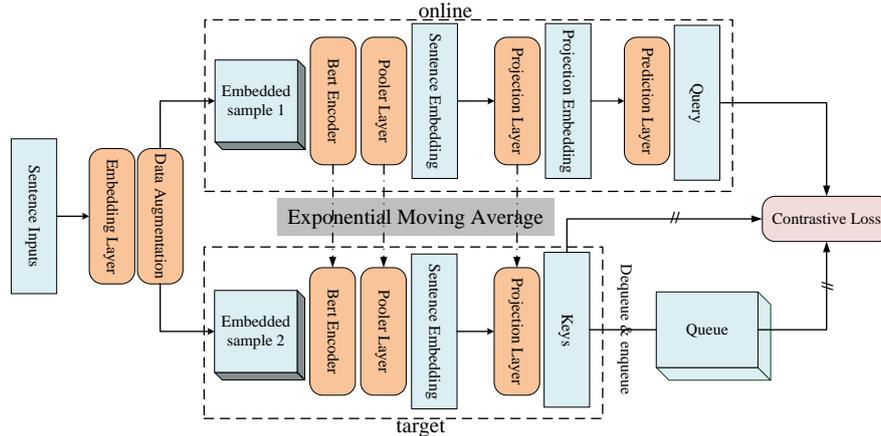


Figure 1: The model structure of MoCoSE. The embedding layer consists of a BERT embedding layer with additional data augmentation. The pooler, projection, and predictor layers all keep the same dimensions with the encoder layer. The MoCoSE minimizes contrastive loss between query, queue and keys (i.e. InfoNCE loss).

Contrastive Loss Similar to MoCo, we also use InfoNCE (Oord et al., 2018) as contrastive loss, as shown in eq.(3).

$$\mathcal{L} = -\log \frac{\exp(q \cdot k/\tau)}{\exp(q \cdot k/\tau) + \sum_l \exp(q \cdot l/\tau)} \quad (3)$$

Where, q refers to the query vectors obtained by the online branch; k refers to the key vectors obtained by the target branch; and l is the negative samples in the queue; τ is the temperature parameter.

4 Experiments

4.1 Settings

We train with a randomly selected corpus of 1 million sentences from the English Wikipedia, and we conduct experiments on seven standard semantic text similarity (STS) tasks, including STS 2012—2016 (Agirre et al., 2012, 2013, 2014, 2015, 2016), STSBenchmark (Cer et al., 2017) and SICK-Relatedness (Wijnholds and Moortgat, 2021). The SentEval¹ toolbox is used to evaluate our model, and we use the Spearman’s correlation to measure the performance. We start our training by loading pre-trained BERT checkpoints² and use the $[CLS]$ token embedding from the model output as the sentence embedding. In addition to the semantic similarity task, we also evaluate on seven transfer learning tasks to test the generalization performance of the model. For text augmentation, we tried several vector-level methods mentioned in ConSERT,

¹<https://github.com/facebookresearch/SentEval>

²<https://huggingface.co/models>

including position shuffle, token dropout, feature dropout. In addition, we also tried several text-level methods from the nlpaug³ toolkit, including synonym replace, typo, back translation and paraphrase.

Training Details The learning rate of MoCoSE-BERT-base is set to $3e-5$, and for MoCoSE-BERT-large is $1e-5$. With a weight decay of $1e-6$, the batch size of the base model is 64, and the batch size of the large model is 32. We validate the model every 100 step and train for one epoch. The EMA decay weight η is incremented from 0.75 to 0.95 by the cosine function. The negative queue size is 512. For more information please refer to Appendix A.1.

4.2 Main Results

We compare the proposed MoCoSE with several commonly used unsupervised methods and the current state-of-the-art contrastive learning method on the text semantic similarity (STS) task, including average GloVe embeddings (Pennington et al., 2014), average BERT or RoBERTa embeddings, BERT-flow, BERT-whitening, ISBERT (Zhang et al., 2020a), DeCLUTR (Giorgi et al., 2021), CT-BERT (Carlsson et al., 2021) and SimCSE.

As shown in Table 1, the average Spearman’s correlation of our best model is 77.27%, outperforming unsupervised SimCSE with BERT-base. Our model outperforms SimCSE on STS2012, STS2015, and STS-B, and SimCSE perform better on the STS2013 task. Our MoCoSE-BERT-

³<https://github.com/makcedward/nlpaug>

Model	STS12	STS13	STS14	STS15	STS16	STS-B	SICK-R	Avg.
Unsupervised Models (Base)								
GloVe (avg.)	55.14	70.66	59.73	68.25	63.66	58.02	53.76	61.32
BERT (first-last avg.)	39.70	59.38	49.67	66.03	66.19	53.87	62.06	56.70
BERT-flow	58.40	67.10	60.85	75.16	71.22	68.66	64.47	66.55
BERT-whitening	57.83	66.90	60.90	75.08	71.31	68.24	63.73	66.28
IS-BERT	56.77	69.24	61.21	75.23	70.16	69.21	64.25	66.58
CT-BERT	61.63	76.80	68.47	77.50	76.48	74.31	69.19	72.05
RoBERTa (first-last avg.)	40.88	58.74	49.07	65.63	61.48	58.55	61.63	56.57
RoBERTa-whitening	46.99	63.24	57.23	71.36	68.99	61.36	62.91	61.73
DeCLUTR-RoBERT	52.41	75.19	65.52	77.12	78.63	72.41	68.62	69.99
SimCSE	68.40	82.41	74.38	80.91	78.56	76.85	72.23	76.25
MoCoSE	71.48	81.40	74.47	83.45	78.99	78.68	72.44	77.27
Unsupervised Models (Large)								
SimCSE-RoBERTa	72.86	83.99	75.62	84.77	81.80	81.98	71.26	78.90
SimCSE-BERT	70.88	84.16	76.43	84.50	79.76	79.26	73.88	78.41
MoCoSE-BERT	74.50	84.54	77.32	84.11	79.67	80.53	73.26	79.13

Table 1: Spearman correlation of MoCoSE on seven semantic text similarity tasks. We compared with the state-of-the-art method SimCSE. MoCoSE achieves the best results with both BERT-base and BERT-large pre-trained models.

Model	MR	CR	SUBJ	MPQA	SST	TREC	MRPC	Avg.
Unsupervised Model (Base)								
GloVe (avg.)	77.25	78.30	91.17	87.85	80.18	83.00	72.87	81.52
Skip-thought	76.50	80.10	93.60	87.10	82.00	92.20	73.00	83.50
Avg. BERT embeddings	78.66	86.25	94.37	88.66	84.40	92.80	69.54	84.94
BERT-[CLS]embedding	78.68	84.85	94.21	88.23	84.13	91.40	71.13	84.66
SimCSE-RoBERTa	81.04	87.74	93.28	86.94	86.60	84.60	73.68	84.84
SimCSE-BERT	81.18	86.46	94.45	88.88	85.50	89.80	74.43	85.81
MoCoSE-BERT	81.07	86.43	94.76	89.70	86.35	84.06	75.86	85.46
Unsupervised Model (Large)								
SimCSE-RoBERTa	82.74	87.87	93.66	88.22	88.58	92.00	69.68	86.11
MoCoSE-BERT	83.71	89.07	95.58	90.26	87.96	84.92	76.81	86.90

Table 2: Performance of MoCoSE on the seven transfer tasks. We compare the performance of MoCoSE and other models on the seven transfer tasks evaluated by SentEval, and MoCoSE remains at a comparable level with the SimCSE.

large model outperforms SimCSE-BERT-Large by about 0.7 on average, mainly on STS12, STS13, and STS14 tasks, and maintains a similar level on other tasks.

Furthermore, we also evaluate the performance of MoCoSE on the seven transfer tasks provided by SentEval. As shown in Table 2, MoCoSE-BERT-base outperforms most of the previous unsupervised method, and is on par with SimCSE-BERT-base.

5 Empirical Study

To further explore the performance of the MoCo-like contrasting model on learning sentence embedding, we set up the following ablation trials.

5.1 EMA Decay Weight

We use EMA to update the model parameters for the target branch and find that EMA decay weight affects the performance of the model. The EMA decay weight affects the update process of the model, which further affects the vectors involved in the contrastive learning process. Therefore, we set different values of EMA decay weight and train the model with other hyperparameters held constant. As shown in Table 3 and Appendix A.5, the best result is obtained when the decay weight of EMA is set to 0.85. Compared to the choice of EMA decay weight in CV (generally as large as 0.99), the value of 0.85 in our model is smaller, which means that the model is updated faster. We speculate that this is because the NLP model is more sensitive in the fine-tuning phase and the model weights change

EMA	0.5	0.8	0.85	0.9	0.95	0.99
Avg.	75.76	75.19	76.49	76.05	76.08	75.12

Table 3: Effect of EMA decay weight on model performance. The best results are obtained with the EMA decay weight at 0.85.

	Proj.	Pred.	Corr.		Proj.	Pred.	Corr.
		1	60.46			1	66.96
0		2	62.67	2	2	66.29	
		3	63.62		3	61.57	
		1	76.74		1	31.51	
1	2	76.89		3	2	43.97	
		3	76.24		3	39.13	

Table 4: The impact of different combinations of projection and predictor on the model.

more after each step of the gradient, so a faster update speed is needed.

5.2 Projection and Prediction

Several papers have shown (e.g. Section F.1 in BYOL (Grill et al., 2020)) that the structure of projection and prediction layers in a contrastive learning framework affects the performance of the model. We combine the structure of projection and prediction with different configurations and train them with the same hyperparameters. As shown in Table 4, the best results are obtained when the projection is 1 layer and the prediction has 2 layers. The experiments also show that the removal of projection layers degrades the performance of the model.

5.3 Data Augmentation

We investigate the effect of some widely-used data augmentation methods on the model performance. As shown in Table 5, cut off and token shuffle do not improve, even slightly hurt the model’s performance. Only the adversarial attack (FGSM) has slight improvement on the performance. Therefore, in our experiments, we added FGSM as a default data augmentation of our model in addition to dropout. Please refer to Appendix A.7 for more FGSM parameters results. We speculate that the reason token cut off is detrimental to the model results is that the cut off perturbs too much the vector formed by the sentences passing through the embedding layer. Removing one word from the text may have a significant impact on the semantics. We tried two parameters 0.1 and 0.01 for the feature cut off, and with these two parameters, the results of using the feature cut off is at most the same as

Augmentation Methods	Avg.
Dropout only	76.76
+ FGSM ($\epsilon=5e-9$)	77.04
+ Position_shuffle (True)	73.80
+ Token dropout (prob=0.1)	41.32
+ Feature dropout (prob=0.01)	76.33
+ Feature dropout (prob=0.1)	71.62
+ Typos	22.32
+ Synonym replace (roberta-base)	28.70
+ Paraphrasing (xlnet-base-cased)	60.45
+ Backtranslation (en->de->en)	69.35

Table 5: The effect of different data augmentation methods.

without using feature the cut off, so we discard the feature cut off method. More results can be found in Appendix A.6.

The token shuffle is slightly, but not significantly, detrimental to the results of the model. This may be due to that BERT is not sensitive to the position of token. In our experiment, the sentence-level augmentation methods also failed to outperform than the drop out, FGSM and position shuffle.

Among the data augmentation methods, only FGSM together with dropout improves the results, which may due to the adversarial attack slightly enhances the difference between the two samples and therefore enables the model to learn a better representation in more difficult contrastive samples.

5.4 Predictor Mapping Dimension

The predictor maps the representation to a feature space of a certain dimension. We investigate the effect of the predictor mapping dimension on the model performance. Table 6.a shows that the predictor mapping dimension can seriously impair the performance of the model when it is small, and when the dimension rises to a suitable range or larger, it no longer has a significant impact on the model. This may be related to the intrinsic dimension of the representation, which leads to the loss of semantic information in the representation when the predictor dimension is smaller than the intrinsic dimension of the feature, compromising the model performance. We keep the dimension of the predictor consistent with the encoder in our experiments. More results can be found in Appendix A.8.

5.5 Batch Size

With a fixed queue size, we investigated the effect of batch size on model performance, the results is in Table 6.b, and the model achieves the best per-

Dim	Avg.	Size	Avg.
256	73.91	32	73.86
512	76.07	64	77.25
768	77.04	128	76.78
1024	77.02	256	76.62
2048	77.03		

(a)

(b)

Table 6: (a) Impact of prediction dimension on model performance. (b) Impact of batch size on the model with fixed queue size. Both table under a batch size setting to 512.

formance when the batch size is 64. Surprisingly the model performance does not improve with increasing batch size, which contradicts the general experience in image contrastive learning. This is one of our motivations for further exploring the effect of the number of negative samples on the model.

5.6 Size of Negative Sample Queue

The queue length determines the number of negative samples, which directly influence performance of the model. We first test the size of negative sample queue to the model performance. With queue size longer than 1024, the results get unstable and worse. We suppose this may be due to the random interference introduced to the training by filling the initial negative sample queue. This interference causes a degradation of the model’s performance when the initial negative sample queue becomes longer. To reduce the drawbacks carried out by this randomness, we changed the way the negative queue is initialized. We initialize a smaller negative queue, then fill the queue to its set length in the first few updates, and then update normally. According to experiments, the model achieves the highest results when the negative queue size set to 512 and the smaller initial queue size set to 128.

According to the experiments of MoCo, the increase of queue length improves the model performance. However, as shown in Table 7, increasing the queue length with a fixed batch size decreases our model performance, which is not consistent with the observation in MoCo. We speculate that this may be due to that NLP models updating faster, and thus larger queue lengths store too much outdated feature information, which is detrimental to the performance of the model. Combined with the observed effect of batch size, we further conjecture that the effect of the negative sample queue on model performance is controlled by the model

Initial Size	Queue Size				
	128	256	512	1024	4096
w.o. init.	76.40	76.19	75.38	76.63	50.17
init. 1/4 queue	75.92	76.34	77.30	76.20	50.42
init. 1/2 queue	76.16	76.39	76.94	76.57	38.74
init. all (normal)	76.87	75.81	76.29	76.45	45.80

Table 7: Correlation performance of initializing different proportion of negative queue with different negative queue size.

Corr.	0~512	256~768	512~1024	Without 256~768	All
Avg.	76.10	77.02	75.71	76.18	76.86

Table 8: The impact of negative samples at different locations in the queue on the model performance.

history information contained in the negative sample in the queue. See Appendix A.9 and A.10 for more results of the effect of randomization size and queue length.

Since the queue is first-in-first-out, to test the hypothesis above, we sliced the negative sample queue and use different parts of the queue to participate in loss calculation. Here, we set the negative queue length to 1024, the initial queue size to 128, and the batch size to 256. Thus, 256 negative samples will be pushed into the queue for each iteration. We take 0 ~ 512, 256 ~ 768, 512 ~ 1024, a concatenated slice 0 ~ 256 and 768 ~ 1024, and all negative sample queues respectively for testing. The experiment results are shown in Table 8.

The experiments show that the model performs best when using the middle part of the queue. So we find that the increase in queue length affects the model performance not only because of the increased number of negative samples, but more because it provides historical information within a certain range.

5.7 Maximum Traceable Distance Metric

To testify there are historical information in negative sample queue influencing the model performance, we define a Maximum Traceable Distance Metric d_{trace} to help explore the phenomenon.

$$d_{trace} = \frac{1}{1 - \eta} + \frac{queue_size}{batch_size} \quad (4)$$

The η refers to the decay weight of EMA. The d_{trace} calculates the update steps between the current online branch and the oldest negative samples in the queue. The first term of the formula represents the traceable distance between target and

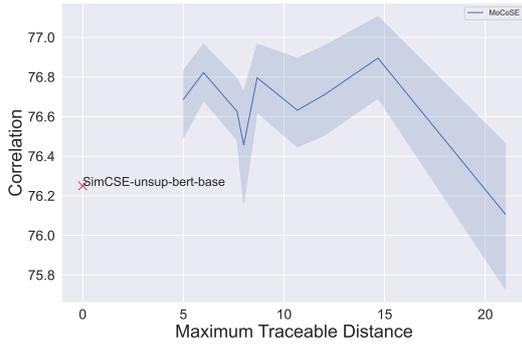


Figure 2: The relationship between traceable distance and model correlation.

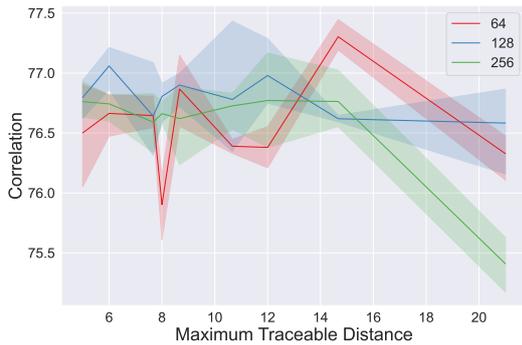


Figure 3: The batch size does not invalidate the traceable distance. The traceable distance needs to be maintained within a reasonable range even for different batch sizes. This explains why increasing the batch size only does not improve the performance, because increasing the batch size only can cause the distance changes into unsuitable regions.

online branch due to the EMA update mechanism. The second term represents the traceable distance between the negative samples in the queue and the current target branch due to the queue’s first-in-first-out mechanism. The longer traceable distance, the wider the temporal range of the historical information contained in the queue. We obtained different value of traceable distance by jointly adjust the decay weight, queue size, and batch size. As shown in Figure 2 and Figure 3, the best result of BERT base is obtained with d_{trace} is set around 14.67. The best result of BERT large shows the similar phenomenon, see Appendix A.11 for details. This further demonstrates that in text contrastive learning, the historical information used should be not too old and not too new, and the appropriate traceable distance between branches is also important. Some derivations about eq.4 can be found in Appendix A.12.

However, for an image contrast learning model, like MoCo, experimental results suggests that

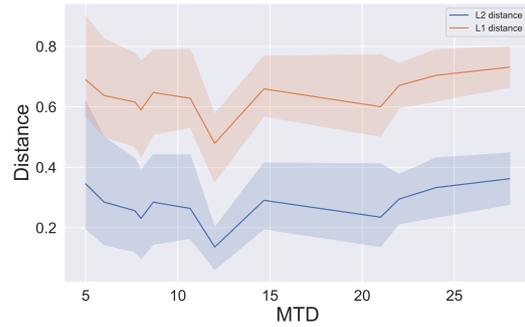


Figure 4: L1 and L2 distances of learned embedding’s uniformity and alignment with a fixed point changes along with MTD.

longer queue size increases the performance. We believe that this is due to the phenomenon of unique anisotropy (Zhang et al., 2020b) of text that causes such differences. The text is influenced by the word frequency producing the phenomenon of anisotropy with uneven distribution, which is different from the near-uniform distribution of pixel points of image data. Such a phenomenon affects the computation of the cosine similarity (Wang and Isola, 2020), and the loss of InfoNCE that we use depends on it, which affects the performance of the model through the accumulation of learning steps. To test such a hypothesis, we use alignment and uniformity to measure the distribution of the representations in space and monitor the corresponding values of alignment and uniformity for different MTDs. As shown in the Figure 4, it can be found that a proper MTD allows the alignment and uniformity of the model to reflects an optimal combination. The change in MTD is reflected in the performance of uniformity and alignment of the learned text embedding, and the increase and decrease of MTD is a considering result of uniformity and alignment moving away from their optimal combination region.

6 Conclusion

In this study, we propose MoCoSE, it applies the MoCo-style contrastive learning model to the empirical study of sentence embedding. We conducted experiments to study every detail of the model to provide some experiences for text contrastive learning. We further delve into the application of the negative sample queue to text contrastive learning and propose a maximum traceable distance metric to explain the relation between the queue size and model performance.

Acknowledgments

Our work is supported by the National Key Research and Development Program of China under grant No.2019YFC1521400, National Natural Science Foundation of China under grant No.62072362 and No.61902229 and International Science and Technology Cooperation Project of Shaanxi (2020KW-006).

References

- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Inigo Lopez-Gazpio, Montse Maritxalar, Rada Mihalcea, German Rigau, Larraitz Uribe, and Janyce Wiebe. 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 M. Cer, Mona T. Diab, Aitor Gonzalez-Agirre, Rada Mihalcea, German Rigau, 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)*, volume 1, 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*, volume 1, pages 32–43.
- Fredrik Carlsson, Magnus Sahlgren, Evangelia Gogoulou, Amaru Cuba Gyllensten, and Erik Ylipää Hellqvist. 2021. Semantic re-tuning with contrastive tension. In *ICLR 2021: The Ninth International Conference on Learning Representations*.
- Daniel M. Cer, Mona T. Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14.
- Pengguang Chen, Shu Liu, and Jiaya Jia. 2021. Jigsaw clustering for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11526–11535.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- Xinlei Chen and Kaiming He. 2021. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina N. Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. 2021. Whitening for self-supervised representation learning. In *ICML 2021: 38th International Conference on Machine Learning*, pages 3015–3024.
- Hongchao Fang and Pengtao Xie. 2020. Cert: Contrastive self-supervised learning for language understanding. *arXiv preprint arXiv:2005.12766*.
- Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tiejun Liu. 2019. Representation degeneration problem in training natural language generation models. *arXiv preprint arXiv:1907.12009*.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.
- 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.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *ICLR 2015: International Conference on Learning Representations 2015*.

- Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. 2020. Bootstrap your own latent: A new approach to self-supervised learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 21271–21284.
- R. Hadsell, S. Chopra, and Y. LeCun. 2006. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738.
- Chih-Hui Ho and Nuno Nvasconcelos. 2020. Contrastive learning with adversarial examples. In *Advances in Neural Information Processing Systems*, volume 33, pages 17081–17093.
- Quentin Lhoest, Albert Villanova del Moral, Patrick von Platen, Thomas Wolf, Yacine Jernite, Abhishek Thakur, Lewis Tunstall, Suraj Patil, Mariama Drame, Julien Chaumond, Julien Plu, Joe Davison, Simon Brandeis, Victor Sanh, Teven Le Scao, Kevin Canwen Xu, Nicolas Patry, Steven Liu, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Nathan Raw, Sylvain Lesage, Anton Lozhkov, Matthew Carrigan, Théo Matussière, Leandro von Werra, Lysandre Debut, Stas Bekman, and Clément Delangue. 2021. [huggingface/datasets: 1.13.2](#).
- 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.
- Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi. 2021. Prototypical contrastive learning of unsupervised representations. In *ICLR 2021: The Ninth International Conference on Learning Representations*.
- Jovana Mitrovic, Brian McWilliams, and Melanie Rey. 2020. Less can be more in contrastive learning. *"I Can't Believe It's Not Better!" NeurIPS 2020 workshop*.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Jeffrey Pennington, Richard Socher, and Christopher 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, Doha, Qatar. Association for Computational Linguistics.
- Jianlin Su, Jiarun Cao, Weijie Liu, and Yangyiwen Ou. 2021. Whitening sentence representations for better semantics and faster retrieval. *arXiv preprint arXiv:2103.15316*.
- Dong Wang, Ning Ding, Piji Li, and Haitao Zheng. 2021. Cline: Contrastive learning with semantic negative examples for natural language understanding. In *ACL 2021: 59th annual meeting of the Association for Computational Linguistics*, pages 2332–2342.
- Tongzhou Wang and Phillip Isola. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR.
- Gijs Wijnholds and Michael Moortgat. 2021. Sick-nl: A dataset for dutch natural language inference. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1474–1479.
- 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: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Zhuofeng Wu, Sinong Wang, Jiatao Gu, Madian Khabsa, Fei Sun, and Hao Ma. 2020. Clear: Contrastive learning for sentence representation. *arXiv preprint arXiv:2012.15466*.
- 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 *ACL 2021: 59th annual meeting of the Association for Computational Linguistics*, pages 5065–5075.
- Jure Zbontar, Li Jing, Ishan Misra, Yann Lecun, and Stéphane Deny. 2021. Barlow twins: Self-supervised learning via redundancy reduction. In *ICML 2021: 38th International Conference on Machine Learning*, pages 12310–12320.
- Yan Zhang, Ruidan He, Zuozhu Liu, Kwan Hui Lim, and Lidong Bing. 2020a. 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.

Zhong Zhang, Chongming Gao, Cong Xu, Rui Miao, Qinli Yang, and Junming Shao. 2020b. [Revisiting representation degeneration problem in language modeling](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 518–527. Association for Computational Linguistics.

A Appendix

A.1 Experiment Settings

We train our MoCoSE model using a single NVIDIA RTX3090 GPUs. Our training system runs Microsoft Windows 10 with CUDA toolkit 11.1. We use Python 3.8 and PyTorch version v1.8. We build the model with Transformers 4.4.2 (Wolf et al., 2020) and Datasets 1.8.0 (Lhoest et al., 2021) from Huggingface. We preprocess the training data according to the SimCSE to directly load the stored data in training. We compute the uniformity and alignment metrics of embedding on the STS-B dataset according to the method proposed by Wang (Wang and Isola, 2020). The STS-B dataset is also preprocessed. We use the **nlpaug** toolkit in our data augmentation experiments. For synonym replace, we use `'ContextualWordEmbsAug'` function with `'roberta-base'` as parameter. For typo, we use `'SpellingAug'` and back translation we use `'BackTranslationAug'` with parameter `'facebook/wmt19-en-de'` and paraphrase we use `'ContextualWordEmbsForSentenceAug'` with parameter `'xlnet-base-cased'`. All the parameter listing here is default value given by official.

A.2 Symmetric Two-branch Structure

We remove the online branch predictor and set the EMA decay weight to 0, i.e., make the structure and weights of the two branches identical. As shown in Figure 5, it is clear that the model is collapsing at this point. And we find that the model always works best at the very beginning, i.e., training instead hurts the performance of the model. In addition, as the training proceeds, the correlation coefficient of the model approaches 0, i.e., the prediction results have no correlation with the actual labeling. At this point, it is clear that a collapse of the model is observed. We observed such a result for several runs, so we adopted a strategy of double branching with different structures plus EMA momentum updates in our design. Subsequent experiments demonstrated that this allowed the model to avoid from collapsing.



Figure 5: Experiment on a symmetric two-branch structure with EMA decay weight set to 0.



Figure 6: Experiment after adding predictor on the online branch with EMA decay weight set to 0.

We add predictor to the online branch and set the EMA decay weight to 0. We find that the model also appears to collapse and has a dramatic oscillation in the late stage of training, as shown in Figure 6.

A.3 Pseudo-Code for Training MoCoSE

The PyTorch style pseudo-code for training MoCoSE with the negative sample queue is shown in Algorithm 1.

A.4 Distribution of Singular Values

Similar to SimCSE, we plot the distribution of singular values of MoCoSE sentence embeddings with SimCSE and BERT for comparison. As illustrated in Figure 7, our method is able to alleviate the rapid decline of singular values compared to other methods, making the curve smoother, i.e., our model is able to make the sentence embedding more isotropic.

Algorithm 1: Momentum Contrastive Sentence Embedding

Input:

\mathcal{D} : Training data set ;
 \mathcal{Q} : Negative Sample Queue;
 E_a : Embedding with random data augmentation;
 θ_o, θ_t : weights of online branch and target branch;
Optimizer : Adam optimizer
 K, K_s : Queue size, Queue size at initialisation;
 η : ema decay ema and ema scheduling strategy;
 τ Temperature parameters

Output: MoCoSE model θ_o

```
1 Initializing the queue  $\mathcal{Q}$  with size  $K_s$ ;  
2 foreach  $\mathcal{B} \in \mathcal{D}$  do  
3    $v_o, v_t \leftarrow E_a(\mathcal{B}), E_a(\mathcal{B})$  // Using data Augmentation to generate  
   different views  
4    $z_o \leftarrow \theta_o(v_o)$  //  $(N, d)$ ,  $N$  is batch size,  $d$  is dimension of sentence  
   embedding  
5    $z_t \leftarrow \theta_t(v_t)$   
6    $l_{z_o, z_t, \mathcal{Q}} \leftarrow -\log \frac{\exp(z_o \cdot z_t / \tau)}{\exp(z_o \cdot z_t / \tau) + \sum_{x \in \mathcal{Q}} \exp(z_o \cdot x / \tau)}$  // compute contrastive loss  
   using InFoNCE  
7   optimizer( $l_{z_o, z_t, \mathcal{Q}}, \theta_o$ ) // Update only the parameters of the online  
   branch according to the loss gradient;  
8    $\theta_t \leftarrow \eta * \theta_t + (1 - \eta) * \theta_o$  // Update the parameters of the target  
   branch using EMA  
9   enqueue( $\mathcal{Q}, v_t$ ) // Update the negative sample queue  $\mathcal{Q}$   
10  dequeue( $\mathcal{Q}$ )  
11 return  $\theta_o$ 
```

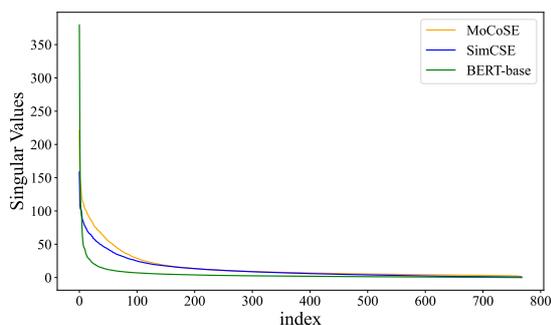


Figure 7: Singular value distributions of sentence embedding matrix from sentences in STS-B.

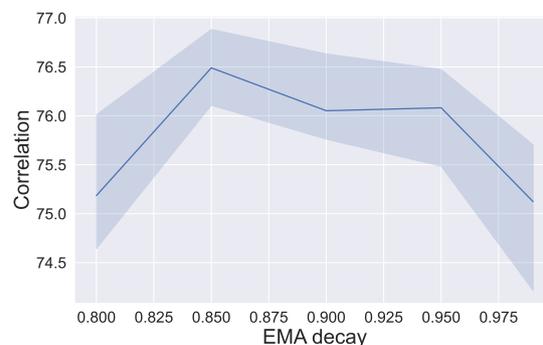


Figure 8: Effect of EMA decay weight on model performance.

A.5 Experiment Details of EMA Hyperparameters

The details of the impact caused by the EMA parameter are shown in the Figure 8. We perform this experiment with all parameters held constant except for the EMA decay weight.

A.6 Details of Different Data Augmentations

We use only dropout as a baseline for the results of data augmentations. Then, we combine dropout with other data augmentation methods and study their effects on model performance. The results are shown in Figure 9.

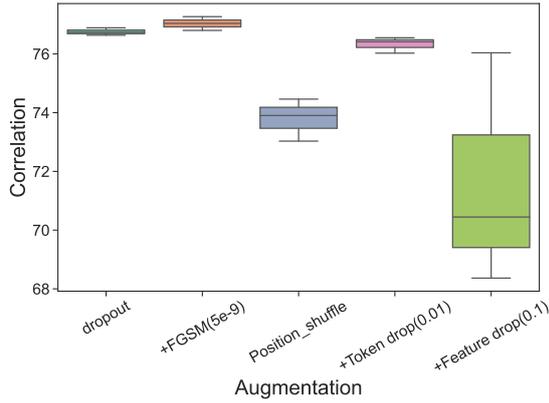


Figure 9: Impact of four additional data enhancements with dropout combinations on the model.

A.7 Experiment Details of FGSM

We test the effect of the intensity of FGSM on the model performance. We keep the other hyperparameters fixed, vary the FGSM parameters (1e-9, 5e-9, 1e-8, 5e-8). As seen in Table 9, the average results of the model are optimal when the FGSM parameter is 5e-9.

Epsilon	1e-9	5e-9	1e-8	5e-8	No
Avg.	75.61	76.64	75.39	76.62	76.26

Table 9: Different parameters of FGSM in data augmentation affect the model results.

A.8 Dimension of Sentence Embedding

In both BERT-whitening (Su et al., 2021) and MoCo (He et al., 2020), it is mentioned that the dimension of embedding can have some impact on the performance of the model. Therefore, we also changed the dimension of sentence embedding in MoCoSE and trained the model several times to observe the impact of the embedding dimension. Because of the queue structure of MoCoSE, we need to keep the dimension of negative examples consistent while changing the dimension of sentence embedding. As shown in the Figure 10, when the dimension of Embedding is low, this causes considerable damage to the performance of the model; while when the dimension rises to certain range, the performance of the model stays steady.

A.9 Details of Random Initial Queue Size

We test the influence of random initialization size of the negative queue on the model performance when queue length and batch size are fixed.

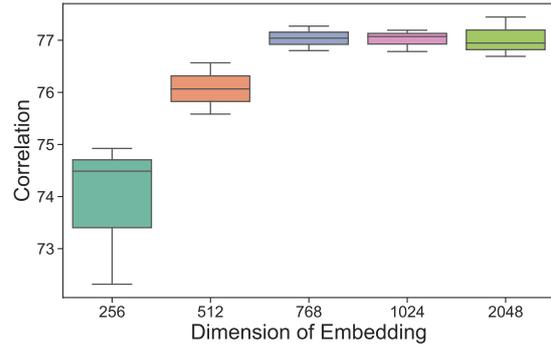


Figure 10: Impact of dimensions of the sentence embedding.

As seen in Figure 11, random initialization does have some impact on the model performance.

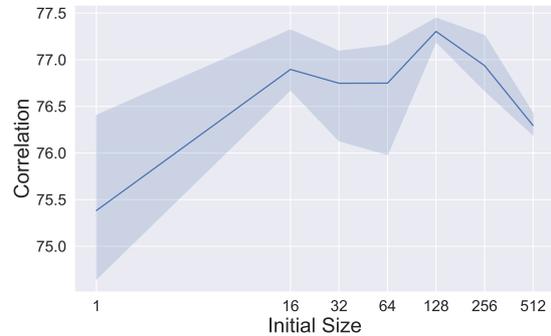


Figure 11: The effect of the initial queue size on the model results when the queue length is 512 and the batch size is 64.

A.10 Queue Size and Initial Size

We explored the effect of different combinations of initial queue sizes and queue length on the model performance. The detailed experiment results are shown in Figure 12. It can be found that model performance rely deeply on initialization queue size. Yet, too large queue size will make the model extremely unstable. This is quite different from the observation of negative sample queue in image contrastive learning.

A.11 Maximum Traceable Distance in BERT-large

We also train mocose with different batch size and queue size on BERT-large. As shown in Figure 13, we observe the best model performance in MoCoSE-BERT-large within the appropriate Maximum Traceable Distance range (around 22). Once again, this suggests that even on BERT-large, the longer queue sizes do not improve the model performance indefinitely. Which also implies that the

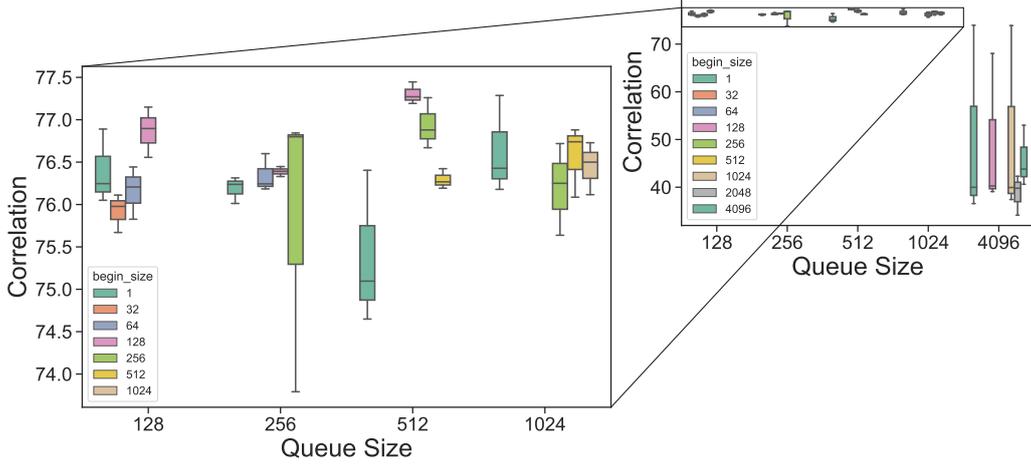


Figure 12: The impact of different initial negative sample queue sizes for different initial sizes on model performance. (left):Zoomed view. (right):Overview with different negative queue size. Results of different initial size under same queue size.

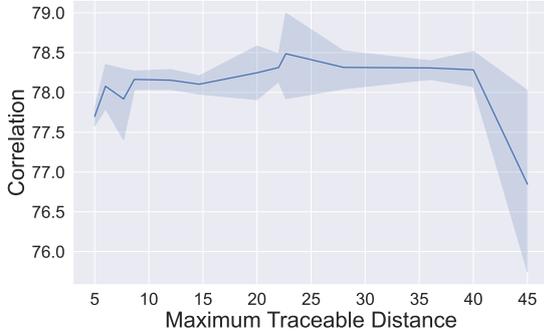


Figure 13: The relationship between MTD and correlation of MoCoSE-BERT-large. It can be seen that even at large model, peaks occur within a certain MTD range.

history information contained in the negative sample queue needs to be kept within a certain range on BERT-large as well.

A.12 Proof of Maximum Traceable Distance

Here, we prove the first term of the formula for Maximum Traceable Distance. Due to the EMA update mechanism, the weight of target branch is a weighted sum of the online weight in update history. The first term of Maximum Traceable Distance calculate the weighted sum of the historical update steps given a certain EMA decay weight η . From the principle of EMA mechanism, we can get the following equation.

$$\mathcal{S}_n = \sum_{i=0}^k (1 - \eta) \cdot \eta^i \cdot (i + 1) \quad (5)$$

\mathcal{S}_n represents the update steps between online and target branch due to the EMA mechanism. Since EMA represents the weighted sum, we need to ask for \mathcal{S}_n to get the weighted sum.

We can calculate \mathcal{S}_n as:

$$\mathcal{S}_n = (-1) * \eta^{k+1} * (k + 1) - \frac{(1 - \eta^{k+1})}{(\eta - 1)} \quad (6)$$

As k tends to infinity, the limit for \mathcal{S}_n can be calculated as following:

$$\lim_{k \rightarrow \infty} \mathcal{S}_n = \lim_{k \rightarrow \infty} \left[(-1) * \eta^{k+1} * (k + 1) - \frac{(1 - \eta^{k+1})}{(\eta - 1)} \right] \quad (7)$$

It is obvious to see that the limit of the equation 7 consists of two parts, so we calculate the limit of these two parts first.

$$\lim_{k \rightarrow \infty} (-1) * \eta^{k+1} * (k + 1) \stackrel{\eta \leq 1}{=} 0 \quad (8)$$

The limit of the first part can be calculated as 0. Next, we calculate the limit of the second part.

$$\lim_{k \rightarrow \infty} \frac{(1 - \eta^{k+1})}{(\eta - 1)} \stackrel{\eta \leq 1}{=} \frac{1}{1 - \eta} \quad (9)$$

We calculate the limit of the second part as $\frac{1}{1-\eta}$. Since the limits of both parts exist, we can obtain

the limit of \mathcal{S}_n by the law of limit operations.

$$\begin{aligned} & \lim_{k \rightarrow \infty} \mathcal{S}_n \\ &= \lim_{k \rightarrow \infty} \left[(-1) * \eta^{k+1} * (k+1) - \frac{(1 - \eta^{k+1})}{(\eta - 1)} \right] \\ &= \lim_{k \rightarrow \infty} (-1) * \eta^{k+1} * (k+1) - \lim_{k \rightarrow \infty} \frac{(1 - \eta^{k+1})}{(\eta - 1)} \\ &= \frac{1}{1 - \eta} \end{aligned} \tag{10}$$