

TNCSE: Tensor’s Norm Constraints for Unsupervised Contrastive Learning of Sentence Embeddings

Tianyu Zong¹, Bingkang Shi², Hongzhu Yi¹, Jungang Xu^{1*}

¹School of Computer Science and Technology, University of Chinese Academy of Sciences

²School of Cyber Security, University of Chinese Academy of Sciences

{zongtianyu20, shibingkang20, yihongzhu23}@mailsucas.ac.cn

xujg@ucas.ac.cn

Abstract

Unsupervised sentence embedding representation has become a hot research topic in natural language processing. As a tensor, sentence embedding has two critical properties: direction and norm. Existing works have been limited to constraining only the orientation of the samples’ representations while ignoring the features of their module lengths. To address this issue, we propose a new training objective that optimizes the training of unsupervised contrastive learning by constraining the module length features between positive samples. We combine the training objective of Tensor’s Norm Constraints with ensemble learning to propose a new Sentence Embedding representation framework, TNCSE. We evaluate seven semantic text similarity tasks, and the results show that TNCSE and derived models are the current state-of-the-art approach; in addition, we conduct extensive zero-shot evaluations, and the results show that TNCSE outperforms other baselines.

Code — <https://github.com/tianyuzong/TNCSE>

Introduction

Unsupervised sentence embedding representations have always been a research focus in natural language processing. In recent years, with the proposal of excellent pre-trained language models such as BERT(Devlin et al. 2019), RoBERTa(Liu et al. 2019), Sentence-BERT(Reimers and Gurevych 2019), and so on, many works have been done to optimize further the quality of sentence embedding representations based on them. Sentence embedding for unsupervised contrastive learning represented by SimCSE(Gao, Yao, and Chen 2021) has attracted extensive research interest, which uses InfoNCE(van den Oord, Li, and Vinyals 2018) as the loss function, the dropout function of BERT-like models as the positive sample generation method, other samples in the batch as the soft-negative samples, and cosine similarity as the semantic similarity metric function for the unsupervised training; on this basis, ESIMCSE(Wu et al. 2022b) introduces near-antonym data augmentation, InfoCSE(Wu et al. 2022a) introduces auxiliary networks and adds masked language models, ArcCSE(Zhang et al. 2022)

introduces margin for contrastive learning, EDFSE(Zong and Zhang 2023) employs the data augmentation strategy for multilingual round-trip translation(RTT) and ensemble learning, and RankCSE(Liu et al. 2023), the current SOTA method, employs knowledge distillation of naive BERT-like models with existing checkpoints.

Although the above approaches have demonstrated in different directions that fine-tuning based on InfoNCE can further optimize SimCSE, they have all neglected an essential drawback of an InfoNCE with cosine similarity as a metric, i.e., cosine similarity can only measure the angle between the embedding representations of a pair of sentences and ignores the critical attribute of module of the embedding representations. In fact, if the embedding representations of two sentences have similar angles in the high-dimensional space, but the difference in the module is significant, their semantics may still be very different, and the cosine similarity only cannot differentiate the semantic difference between these two sentences. Based on this intuitionistic phenomenon, it is necessary to introduce the difference in the module between the representations to optimize SimCSE further.

Meanwhile, ensemble learning has been shown to improve sentence embedding representations’ quality significantly. For example, EDFSE uses six SimCSE-BERTs that have been fine-tuned by multilingual RTT data augmentation as submodels and directly sums the outputs of the six submodels to obtain an ensemble model, but such a large union brings a huge inference overhead; RankCSE, on the other hand, uses the ensemble learning of two existing checkpoints as ground-truth to distill the knowledge of a naive BERT-like model. So, is it possible to introduce norm constraints on sentence embedding representations into ensemble learning, and the model can train autonomously without relying on other checkpoints of the same type?

The answer is yes, and in this paper, we propose a new and intuitive training objective: for a pair of positive samples of the embedding tensor, the cosine similarity between them should be high, and both norms should be similar. In other words, the alignment between positive samples should constrain the embedding’s angle and modulus simultaneously. In terms of model structure, the multi-encoder structure can model richer sample features from multiple spaces compared to a single encoder. Considering the inference cost, we ensemble two encoders that have been unsupervised

*Corresponding author.

fine-tuned. For the semantic representation of a sample, we perform self-supervised cosine constraints within each encoder and implement mutually supervised norm constraints between the two encoders to achieve our proposed training objective.

Then, we evaluate our proposed method on seven semantic text similarity tasks. The experimental results show that TNCSE, added by norm constraints, outperforms EDFSE on BERT-base, which is three times as large as TNCSE and becomes the new state-of-the-art method. For a fair comparison, we also do ensemble learning on other baselines and comparisons, and the results show that TNCSE is still the best. In addition, we evaluate TNCSE in a wider zero-shot on the MTEB(Muennighoff et al. 2023) list, and the results show that TNCSE also outperforms the baseline. We conduct a series of ablation experiments to evaluate the role of each part of the model, which illustrates the mechanism of TNCSE’s action intuitively. We summarize the main contributions of our work as follows:

- We are the first to intuitively propose a positive sample tensor modulus-constrained training objective based on unsupervised contrastive learning and demonstrate that it is effective.
- We combine ensemble learning with a tensor norm constraint strategy to propose a new sentence embedding representation framework, TNCSE, and generalize a series of derived models.
- We evaluate TNCSE and derived models on seven STS tasks, and the results show that TNCSE has become a new SOTA method for STS tasks. In addition, we have done a series of zero-shot evaluations, such as multilingual and cross-linguistic STS tasks, sentence classification, retrieval, and reranking, and the results show that TNCSE outperforms the baselines.

Related Work

InfoNCE(Noise Contrastive Estimation Loss(van den Oord, Li, and Vinyals 2018)) is a loss function for self-supervised learning, usually used to learn feature representations, which is based on the idea of information theory and learns the model parameters by comparing the similarity (usually cosine similarity) between positive and negative samples. There have been some significant works combining InfoNCE loss and pre-trained language models BERT(Devlin et al. 2019), RoBERTa(Liu et al. 2019), etc. to achieve training of sentence embedding representations for unsupervised contrastive learning. Both SimCSE(Gao, Yao, and Chen 2021) and ConSERT(Yan et al. 2021) use the idea of dropout to generate positive samples, and they have in common the use of cosine similarity as the only metric to discriminate between positive and negative samples. Subsequent work further optimizes the training of unsupervised sentence embeddings based on SimCSE, e.g., ESimCSE(Wu et al. 2022b) provides the idea of proxemic data augmentation while associating the idea of MoCo’s(He et al. 2020) Momentum Queue to improve the quality of SimCSE’s representations; DiffCSE(Chuang et al. 2022) introduces additional discriminator knowledge to assist the pre-trained models to mask lin-




Figure 1: The figure denotes the subtraction of the semantic representation tensor in 3D space.




Figure 2: The figure denotes the binary function image of the constraint loss of the tensor norm for independent variables k and t .

guistic modeling. SNCSE(Wang and Dou 2023) employs a contrastive learning strategy that combines soft and negative samples with a bi-directional margin loss; and InfoCSE(Wu et al. 2022a) introduces an additional network for constructing a masked language model; EDFSE(Zong and Zhang 2023) constructs huge ensemble models by training multiple encoders through RTT data augmentation; the current SOTA approach, RankCSE(Liu et al. 2023), uses dual-teacher ensemble learning and distillation to train encoder. The commonality among these existing works is that they all introduce the InfoNCE loss, which employs the cosine angle between embeddings of individual samples to determine similarity. The crucial factor of the norm of the embedding tensor has been overlooked, so we propose an unsupervised ensemble learning framework with embedding representation norm constraints, to enhance the model’s ability to discriminate the positive and negative samples.

Method

In this section, we start by presenting the norm constraints imposed on semantic representation tensors. Subsequently,

we incorporate ensemble learning strategies, leading to the elaboration of our proposed framework, the Tensor Norm Constrained Semantic Embedding (TNCSE).

Norm Constraints on Representations

Traditional unsupervised text embedding representation training typically uses the cosine similarity between representation tensors as the primary metric for determining the degree of similarity between sentence pairs, but this only considers the “direction” of the tensor properties and ignores the “magnitude”. Wang and Isola (2020) believes that semantic representations can be projected onto a hypersphere. The distribution of positive samples should be aligned as much as possible. The soft negative samples should be distributed uniformly on the hypersphere. However, to distinguish semantic similarity among many texts in a fine-grained way, we should consider both the direction and the magnitude of the representation tensor. For a semantic representation of a pair of positive samples, intuitively, the angle between them should be as slight as possible, and the magnitudes should be identical. Therefore, in our approach, we model the 2-paradigm (denoted hereafter as the **norm**) of the semantic representation tensor as its “magnitude”, to perform constraints, and propose a loss based on the norm constraints as shown in Eq. 1:

$$l_{TN}(h, h^+) = \frac{\|h - h^+\|}{\|h\| + \|h^+\|}. \quad (1)$$

In Eq. 1, h and h^+ denote the embedding representations of a sample and its positive samples, $\|\cdot\|$ denotes the norm of tensor, and we come to demonstrate the rationality of this loss. For visualization, we report this process in a 3-dimensional space. As shown in Fig. 1, the two semantic representation tensors h and h^+ of positive samples have difference tensor $h - h^+$. These three tensors construct a triangle with sides $\|h\|$, $\|h^+\|$, and $\|h - h^+\|$, and let the angle between h and h^+ be γ . Then, according to the cosine theorem, we have:

$$l_{TN}(h, h^+) = \frac{\sqrt{\|h\|^2 + \|h^+\|^2 - 2\|h\|\|h^+\|\cos\gamma}}{\|h\| + \|h^+\|}. \quad (2)$$

Since the norm of semantic representation are all nonzero, we make $\|h^+\| = k \cdot \|h\|$, $t = \cos\gamma$, where $k \in (0, +\infty)$ and $t \in [-1, 1]$. Bringing above into the Eq. 2, there is:

$$l_{TN}(k, t) = \frac{\sqrt{1 + k^2 - 2 \cdot kt}}{1 + k}. \quad (3)$$

We plot the image of the binary function l_{TN} for the independent variables k and t over part of the domain of definition, as shown in Fig. 2. l_{TN} takes a minimal value when k and t are equal to 1, respectively, in other word, when $\cos\gamma = 1$ and $\|h^+\| = \|h\|$. This result aligns with our intuition, our proposed l_{TN} can simultaneously constrain the two metrics of cosine similarity and modulus length between positive samples of the semantic representation tensor.

Model Structure Design

This subsection describes our proposed sentence embedding representation framework TNCSE based on tensor norm constraints, as shown in Fig. 3b. The traditional large-scale ensemble model EDFSE (Fig. 3a) uses six encoders, each of which has been pretrained with multilingual RTT data augmentation and unsupervised SimCSE. This ensemble strategy endows the model with the advantage of diverging the embeddings from different encoders as much as possible, thereby increasing the “intrinsic rank” of the joint representation, leading to enhanced capability in discriminating between semantically similar sentences. We follow EDFSE’s data augmentation approach, but only employ two encoders that have been data augmented and fine-tuned by unsupervised SimCSE for reducing EDFSE’s inference overhead. For each of the two encoders of TNCSE-BERT and TNCSE-RoBERTa, we randomly chose four target languages for RTT Data Augmentation (EN-Target-EN). Each of these encoders ends up with its pooler layer. We denote these two encoders as Encoder I and Encoder II, respectively.

When the TNCSE is training, due to the dropout in the encoder, each sentence is fed into Encoder I and Encoder II twice to obtain positive samples, respectively. The last hidden state of both encoders contains two similar embeddings, and these four representations are denoted as h_I^L , h_I^{L+} , h_{II}^L , and h_{II}^{L+} , respectively. Since both are representations of the same sentence sample, two of these four hidden states are positive samples of each other, so we set the respective InfoNCE and the interaction-constrained InfoNCE (ICNCE), InfoNCE is denoted as Eq. 4:

$$L_{NCE}(h_i, h_i^+) = -\log \frac{e^{\frac{\text{sim}(h_i, h_i^+)}{\tau}}}{e^{\frac{\text{sim}(h_i, h_i^+)}{\tau}} + \sum e^{\frac{\text{sim}(h_i, h_j^-)}{\tau}}}. \quad (4)$$

Here, $i \in \{I, II\}$ indicates from Encoder I or Encoder II. h_i and h_i^+ are mutually positive samples, and h_j^- is the other samples within the current *batch size* as soft-negative samples. We follow the SimCSE setting with $\tau = 0.05$; $\text{sim}(\cdot)$ is the cosine similarity. ICNCE is then denoted as Eq. 5:

$$L_{ICNCE} = L_{NCE}(h_I^L, h_{II}^L). \quad (5)$$

InfoNCE serves to continue training and maintain the general direction of the encoder’s gradient update, while ICNCE is designed to precondition subsequent tensor norm loss.

Due to the norms of the last hidden states in BERT-like models converging to a uniform value, we cannot effectively apply tensor norm constraints to the last hidden states of positive sample pairs. However, we notice that the last pooler layer in a BERT-like model is a single feedforward neural network, which can be endowed with normative features by projecting the last hidden state onto a space of the same dimension. Accordingly, we naturally employ the output of the pooler layer for tensor norm constraints. Moreover, we find that the cosine similarity of the last hidden state of the positive sample pairs can be used as a nondeterministic prompt for the l_{TN} . Thus, we amend the l_{TN} to be Eq. 6:

Model	STS12	STS13	STS14	STS15	STS16	STSB	SICKR	Avg.
BERT-base								
SimCSE(Gao, Yao, and Chen 2021)*	68.40	82.41	74.38	80.91	78.56	76.85	72.23	76.25
DiffCSE(Chuang et al. 2022)*	72.28	84.43	76.47	83.90	80.54	80.59	71.23	78.49
ESimCSE(Wu et al. 2022b)*	73.40	83.27	77.25	82.66	78.81	80.17	72.30	78.27
ArcCSE(Zhang et al. 2022)*	72.08	84.27	76.25	82.32	79.54	79.92	72.39	78.11
InfoCSE(Wu et al. 2022a)*	70.53	84.59	76.40	85.10	81.95	82.00	71.37	78.85
PromptBERT(Jiang et al. 2022)*	71.56	84.58	76.98	84.47	80.60	81.60	69.87	78.54
SNCSE(Wang and Dou 2023)*	70.67	84.79	76.99	83.69	80.51	81.35	74.77	78.97
WhitenedCSE(Zhuo et al. 2023)*	74.03	84.90	76.40	83.40	80.23	81.14	71.33	78.78
EDFSE(Zong and Zhang 2023)*	74.48	83.14	76.39	84.45	80.02	81.97	72.83	79.04
TNCSE	75.52	83.91	77.57	84.97	80.42	81.72	72.97	79.58
EDFSE D(Zong and Zhang 2023)*	74.50	83.61	76.24	84.02	80.44	81.94	74.16	79.27
TNCSE D	75.42	84.64	77.62	74.92	80.50	81.79	73.52	79.77
RankCSE(Liu et al. 2023)♣	74.61	85.70	78.09	84.64	81.36	81.82	74.51	80.10
RankCSE+UC	73.29	85.90	78.16	85.90	82.52	83.13	73.36	80.32
TNCSE+UC	75.79	85.27	78.67	85.99	82.01	83.16	73.01	80.56
RankCSE+UC D	72.99	85.72	77.73	84.93	81.86	82.43	74.35	80.00
TNCSE+UC D	75.95	85.31	78.50	85.69	81.86	83.03	73.89	80.60
RoBERTa-base								
SimCSE(Gao, Yao, and Chen 2021)*	70.16	81.77	73.24	81.36	80.65	80.22	68.56	76.57
DiffCSE(Chuang et al. 2022)*	70.05	83.43	75.49	82.81	82.12	82.38	71.19	78.21
ESimCSE(Wu et al. 2022b)*	69.90	82.50	74.68	83.19	80.30	80.99	70.54	77.44
PromptBERT(2022)*	73.94	84.74	77.28	84.99	81.74	81.88	69.50	79.15
SNCSE(Wang and Dou 2023)*	70.62	84.42	77.24	84.85	81.49	83.07	72.92	79.23
WhitenedCSE(Zhuo et al. 2023)*	70.73	83.77	75.56	81.85	83.25	81.43	70.96	78.22
IS-CSE(He et al. 2023)*	71.39	82.58	74.36	82.75	81.61	81.40	69.99	77.73
EDFSE(Zong and Zhang 2023)	72.67	83.00	75.69	84.07	82.01	82.53	71.92	78.84
TNCSE	74.11	84.00	76.06	84.80	81.61	82.68	73.47	79.53
EDFSE D(Zong and Zhang 2023)	71.04	81.08	77.04	83.08	81.96	82.36	74.54	78.73
TNCSE D	74.56	84.74	76.30	84.89	81.70	83.01	74.18	79.91
RankCSE(Liu et al. 2023)♣	69.09	81.15	73.62	81.31	81.43	81.22	70.08	76.84
RankCSE+UC	74.18	84.06	77.72	83.26	79.81	81.25	72.58	78.98
TNCSE+UC	74.52	85.26	77.63	85.85	82.62	83.65	73.35	80.41
RankCSE+UC D	68.55	82.23	73.61	81.28	81.28	80.98	71.01	76.99
TNCSE+UC D	74.14	83.86	76.09	84.07	81.59	82.90	73.55	79.46

Table 1: This table reports the results of the TNCSE and baseline evaluation on the seven STS tasks. * denotes results derived from the original paper. indicates a fusion with existing unsupervised checkpoint knowledge. indicates the best result on the main metric **Avg.** Since RankCSE(Liu et al. 2023) has not officially open-sourced any code or checkpoints, ♣ denotes the result of a third-party open-source code replication¹. **D** denotes distillation to a single encoder.

model containing two encoders, we follow EDFSE’s distillation approach and distill the knowledge from TNCSE to a naive BERT/RoBERTa, and obtain TNCSE-D, which also outperforms baselines as a single-encoder model. Compared to a traditional ensemble model like EDFSE, we propose an efficient training objective that makes the inference overhead of TNCSE only one-third of EDFSE-BERT, but with better results on the 7 STS task.

Since RankCSE relies on the knowledge of the existing powerful checkpoints, SimCSE-base and SimCSE-large, to train the encoder in a knowledge distillation approach, and RankCSE is not trained independently, we consider that it is inappropriate to compare RankCSE with TNCSE and other baselines directly. We ensemble and distill RankCSE

and TNCSE with unsupervised checkpoint InfoCSE in the same way, respectively, and get the ensemble model (UC) and distillation model (UC D) to evaluate; in other words, we prove that TNCSE is better than RankCSE by proving that TNCSE+InfoCSE is better than RankCSE+InfoCSE. In summary, TNCSE and derived models are the current SOTA methods for unsupervised STS tasks. Due to computational resource constraints and the unavailability of some test sets, we randomly select 30 sentence classification tasks and 30 retrieval tasks, 29 multilingual and cross-linguistic STS tasks (STS17, STS22), and all 9 reranking tasks for which test sets are available on the MTEB, and report the experimental results in Table 2, where the overall results show that the TNCSE and derived models outperform the baseline

Tasks(Avg.)	SimCSE	ESimCSE	DiffCSE	InfoCSE	SNCSE	RankCSE	TNCSE	+D	+UC D
STS17(11 Acc)	34.22	35.64	31.61	35.95	22.64	37.59	36.32	35.41	37.73
STS22(18 Acc)	32.78	36.75	34.33	26.48	23.61	37.01	39.34	38.95	39.26
Classification(30 Acc)	58.15	58.26	57.44	58.15	58.05	58.15	58.38	58.59	58.84
Reranking(9 Map)	36.78	38.12	37.13	38.36	31.95	37.11	38.15	38.93	38.29
Retrieval(30 Map@10)	18.98	21.13	19.09	21.54	16.88	18.33	22.10	23.16	22.14

Table 2: This table reports the results of the zero-shot evaluation of TNCSE and baseline on MTEB, with the number of tasks and metric indicated in parentheses. Details of the results are reported in Appendix IV.

Model(Dual Encoder)	SimCSE	ESimCSE	DiffCSE	InfoCSE	RankCSE	TNCSE(Untrained)	TNCSE
7 STS Avg.	77.97	78.51	77.89	77.05	78.22	78.27	79.58

Table 3: We reports the results of ensemble learning done by a set of encoders obtained by adding unlabelled SICKR dataset and the 2 RTT augmented datasets to which each baseline is trained(Default hyperparameters) and compares them with TNCSE and the two encoders untrained by TNCSE. The inconsistency between the dual SimCSE and TNCSE-Untrained results is due to the fine-tuning of the learning rate we have done during pre-training, but the dependence on external knowledge is the same.




Figure 4: The figure reports the results of the significance test. We specify the random seeds are 1 to 5, other hyperparameters are defaulted, training set is uniformly Wiki1M and unlabelled SICKR.

approach.

Significance Test Since the BERT-like model is single encoder and training is susceptible to random seeds, we design TNCSE with dual encoder structure, which is more stable than single encoder. We report the significance test of TNCSE and baselines in Figure 4.

Ablation Studies and Discussion

Alignment and Uniformity

Wang and Isola (2020) proposes that two critical metrics for measuring sentence embedding representations are unifor-




Figure 5: The bar and line graphs separately represent the norm mean of the Last hidden state and the model’s performance on the STS-B validation set for different LayerNorm settings.

metry and alignment, respectively, which we have already introduced in the model section, both of which are as small as possible. We report the scores of these two metrics for TNCSE and baselines in Figure 6. Due to the introduction of the tensor-norm alignment strategy in TNCSE, the loss of TNCSE concerning conventional angle alignment is superior. Because TNCSE unites ensemble learning, uniformity is also superior. We report on the detailed formulation of alignment and uniformity in Appendix III.

Loss Choice	7 STS Avg.
None (Dual Encoder Untrained)	78.27
L_{NCE}	78.42
L_{ICNCE}	78.71
L_{ICTN}	78.38
$L_{NCE} + L_{ICNCE}$	78.71
$L_{NCE} + L_{ICTN}$	79.53
$L_{ICNCE} + L_{ICTN}$	79.43
$L_{NCE} + L_{ICNCE} + L_{ICTN}$ (Ours)	79.58

Table 4: This table combines each of the loss functions to explore the contribution of each to model training. **None** denotes a direct ensemble of two SimCSE-trained encoders. All experiments use CLS pooling method.

Model	SimCSE	TNCSE-Single
STS-B testset	77.75	78.37

Table 5: This table compares the training results of SimCSE and TNCSE with only a single encoder under the same training set.

Ablation of The Loss Function

Since our loss function contains three terms, InfoNCE, ICNCE, and ICTN, this subsection investigates the performance gain from each term and analyses the reasons. The results of the ablation experiments are reported in Table 4.

L_{NCE} Only Since TNCSE employs a dual encoder structure, the sentence embedding is modeled through ensemble learning. If only unsupervised training of InfoNCE is performed for each encoder, which does not directly improve the ensemble learning, and result is improved insignificantly.

L_{ICNCE} Only ICNCE is a boosted version of InfoNCE. ICNCE optimizes the tensor angle of positive and negative samples between two encoders, which works in ensemble learning. However, since this does not introduce a new training objective, it still belongs to the incremental training of InfoNCE, so the improvement is limited.

L_{ICTN} Only Since L_{ICTN} (Eq. 7) accomplishes norm constraint on the output of the encoder’s pooler layer, but in the pre-training of the encoder, only the last hidden state is taken out to do the self-supervised training, and it does not go through the pooler layer, there is a margin between the last hidden state and the pooler output. We will analyze in the discussion subsection that the last hidden state cannot directly use norm loss because it lacks norm features, so using only the pooler output to constrain the tensor norm in ensemble learning is almost ineffective with the CLS pooling method.

$L_{NCE} + L_{ICNCE}$ Under the joint constraints of L_{NCE} and L_{ICNCE} , the two encoders perform self-supervised continuation training and, at the same time, mutually supervise each other’s hidden states, which facilitates fine-grained




Figure 6: The alignment and uniformity metrics for TNCSE and baselines; the better the distribution is down the left. Darker the circle color, better the model performs on 7 STS.

representations of samples by ensemble learning, but still falls under the tensor’s angle constraints. Therefore, the improvement is minor.

$L_{NCE} + L_{ICTN}$ and $L_{ICNCE} + L_{ICTN}$ The combination of the tensor’s norm constraints and either of the angle constraints resulted in a significant improvement in training. This is due to the two properties of the trained tensor, allowing the two encoders to more clearly distinguish between two soft negative samples with similar representations after ensemble learning, improving the model performance.

Ablation of Ensemble Learning

Since TNCSE is dual-encoder structured, although we have distilled its knowledge to a single encoder, we still ablate its structure. We keep only one naive encoder, BERT-base, whose pooler layer outputs a pair of positive samples to accomplish the constraints of L_{TN} (Eq. 6) and the last hidden state to accomplish the constraints of self-supervised InfoNCE. In other words, this structure only increases the L_{TN} loss relative to Unsup-SimCSE, and we retrained unsupervised SimCSE in the same experimental setting with the same training set for both Wiki1M and unlabelled SICKR. We report the model’s performance on the STS-B test set in Table 5. Compared with the Unsup-SimCSE, which includes the same baseline knowledge, the model with the added L_{TN} loss is more effective than the unsupervised SimCSE. In addition, we put several representative baselines through the same two different RTT pre-training to get a dual encoder, and then compare it with TNCSE. The results are reported in Table 3.

Impact of Unlabelled SICKR

We have found that when reproducing SimCSE, ESimCSE, and DiffCSE, we cannot reproduce the results reported in

Datasets	Wiki1M	+Unlabelled SICKR
SimCSE	74.3	75.9(+1.6)
ESimCSE	75.8	76.7(+0.9)
DiffCSE	75.2	78.0(+2.8)
InfoCSE	75.8	77.4(+1.6)
RankCSE	77.2	77.5(+0.3)
TNCSE	79.3	79.6(+0.3)

Table 6: The impact of adding unlabelled SICKR datasets on model training.

the paper using the official open-source code and default hyperparameters. For example, our reproduction of SimCSE is only about 74%, which is far from the reported 76.25%, so we can only improve the reproduction level by adding an unsupervised dataset, based on which we roughly improve the results of SimCSE to about 76% in order to be fair enough to conduct the subsequent experiments. In Table 6, we report the results we obtained by training the model with the default hyperparameters and the results by adding the SICKR dataset. In addition, we report the results of training TNCSE using the original Wiki1M to demonstrate that adding the unlabelled dataset does not significantly improve the model performance.

Discussion: Why Choose Pooler Output Over Last Hidden State?

In our framework, pooler output is chosen for norm constraints; in this subsection, we discuss the reason. BERT and RoBERTa, pre-trained models based on the structure of the Transformer encoder, include two LayerNorms in each of their encoder layers, which are located after the attention layer and after the FNN, respectively, which makes their last hidden state lose the norm feature. No matter the text input to SimCSE-BERT, their last hidden state’s norms are almost always distributed between 14 and 16, which cannot do norm constraints. We remove the last few LayerNorms in the encoder layer of SimCSE-BERT to make the last hidden state get the norm features and apply the Eq. 8 on the last hidden state. However, the training results we get are instead worse, and the more LayerNorms we remove, the worse the result is. We consider that this destroys the pre-training information of SimCSE-BERT, which is unfavorable to the TNCSE training. We report this finding in Figure 5. We first randomly select 100 sentences from Wiki1M. Then we remove the last several layers of LayerNorm of the encoder, and report the average norm of the last hidden state, as shown in the bar graph. We train two encoders with the LayerNorm removed using the TNCSE method and report the performance on the STS-B test set. As can be seen from the line graphs, the model almost loses its semantic discriminative ability when all the last six layers of LayerNorm are removed. Overall, the more LayerNorms are preserved, the better the model works, but both are not as good as our proposed scheme employing Pooler output for tensor constraints. The above experiment demonstrates the ingenuity of our design.

Conclusion

In this work, we propose a new training objective of unsupervised sentence embedding, which can constrain the module length of sentence embedding representations between positive samples; we model it as a tensor’s norm and jointly train with ensemble learning. Based on these methods, we propose a new unsupervised sentence embedding representation framework, TNCSE, which becomes a new SOTA method on seven STS tasks. Moreover, nearly one hundred zero-shot tasks are used to evaluate, and the results show that TNCSE and its derived models outperform other baselines. In addition, we conduct a series of ablation experiments to explore the effects of individual components.

Acknowledgments

We thank the anonymous reviewers and area chair for their contributions to this work.

References

- Agirre, E.; Banea, C.; Cardie, C.; Cer, D.; Diab, M.; Gonzalez-Agirre, A.; Guo, W.; Lopez-Gazpio, I.; Maritxalar, M.; Mihalcea, R.; Rigau, G.; Uria, L.; and Wiebe, J. 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)*, 252–263. Denver, Colorado: Association for Computational Linguistics.
- Agirre, E.; Banea, C.; Cardie, C.; Cer, D.; Diab, M.; Gonzalez-Agirre, A.; Guo, W.; Mihalcea, R.; Rigau, G.; and Wiebe, J. 2014. SemEval-2014 Task 10: Multilingual Semantic Textual Similarity. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, 81–91. Dublin, Ireland: Association for Computational Linguistics.
- Agirre, E.; Banea, C.; Cer, D.; Diab, M.; Gonzalez-Agirre, A.; Mihalcea, R.; Rigau, G.; and Wiebe, J. 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)*, 497–511. San Diego, California: Association for Computational Linguistics.
- Agirre, E.; Cer, D.; Diab, M.; and Gonzalez-Agirre, A. 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)*, 385–393. Montréal, Canada: Association for Computational Linguistics.
- Agirre, E.; Cer, D.; Diab, M.; Gonzalez-Agirre, A.; and Guo, W. 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*, 32–43. Atlanta, Georgia, USA: Association for Computational Linguistics.

- Cer, D.; Diab, M.; Agirre, E.; Lopez-Gazpio, I.; and Specia, L. 2017a. SemEval-2017 Task 1: Semantic Textual Similarity Multilingual and Crosslingual Focused Evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, 1–14. Vancouver, Canada: Association for Computational Linguistics.
- Cer, D. M.; Diab, M. T.; Agirre, E.; Lopez-Gazpio, I.; and Specia, L. 2017b. SemEval-2017 Task 1: Semantic Textual Similarity Multilingual and Crosslingual Focused Evaluation. In Bethard, S.; Carpuat, M.; Apidianaki, M.; Mohammad, S. M.; Cer, D. M.; and Jurgens, D., eds., *Proceedings of the 11th International Workshop on Semantic Evaluation, SemEval@ACL 2017, Vancouver, Canada, August 3-4, 2017*, 1–14. Association for Computational Linguistics.
- Chen, X.; Zeynali, A.; Camargo, C.; Flöck, F.; Gaffney, D.; Grabowicz, P.; Hale, S.; Jurgens, D.; and Samory, M. 2022. SemEval-2022 Task 8: Multilingual news article similarity. In Emerson, G.; Schluter, N.; Stanovsky, G.; Kumar, R.; Palmer, A.; Schneider, N.; Singh, S.; and Ratan, S., eds., *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, 1094–1106. Seattle, United States: Association for Computational Linguistics.
- Chuang, Y.; Dangovski, R.; Luo, H.; Zhang, Y.; Chang, S.; Soljagic, M.; Li, S.; Yih, S.; Kim, Y.; and Glass, J. R. 2022. DiffCSE: Difference-based Contrastive Learning for Sentence Embeddings. In Carpuat, M.; de Marneffe, M.; and Ruíz, I. V. M., eds., *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, 4207–4218. Association for Computational Linguistics.
- Conneau, A.; and Kiela, D. 2018. SentEval: An Evaluation Toolkit for Universal Sentence Representations. In Calzolari, N.; Choukri, K.; Cieri, C.; Declerck, T.; Goggi, S.; Hasida, K.; Isahara, H.; Maegaard, B.; Mariani, J.; Mazo, H.; Moreno, A.; Odijk, J.; Piperidis, S.; and Tokunaga, T., eds., *Proceedings of the Eleventh International Conference on Language Resources and Evaluation, LREC 2018, Miyazaki, Japan, May 7-12, 2018*. European Language Resources Association (ELRA).
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, 4171–4186. Association for Computational Linguistics.
- Gao, T.; Yao, X.; and Chen, D. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In Moens, M.; Huang, X.; Specia, L.; and Yih, S. W., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, 6894–6910. Association for Computational Linguistics.
- He, H.; Zhang, J.; Lan, Z.; and Zhang, Y. 2023. Instance Smoothed Contrastive Learning for Unsupervised Sentence Embedding. In Williams, B.; Chen, Y.; and Neville, J., eds., *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, 12863–12871. AAAI Press.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. B. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, 9726–9735. Computer Vision Foundation / IEEE.
- Jiang, T.; Jiao, J.; Huang, S.; Zhang, Z.; Wang, D.; Zhuang, F.; Wei, F.; Huang, H.; Deng, D.; and Zhang, Q. 2022. PromptBERT: Improving BERT Sentence Embeddings with Prompts. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, 8826–8837. Association for Computational Linguistics.
- Liu, J.; Liu, J.; Wang, Q.; Wang, J.; Wu, W.; Xian, Y.; Zhao, D.; Chen, K.; and Yan, R. 2023. RankCSE: Unsupervised Sentence Representations Learning via Learning to Rank. In Rogers, A.; Boyd-Graber, J. L.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, 13785–13802. Association for Computational Linguistics.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR*, abs/1907.11692.
- Marelli, M.; Menini, S.; Baroni, M.; Bentivogli, L.; Bernardi, R.; and Zamparelli, R. 2014. A SICK cure for the evaluation of compositional distributional semantic models. In Calzolari, N.; Choukri, K.; Declerck, T.; Loftsson, H.; Maegaard, B.; Mariani, J.; Moreno, A.; Odijk, J.; and Piperidis, S., eds., *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014, Reykjavik, Iceland, May 26-31, 2014*, 216–223. European Language Resources Association (ELRA).
- Muennighoff, N.; Tazi, N.; Magne, L.; and Reimers, N. 2023. MTEB: Massive Text Embedding Benchmark. In Vlachos, A.; and Augenstein, I., eds., *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2023, Dubrovnik, Croatia, May 2-6, 2023*, 2006–2029. Association for Computational Linguistics.
- Reimers, N.; and Gurevych, I. 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)*, 3982–3992. Hong Kong, China: Association for Computational Linguistics.

- van den Oord, A.; Li, Y.; and Vinyals, O. 2018. Representation Learning with Contrastive Predictive Coding. *CoRR*, abs/1807.03748.
- Wang, H.; and Dou, Y. 2023. SNCSE: Contrastive Learning for Unsupervised Sentence Embedding with Soft Negative Samples. In Huang, D.; Premaratne, P.; Jin, B.; Qu, B.; Jo, K.; and Hussain, A., eds., *Advanced Intelligent Computing Technology and Applications - 19th International Conference, ICIC 2023, Zhengzhou, China, August 10-13, 2023, Proceedings, Part IV*, volume 14089 of *Lecture Notes in Computer Science*, 419–431. Springer.
- Wang, T.; and Isola, P. 2020. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 9929–9939. PMLR.
- Wu, X.; Gao, C.; Lin, Z.; Han, J.; Wang, Z.; and Hu, S. 2022a. InfoCSE: Information-aggregated Contrastive Learning of Sentence Embeddings. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Findings of the Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, 3060–3070. Association for Computational Linguistics.
- Wu, X.; Gao, C.; Zang, L.; Han, J.; Wang, Z.; and Hu, S. 2022b. ESIMCSE: Enhanced Sample Building Method for Contrastive Learning of Unsupervised Sentence Embedding. In Calzolari, N.; Huang, C.; Kim, H.; Pustejovsky, J.; Wanner, L.; Choi, K.; Ryu, P.; Chen, H.; Donatelli, L.; Ji, H.; Kurohashi, S.; Paggio, P.; Xue, N.; Kim, S.; Hahm, Y.; He, Z.; Lee, T. K.; Santus, E.; Bond, F.; and Na, S., eds., *Proceedings of the 29th International Conference on Computational Linguistics, COLING 2022, Gyeongju, Republic of Korea, October 12-17, 2022*, 3898–3907. International Committee on Computational Linguistics.
- Yan, Y.; Li, R.; Wang, S.; Zhang, F.; Wu, W.; and Xu, W. 2021. ConSERT: A Contrastive Framework for Self-Supervised Sentence Representation Transfer. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, 5065–5075. Association for Computational Linguistics.
- Zhang, Y.; Zhu, H.; Wang, Y.; Xu, N.; Li, X.; and Zhao, B. 2022. 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)*, 4892–4903. Dublin, Ireland: Association for Computational Linguistics.
- Zhuo, W.; Sun, Y.; Wang, X.; Zhu, L.; and Yang, Y. 2023. WhitenedCSE: Whitening-based Contrastive Learning of Sentence Embeddings. In Rogers, A.; Boyd-Graber, J. L.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, 12135–12148. Association for Computational Linguistics.
- Zong, T.; and Zhang, L. 2023. An Ensemble Distillation Framework for Sentence Embeddings with Multilingual Round-Trip Translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 14074–14082.