

InfoCSE: Information-aggregated Contrastive Learning of Sentence Embeddings

Xing Wu^{1,2,3}, Chaochen Gao^{1,2*}, Zijia Lin³, Jizhong Han¹, Zhongyuan Wang³, Songlin Hu^{1,2†}

¹Institute of Information Engineering, Chinese Academy of Sciences

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

³Kuaishou Technology

{gaochaochen,zangliangjun,hanjizhong,husonglin}@iie.ac.cn

{wuxing,wangzhongyuan}@kuaishou.com, linzijia07@tsinghua.org.cn

Abstract

Contrastive learning has been extensively studied in sentence embedding learning, which assumes that the embeddings of different views of the same sentence are closer. The constraint brought by this assumption is weak, and a good sentence representation should also be able to reconstruct the original sentence fragments. Therefore, this paper proposes an **information-aggregated contrastive learning** framework for learning unsupervised sentence embeddings, termed InfoCSE. InfoCSE forces the representation of [CLS] positions to aggregate denser sentence information by introducing an additional Masked language model task and a well-designed network. We evaluate the proposed InfoCSE on several benchmark datasets w.r.t the semantic text similarity (STS) task. Experimental results show that InfoCSE outperforms SimCSE by an average Spearman correlation of 2.60% on BERT-base, and 1.77% on BERT-large, achieving state-of-the-art results among unsupervised sentence representation learning methods. Our code are available at github.com/caskcsg/sentemb/tree/main/InfoCSE.

1 Introduction

Sentence embeddings aim to capture rich semantic information to be applied in many downstream tasks (Zhang et al.; Wu et al.; Liu et al., 2021). Recently, researchers have started to use contrastive learning to learn better unsupervised sentence embeddings (Gao et al., 2021; Yan et al., 2021; Wu et al., 2021a; Zhou et al., 2022; Wu et al., 2021b; Chuang et al., 2022). Contrastive learning methods assume that effective sentence embeddings should bring similar sentences closer and dissimilar sentences farther. These methods use various data augmentation methods to randomly generate different

Model	STS-B
SimCSE-BERT _{base}	86.2
w/ MLM	
$\lambda = 0.01$	85.7
$\lambda = 0.1$	85.7
$\lambda = 1$	85.1

Table 1: Table from SimCSE (Gao et al., 2021). The masked language model (MLM) objective brings a consistent drop to the SimCSE model in semantic textual similarity tasks. “w/” means “with”, λ is the balance hyperparameter for MLM loss.

views for each sentence and constrain one sentence semantically to be more similar to its augmented counterpart than any other sentence. SimCSE (Gao et al., 2021) is the representative work of contrastive sentence embedding, which uses dropout acts as minimal data augmentation. SimCSE encodes the same sentence twice into embeddings to obtain “positive pairs” and takes other sentence embeddings in the same mini-batch as “negatives”. There have been many improvements to SimCSE, including enhancing positive and negative sample building methods (Wu et al., 2021a), alleviating the influence of improper mini-batch negatives (Zhou et al., 2022), and learning sentence representations that are aware of direct surface-level augmentations (Chuang et al., 2022).

Although SimCSE and its variants have achieved good results and can learn sentence embeddings that can distinguish different sentences, this is not enough to indicate sentence embeddings already contain the semantics of sentences well. If a sentence embedding is sufficiently equivalent to the semantics of the sentence, it should also be able to reconstruct the original sentence to a large extent (Lu et al., 2021). However, as shown in Table 1 from (Gao et al., 2021), experiments show that “the masked language model objective brings a consistent drop in semantic textual similarity tasks” in SimCSE. This is due to the gradients of the MLM

†The first two authors contribute equally.

*Corresponding author.

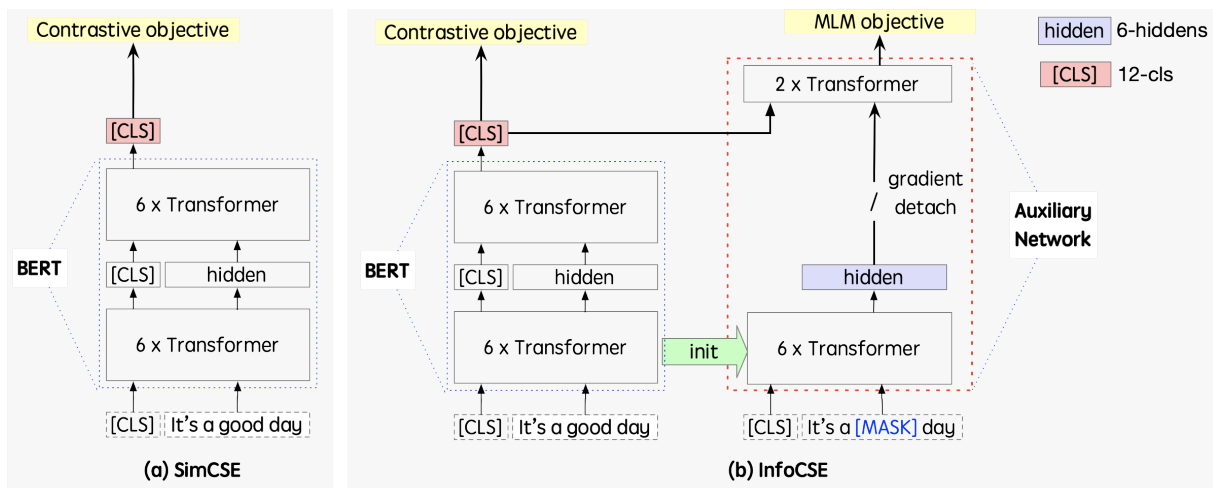


Figure 1: Comparison of InfoCSE and SimCSE structures. SimCSE learns sentence representations through contrastive learning on the [CLS] output embeddings of the BERT model. In addition to contrastive learning, InfoCSE designs an auxiliary network for sentence reconstruction with the [CLS] embeddings, enabling to learn better sentence representations.

objective optimization will easily over-update the parameters of the encoder network, thus causing disturbance to the contrastive learning task. Therefore, it is not an easy job to incorporate sentence reconstruction task in contrastive sentence embedding learning.

To improve contrastive sentence embedding learning with the sentence reconstruction task, we propose an information-aggregated contrastive learning framework, termed InfoCSE. The previous work (Gao et al., 2021) shares the encoder when jointly optimizing and contrastive learning objective and MLM objective. Unlike (Gao et al., 2021), we design an auxiliary network to optimize the MLM objective, as shown in Figure 1-(b). The auxiliary network is an 8-layer transformer network consisting of a frozen lower six layers and an additional two layers. The auxiliary network takes two inputs. One is the sentence embedding of the original text, and the other is the masked text. The sentence embedding is a vector of the [CLS] position encoded by 12 layers of BERT, abbreviated as *12-cl*s. The lower six Transformer layers encode the masked text to output hidden states at each position. We collectively refer to non-[CLS] positions' representations as *6-hiddens*. Then, we feed the concatenation of the *12-cl*s and *6-hiddens* into the additional two layers to perform token prediction for the masked positions. Such a design brings two benefits.

- Since the auxiliary network only contains 8 Transformer layers and the frozen lower six lay-

ers cannot be optimized, the sentence reconstruction ability is limited. Also, the *non*-[CLS] embeddings are the outputs of the 6th Transformer layer, with insufficient semantic information learned. So the MLM task is forced to rely more on the *12-cl*s embedding, encouraging the *12-cl*s embedding to encode richer semantic information.

- The gradient update of the auxiliary network is only back-propagated to the BERT network through the *12-cl*s embedding. Compared to performing the MLM task directly on the BERT, the effect of gradient updates using the *12-cl*s embedding will be much smaller and will not cause large perturbations to the contrastive learning task.

Therefore, under the InfoCSE framework, the *12-cl*s embedding learned can be distinguished from other sentence embeddings through contrastive learning and reconstructing sentences through auxiliary MLM training, while avoiding the disadvantage that the gradient of MLM objective will over-update the parameters of the encoder network.

Experiments on the semantic text similarity (STS) tasks show that InfoCSE outperforms SimCSE by an average Spearman correlation of 2.60%, 1.77% on the base of BERT-base, BERT-large, respectively. InfoCSE also significantly outperforms SimCSE on the open-domain information retrieval benchmark BEIR. We also conduct a set of ablation studies to illustrate the soundness of InfoCSE design.

2 Backgrounds

Definitions Let’s define some symbols first. Suppose we have a set of sentences $x \in \mathbb{X}$, and a 12-layer Transformer blocks BERT model Enc . For a sentence x of length l , we append a special [CLS] token to it, and then feed it into BERT for encoding. The output of each layer is a vector list of length $l + 1$, also called hidden states. We use H to denote the last layer’s hidden states and h to refer to the vector at [CLS] position. In addition, we use M to represent the hidden state of the 6th layer, $M^{>0}$ to represent the other hidden states of the 6th layer except the [CLS] position,

SimCSE As shown in subfigure (a) of Figure 1, we plot the structure diagram of SimCSE-BERT_{base}. The model’s output is the [CLS] position’s vector, which is used as the semantic representation of the sentence. SimCSE uses the same sentence to construct semantically related positive pairs $\langle x, x^+ \rangle$, i.e. $x^+ = x$. Specifically, SimCSE uses dropout as the minimum data augmentation, feeding the same input x twice to the encoder with different dropout masks z and z^+ , and outputs the hidden states of the last layer:

$$H = Enc(x, z), H^+ = Enc(x, z^+) \quad (1)$$

A pooler layer $Pooler$ is further applied to the hidden states as follows:

$$h = Pooler(H), h^+ = Pooler(H^+) \quad (2)$$

Then, for a mini-batch B , the contrastive learning objective w.r.t x is formulated as:

$$\mathcal{L}^{cl} = -\log \frac{\exp(\text{sim}(h, h^+) / \tau)}{\sum_{h' \in B} \exp(\text{sim}(h, h') / \tau)} \quad (3)$$

, where τ is a temperature hyperparameter and $\text{sim}(h, h^+)$ is the similarity metric, which is typically the cosine similarity function.

BERT’s MLM objective MLM randomly masks out a subset of input tokens and requires the model to predict them. Given a sentence x , following BERT (Devlin et al., 2018), we randomly replace 15% of the tokens with [MASK] and get a masked sentence \hat{x} . Then \hat{x} will be feed into the BERT, and the hidden states of the last layer H will be projected through a matrix W to predict the original token of each masked position. The process uses

the cross entropy loss function CE for optimization:

$$\mathcal{L}^{mlm} = \sum_{j \in \text{masked}} CE(H^j W, \hat{x}^j) \quad (4)$$

, where masked denotes the masked positions.

3 InfoCSE: Information-aggregated Contrastive Learning

In this section, we first introduce the MLM pre-training of the auxiliary network, and then describe how to jointly train the contrastive learning objective and the MLM objective with the auxiliary network.

3.1 Pre-training of The Auxiliary Network

As shown in Figure 1-(b). the auxiliary network is an 8-layer transformer network consisting of a lower six layers and an additional two layers. In the pre-training phase of the auxiliary network, the auxiliary network and the BERT network share the lower six layers. We optimize two MLM objectives simultaneously using the same input. The first is BERT’s MLM objective, which we have already covered. The second is the MLM objective of the auxiliary network. We concatenate the vector at [CLS] position of BERT’s 12th layer (h) and the hidden states of the 6th layer except for the [CLS] position ($M^{>0}$):

$$\tilde{H} = [h, M^{>0}]$$

Then \tilde{H} is fed into the additional two layers, and the output hidden state will be used to calculate the cross-entropy loss:

$$\mathcal{L}^{\text{aux}} = \sum_{j \in \text{masked}} CE(\tilde{H}^j W, \hat{x}^j) \quad (5)$$

Therefore, the pre-training loss of the entire auxiliary network is defined as the sum of the two MLM losses:

$$\mathcal{L}^{\text{pretrain}} = \mathcal{L}^{\text{aux}} + \mathcal{L}^{\text{mlm}} \quad (6)$$

The output projection matrix W is shared between the two MLM losses.

3.2 Joint Training of MLM and Contrastive Learning

When jointly training the contrastive learning and MLM objectives, the auxiliary network no longer

shares the lower six layers with BERT. But they are both initialized with the parameters from the first stage. While we still optimize the MLM objective for the auxiliary network, the lower six layers’ parameters are **frozen**. For a sentence x , it will be copied twice. One copy is used to form a positive pair $\langle x, x^+ \rangle$ for contrastive learning. The other is used to generate the masked input \hat{x} , which will be used in the MLM task. The loss of contrastive learning is the same as SimCSE, which we have already introduced. The MLM loss of the auxiliary network has changed a little, here the $M^{>0}$ is the output from the lower six layers of the auxiliary network:

$$C = [h, \text{Detach}(M^{>0})]$$

, where *Detach* means that the gradient will not back-propagate through $M^{>0}$. The loss for joint training is defined as the sum of the contrastive learning loss and the auxiliary network’s MLM loss:

$$\mathcal{L}^{\text{joint}} = \mathcal{L}^{\text{cl}} + \mathcal{L}^{\text{aux}} * \lambda \quad (7)$$

, where λ is the balanced hyperparameter for the MLM loss of the auxiliary network.

4 Experiment

4.1 Setup

When pretraining the auxiliary network, following (Gao and Callan, 2021) we use Bookcorpus (Zhu et al., 2015) and Wikipedia¹ as our datasets. When jointly optimizing contrastive objective and the auxiliary MLM objective, for a fair comparison, our experimental setup mainly follows SimCSE. We use 1-million sentences randomly drawn from English Wikipedia for training². All evaluations are directly based on the sentence embeddings output by the BERT network, and the auxiliary network will not be used.

Semantic Textual Similarity Tasks The semantic textual similarity task measures the capability of sentence embeddings, and we conduct our experiments on seven standard semantic textual similarity (STS) datasets. STS12-STS16 datasets (Agirre et al., 2012, 2013, 2014, 2015, 2016) do not have train or development sets, and thus we evaluate the models on the development set of STS-B (Cer

¹<https://dumps.wikimedia.org/enwiki/latest/enwiki-latest-pages-articles.xml.bz2>

²https://huggingface.co/datasets/princeton-nlp/datasets-for-simcse/resolve/main/wiki1m_for_simcse.txt

et al., 2017) to search for better settings of the hyper-parameters. The SentEval toolkit³ is used for evaluation, and Spearman correlation coefficient⁴ is used to report the model performance.

Open-domain Retrieval Tasks Sentence embeddings are also commonly used on the retrieval tasks, so we evaluate the zero-shot performance of embeddings on the information retrieval benchmark BEIR (Thakur et al., 2021). BEIR contains 9 open-domain information retrieval tasks (fact checking, citation prediction, duplicate question retrieval, parameter retrieval, news retrieval, question answering, tweet retrieval, biomedical IR, entity retrieval) from 18 different datasets. We evaluate InfoCSE on the 14 publicly available datasets.

4.2 Training Details

We start from the publicly pre-trained checkpoint of BERT-base or BERT-large. The hyperparameters basically follow Condenser (Gao and Callan, 2021) and SimCSE (Gao et al., 2021). When pre-training the auxiliary network, following (Gao and Callan, 2021), we train 8 epochs using the Adam optimizer with learning rate = $1e - 4$, global batch size = 1024 on 8 Nvidia V100 GPUs⁵. The hyperparameter $\lambda^{\text{mlm}} = 1.0$ to balance the two MLM losses. When jointly optimizing contrastive objective and the auxiliary MLM objective, following (Gao et al., 2021), we train one epoch using the Adam optimizer with learning rate = $3e - 5$, batch size = 64 on a single Nvidia 3090 GPU. The hyperparameter $\lambda^{\text{mlm}} = 0.005$ to balance the MLM loss and contrastive learning loss. We evaluate the model every 125 training steps on the development set of STS-B and keep the best checkpoint for the final evaluation on test sets.

4.3 Main Results

Baselines We compare our model with many strong unsupervised baselines including SimCSE (Gao et al., 2021), IS-BERT (Zhang et al.), CT-BERT (Carlsson et al., 2021), ConSERT (Yan et al., 2021), SG-OPT (Kim et al., 2021), MirrorBERT (Liu et al., 2021), ESIMCSE (Wu et al., 2021a), DiffCSE (Chuang et al., 2022) and some post-processing methods like BERT-flow (Li et al.,

³<https://github.com/facebookresearch/SentEval>

⁴https://en.wikipedia.org/wiki/Spearman%27s_rank_correlation_coefficient

⁵The pre-training stage adopts 8 training epochs without additional convergence criteria, so no extra validation set is required.

Model	STS12	STS13	STS14	SICK15	STS16	STS-B	SICK-R	Avg.
GloVe embeddings(avg.)	55.14	70.66	59.73	68.25	63.66	58.02	53.76	61.32
BERT _{base} (first-last avg.)	39.70	59.38	49.67	66.03	66.19	53.87	62.06	56.70
BERT _{base} -flow	58.40	67.10	60.85	75.16	71.22	68.66	64.47	66.55
BERT _{base} -whitening	57.83	66.90	60.90	75.08	71.31	68.24	63.73	66.28
IS-BERT _{base} △	56.77	69.24	61.21	75.23	70.16	69.21	64.25	66.58
CT-BERT _{base} △	61.63	76.80	68.47	77.50	76.48	74.31	69.19	72.05
ConSERT _{base} ♥	64.64	78.49	69.07	79.72	75.95	73.97	67.31	72.74
BERT _{base} -flow◇	63.48	72.14	68.42	73.77	75.37	70.72	63.11	69.57
SG-OPT-BERT _{base} ♠	66.84	80.13	71.23	81.56	77.17	77.23	68.16	74.62
Mirror-BERT _{base} ‡	69.10	81.10	73.00	81.90	75.70	78.00	69.10	75.40
SimCSE-BERT _{base} ♣	68.40	82.41	74.38	80.91	78.56	76.85	72.23	76.25
ESimCSE-BERT _{base} ★	73.40	83.27	77.25	82.66	78.81	80.17	72.30	78.27
DiffCSE-BERT _{base} ¶	72.28	84.43	76.47	83.90	80.54	80.59	71.23	78.49
InfoCSE-BERT _{base}	70.53	84.59	76.40	85.10	81.95	82.00	71.37	78.85
ConSERT _{large} ♥	70.69	82.96	74.13	82.78	76.66	77.53	70.37	76.45
BERT _{large} -flow◇	65.20	73.39	69.42	74.92	77.63	72.26	62.50	70.76
SG-OPT-BERT _{large} ♠	67.02	79.42	70.38	81.72	76.35	76.16	70.20	74.46
SimCSE-BERT _{large} ♣	70.88	84.16	76.43	84.50	79.76	79.26	73.88	78.41
ESimCSE-BERT _{large} ★	73.21	85.37	77.73	84.30	78.92	80.73	74.89	79.31
DiffCSE-BERT _{large} †	72.11	84.99	76.19	85.09	78.65	80.34	73.93	78.76
InfoCSE-BERT _{large}	71.89	86.17	77.72	86.20	81.29	83.16	74.84	80.18

Table 2: Sentence embedding performance on 7 semantic textual similarity (STS) test sets. ♣ : results from official published model by (Gao et al., 2021). ♥ : results from (Yan et al., 2021). ♠ : results from (Kim et al., 2021). ◇ : results from (Li et al., 2020). △ : results are reproduced and reevaluated by (Gao et al., 2021). ‡ : results from (Liu et al., 2021). ★ : results from (Wu et al., 2021a). ¶ : results from (Chuang et al., 2022). † : The original paper does not report the results of BERT-large, so we use the official public code to perform a grid search on important hyperparameters for the best results.

2020) and BERT-whitening (Su et al., 2021) along with some naive baselines like averaged GloVe embeddings (Pennington et al., 2014) and averaged first and last layer BERT (Devlin et al., 2018) embeddings. Some baseline methods are evaluated on both the base and large parameter scales, while others are only evaluated on the base parameter scale.

Semantic Textual Similarity (STS) Table 2 shows different methods’ performances on seven semantic textual similarity (STS) test sets. It can be seen that InfoCSE improves the measurement of semantic textual similarity in different parameter scale settings over previous methods. Specifically, InfoCSE-BERT_{base} outperforms SimCSE-BERT_{base} by +2.60%, InfoCSE-BERT_{large} outperforms SimCSE-BERT_{large} by +1.77%. Compared with other recent improvements to SimCSE, InfoCSE also achieves better results.

Open-domain Retrieval Task Table 3 shows the zero-shot performance of different methods on the BEIR benchmark. On the average performance of 14 datasets, InfoCSE outperforms other methods substantially, and InfoCSE achieves the best performance on 10 datasets. Although the recent

improvements to SimCSE have achieved better results than SimCSE on STS tasks, their performance on open-domain information retrieval is similar to SimCSE. In contrast, InfoCSE significantly improves on both the STS and open-domain information retrieval tasks over SimCSE, which shows that InfoCSE has better generalization ability.

5 Ablation Studies

We perform an extensive series of ablation studies to InfoCSE with BERT_{base} scale on the development set of STS-B.

5.1 The Impact of Auxiliary Network

Pre-training of The Auxiliary Network We compared the effect of directly using the parameters of the lower six layers and the upper two layers of BERT as the auxiliary network without pre-training. As shown in Table 4, the joint training effect using the auxiliary network without pre-training is significantly worse, illustrating the importance of the pre-training process. Pre-training enables the additional two-layer transformer to fit the MLM task well, avoiding large gradient oscillations during the joint learning.

Dataset	SimCSE		ESimCSE		DiffCSE		InfoCSE	
	base	large	base	large	base	large	base	large
trec-covid	0.2750	0.2264	0.2291	0.2829	0.2368	0.2291	0.3937	<u>0.3166</u>
nfcopus	0.1048	0.1356	0.1149	0.1483	0.1204	0.1470	0.1358	0.1576
nq	0.1628	0.1671	0.0935	0.1705	0.1188	0.1556	0.2023	<u>0.1790</u>
fiqa	0.0985	0.0975	0.0731	0.1117	0.0924	<u>0.1027</u>	0.0991	0.1000
arguana	0.2796	0.2078	<u>0.3376</u>	0.2604	0.2500	0.2572	0.3244	0.4133
webis-touche2020	0.1342	0.0878	0.0786	0.1057	0.0912	0.0781	<u>0.0935</u>	0.0920
quora	0.7375	0.7511	0.7411	0.7615	0.7491	0.7471	<u>0.8241</u>	0.8268
cqadupstack	0.1349	0.1082	0.1276	0.1196	0.1197	0.1160	0.2097	<u>0.1881</u>
dbpedia-entity	0.1662	0.1495	0.1260	0.1650	0.1537	0.1571	0.2101	<u>0.1838</u>
scidocs	0.0611	0.0688	0.0657	0.0796	0.0673	0.0699	<u>0.0837</u>	0.0859
climate-fever	0.1420	0.1065	0.0796	0.1302	0.1019	<u>0.1087</u>	0.0937	0.0840
scifact	0.2492	0.2541	0.3013	0.2875	0.2666	0.2811	<u>0.3269</u>	0.3801
hotpotqa	0.2382	0.1896	0.1213	0.1970	0.1730	0.2068	0.3177	<u>0.2781</u>
fever	0.2916	0.1776	0.0756	0.1689	0.1416	0.1849	0.1978	0.1252
average	0.2197	0.1948	0.1832	0.2135	0.1916	0.2030	0.2509	<u>0.2436</u>

Table 3: Zero-shot evaluation results on the BEIR benchmark. All scores denote **nDCG@10**. The best score on a given dataset is marked in **bold**, and the second best is underlined.

Joint Training of MLM and Contrastive Learning

We remove different loss terms to study their importance in joint training. We list the results in Table 5. When we remove the MLM loss, the InfoCSE model degenerates the SimCSE model, and the performance on STS-B drops by 3%. When we remove the contrastive loss, the InfoCSE model degenerates the pre-trained auxiliary network, the performance drops sharply. This result shows that both the contrastive loss and the auxiliary MLM loss are crucial in InfoCSE.

5.2 The Impact of Gradient Detach in Joint Training

During joint training, we freeze the parameters of the lower six layers of the auxiliary network through the gradient detach operation. In other words, the lower six layers are only used as feature extractors for the non-[CLS] positions. The additional two transformer layers have a limited number of parameters, and the modeling ability is weak. As shown in Table 6, the joint training effect without gradient separation operation is reduced by 1%. Therefore, the gradient detach operation forces the MLM task to rely more on the sentence embedding of *12-cl*s embedding, which is beneficial to the learning of sentence embedding.

Model	STS-B
InfoCSE	85.49
<i>w/o pre-training</i>	83.73

Table 4: Development set results of STS-B for InfoCSE with or without auxiliary network pre-training. “w/o” denotes without.

Model	STS-B
InfoCSE	85.49
<i>w/o MLM loss</i>	82.45
<i>w/o Contrastive loss</i>	40.00

Table 5: Development set results of STS-B for InfoCSE variants, where we vary the objective. “w/o” denotes without.

Model	STS-B
InfoCSE	85.49
<i>w/o gradient detach</i>	84.41

Table 6: Development set results of STS-B for InfoCSE with or without gradient detach. “w/o” denotes without.

5.3 The Impact of Mask Rate

(Wettig et al., 2022) shows that 15% may not always be the optimal mask rate. Therefore, we further explore the impact of different mask rates for

Mask Rate	10%	15%	20%	25%
STS-B	83.97	84.62	84.74	85.08
Mask Rate	30%	35%	40%	45%
STS-B	84.19	84.70	85.49	84.13

Table 7: Development set results of STS-B when we vary the mask rate.

λ	0.	5e-6	1e-5	5e-5
STS-B	82.45	84.50	85.49	84.7
λ	1e-4	5e-3	1e-2	5e-2
STS-B	84.29	80.48	76.15	75.27

Table 8: Development set results of STS-B when we vary the coefficient λ .

Model	cls	cls_before_pooler
SimCSE	81.72	82.45
DiffCSE	83.90	84.56
InfoCSE	85.08	85.49

Table 9: Development set results of STS-B where we vary the pooler choice. “cls” denotes using the representation of [CLS] token; “cls_before_pooler” denotes using the representation of [CLS] token without the extra linear+activation.

MLM task in joint training. We vary the mask rate from 10% to 45%. As shown in 7, when the mask rate increase from 10% to 25%, the model performance steadily increases. When the mask rate is over 25%, the model performance begins to fluctuate, and the optimal mask rate is 40%.

5.4 The Impact of Coefficient λ

Section 3.2 uses the λ coefficient to weight the auxiliary MLM loss and add it with contrastive loss. Because the contrastive learning objective is a relatively easier task, the scale of contrastive loss will be thousands of times smaller than MLM loss. As a result, we need a smaller λ to balance these two loss terms. In the Table 8 we show the STS-B result under different λ values. Note that when λ goes to zero, the model becomes a SimCSE model. We find that using $\lambda = 1e-5$ can give us the best performance.

5.5 The Impact of Pooler

There are two different pooler methods in SimCSE. One is “cls_before_pooler”, which is the [CLS]

representation of BERT’s last layer. The other is “cls”, which is “cls_before_pooler” with additional linear and activation. Here we compare the effects of using two poolers in different models. As shown in Table 9, “cls_before_pooler” is always better than “cls”, and InfoCSE is significantly better than other models when using either pooler.

6 Analysis

In this section, we further analyze the performance of InfoCSE on transfer tasks and in-domain information retrieval task. We also explore whether different auxiliary objectives can coexist.

6.1 Transfer Tasks

Dataset The transfer tasks include: MR (movie review) (Pang and Lee, 2005), CR (product review) (Hu and Liu, 2004), SUBJ (subjectivity status) (Pang and Lee, 2004), MPQA (opinion-polarity) (Wiebe et al., 2005), SST-2 (binary sentiment analysis) (Socher et al., 2013), TREC (question-type classification) (Voorhees and Tice, 2000) and MRPC (paraphrase detection) (Dolan and Brockett, 2005). In these transfer tasks, we will use a logistic regression classifier trained on top of the frozen sentence embeddings, following the standard setup⁶. We show the results of transfer tasks in Table 10. Compared with SimCSE, InfoCSE can improve the average scores from 85.81% to 86.63%. In SimCSE, the authors also propose directly using BERT’s MLM task to further boost the performance of transfer tasks, but it brings a consistent drop in STS tasks. Compared with SimCSE with MLM, InfoCSE achieves comparable results on transfer tasks without a drop in STS tasks. This shows that the proposed auxiliary MLM network can simultaneously improve contrastive sentence embedding’s performance on semantic text similarity tasks and downstream tasks.

6.2 In-domain Retrieval Task

Following (Chuang et al., 2022), we further explored the performance of InfoCSE on the in-domain retrieval task. We use 2758 sentences from the STS-B test set as the in-domain retrieval corpus. There are 97 positive pairs $\langle s_1, s_2 \rangle$ in the STS-B test set, and we use all s_1 to form the query set. Then we use each query s_1 to retrieve its nearest neighbors from the corpus in the sentence embedding space and see whether its corresponding s_2 is

⁶<https://github.com/facebookresearch/SentEval>

Model	MR	CR	SUBJ	MPQA	SST	TREC	MRPC	Avg.
GloVe embeddings (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
IS-BERT _{base} \triangle	81.09	87.18	94.96	88.75	85.96	88.64	74.24	85.83
SimCSE \clubsuit	81.18	86.46	94.45	88.88	85.50	89.80	74.43	85.81
w/MLM \clubsuit	82.92	87.23	95.71	88.73	86.81	87.01	78.07	86.64
InfoCSE	81.76	86.57	94.90	88.86	87.15	90.60	76.58	86.63

Table 10: Results on transfer tasks of different sentence embedding models, in terms of accuracy. \clubsuit : results from official published model by (Gao et al., 2021). \heartsuit : results from (Yan et al., 2021). \triangle : results are reproduced and reevaluated by (Gao et al., 2021). \spadesuit : results from (Chuang et al., 2022).

Model	R@1	R@5	R@10
SimCSE	74.23	94.85	95.88
DiffCSE	79.38	96.91	98.97
InfoCSE	80.41	100.00	100.00

Table 11: In-domain retrieval results. “R@” denotes recall.

in the top-1/5/10 ranking. We show quantitative results in Table 11. We can observe that InfoCSE outperforms SimCSE and DiffCSE on recall@1/5/10, demonstrating the effectiveness of using InfoCSE for retrieval tasks.

6.3 Compatibility of Different Auxiliary Objectives

Likewise, DiffCSE also introduces an auxiliary objective, replaced token detection (RTD), which brings non-trivial improvements to the ability of sentence embeddings. The difference is that RTD is a discriminative objective, while MLM is a generative objective. Therefore, we further explore whether these two different auxiliary objectives can coexist. Specifically, we simultaneously apply the [CLS] sentence embedding to the auxiliary MLM objective and RTD objective for joint optimization. As shown in Table 12, it is encouraging that the two objectives are well compatible and the combined approach achieves further improvements. This suggests that designing auxiliary objectives could be a promising direction for improving sentence embedding ability, and different auxiliary objectives have the potential to be compatible with each other if properly designed.

Model	STSB	Avg.
SimCSE	82.45	76.25
w/ RTD (DiffCSE)	84.56	78.27
w/ MLM (InfoCSE)	85.49	78.49
w/ RTD + MLM	85.83	79.39

Table 12: The comparison of the improvement brought by different auxiliary objectives to SimCSE. “w/” denotes without. “STSB” denotes the best result on the STS-B development set. “Avg.” denotes the corresponding average result on 7 semantic textual similarity (STS) test sets.

7 Related Work

Unsupervised sentence representation learning has been widely studied. (Socher et al.; Hill et al.; Le and Mikolov) propose to learn sentence representation according to the internal structure of each sentence. (Kiros et al.; Logeswaran and Lee) predict the surrounding sentences of a given sentence based on the distribution hypothesis. (Pagliardini et al.) propose Sent2Vec, a simple unsupervised model allowing to compose sentence embeddings using word vectors along with n-gram embeddings. Recently, contrastive learning has been explored in unsupervised sentence representation learning and has become a promising trend (Zhang et al.; Wu et al.; Meng et al.; Liu et al., 2021; Gao et al., 2021; Yan et al., 2021; Wu et al., 2021a; Zhou et al., 2022; Wu et al., 2021b; Chuang et al., 2022). Those contrastive learning based methods for sentence embeddings are generally based on the assumption that a good semantic representation should be able to bring similar sentences closer while pushing away dissimilar ones. The most related ones

are (Gao et al., 2021) and (Chuang et al., 2022). We solve the problem in (Gao et al., 2021) that the MLM task will reduce the model effect on STS through designing an auxiliary network. Moreover, we achieve further improvement on STS via information aggregation. (Chuang et al., 2022) is similar to us in network structure, and we both design an auxiliary network. But the motivations are quite different. (Chuang et al., 2022) aims to introduce equivariant contrastive learning to SimCSE. Then the sentence representations will be aware of, but not necessarily invariant to, direct surface-level augmentations. We aim to simultaneously achieve information aggregation while solving the compatibility problem of MLM and contrastive learning.

8 Conclusion and Future Work

In this paper, we present InfoCSE, an information-aggregated contrastive learning framework for learning unsupervised sentence embeddings. Empirical improvements on different datasets show the effectiveness and transferability of InfoCSE. We also conduct extensive ablation studies to demonstrate the different modeling choices in InfoCSE. We believe that improving sentence embeddings by optimizing auxiliary objectives is a very promising direction. In the future, we will further explore its potential on more tasks, such as dense passage retrieval (Wu et al., 2022).

9 Limitations

One limitation of our work is that we do not explore the supervised setting that uses human-labeled NLI datasets to further boost the performance.

References

- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Inigo Lopez-Gazpio, Montse Maritxalar, Rada Mihalcea, et al. 2015. Semeval-2015 task 2: Semantic textual similarity, english, spanish and pilot on interpretability. In *Proceedings of the 9th international workshop on semantic evaluation (SemEval 2015)*, pages 252–263.
- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel M Cer, Mona T Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2014. Semeval-2014 task 10: Multilingual semantic textual similarity. In *SemEval@ COLING*, pages 81–91.
- Eneko Agirre, Carmen Banea, Daniel Cer, Mona Diab, Aitor Gonzalez Agirre, Rada Mihalcea, German Rigau Claramunt, and Janyce Wiebe. 2016. Semeval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation. In *SemEval-2016. 10th International Workshop on Semantic Evaluation; 2016 Jun 16-17; San Diego, CA. Stroudsburg (PA): ACL; 2016. p. 497-511. ACL (Association for Computational Linguistics)*.
- Eneko Agirre, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. 2012. Semeval-2012 task 6: A pilot on semantic textual similarity. In ** SEM 2012: The First Joint Conference on Lexical and Computational Semantics—Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 385–393.
- Eneko Agirre, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and Weiwei Guo. 2013. * sem 2013 shared task: Semantic textual similarity. In *Second joint conference on lexical and computational semantics (* SEM), volume 1: proceedings of the Main conference and the shared task: semantic textual similarity*, pages 32–43.
- Fredrik Carlsson, Amaru Cuba Gyllensten, Evangelia Gogoulou, Erik Ylipää Hellqvist, and Magnus Sahlgren. 2021. Semantic re-tuning with contrastive tension. In *International Conference on Learning Representations*.
- Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation. *arXiv preprint arXiv:1708.00055*.
- Yung-Sung Chuang, Rumen Dangovski, Hongyin Luo, Yang Zhang, Shiyu Chang, Marin Soljačić, Shang-Wen Li, Wen-tau Yih, Yoon Kim, and James Glass. 2022. Diffcse: Difference-based contrastive learning for sentence embeddings. *arXiv preprint arXiv:2204.10298*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- William B Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Luyu Gao and Jamie Callan. 2021. Condenser: a pre-training architecture for dense retrieval. *arXiv preprint arXiv:2104.08253*.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.
- Felix Hill, Kyunghyun Cho, and Anna Korhonen. Learning distributed representations of sentences from unlabelled data. *arXiv preprint arXiv:1602.03483*.

- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining, pages 168–177.
- Taeuk Kim, Kang Min Yoo, and Sang-goo Lee. 2021. Self-guided contrastive learning for bert sentence representations. arXiv preprint arXiv:2106.07345.
- Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Skip-thought vectors. In Advances in neural information processing systems.
- Quoc Le and Tomas Mikolov. Distributed representations of sentences and documents. In International conference on machine learning.
- Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020. On the sentence embeddings from pre-trained language models. arXiv preprint arXiv:2011.05864.
- Fangyu Liu, Ivan Vulić, Anna Korhonen, and Nigel Collier. 2021. Fast, effective, and self-supervised: Transforming masked language models into universal lexical and sentence encoders. arXiv preprint arXiv:2104.08027.
- Lajanugen Logeswaran and Honglak Lee. An efficient framework for learning sentence representations. arXiv preprint arXiv:1803.02893.
- Shuqi Lu, Di He, Chenyan Xiong, Guolin Ke, Waleed Malik, Zhicheng Dou, Paul Bennett, Tiejun Liu, and Arnold Overwijk. 2021. Less is more: Pre-train a strong text encoder for dense retrieval using a weak decoder. arXiv preprint arXiv:2102.09206.
- Yu Meng, Chenyan Xiong, Payal Bajaj, Saurabh Tiwary, Paul Bennett, Jiawei Han, and Xia Song. Coco-lm: Correcting and contrasting text sequences for language model pretraining. arXiv preprint arXiv:2102.08473.
- Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. Unsupervised learning of sentence embeddings using compositional n-gram features. arXiv preprint arXiv:1703.02507.
- Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. arXiv preprint cs/0409058.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. arXiv preprint cs/0506075.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), pages 1532–1543.
- Richard Socher, Eric Huang, Jeffrey Pennington, Christopher D Manning, and Andrew Ng. Dynamic pooling and unfolding recursive autoencoders for paraphrase detection. Advances in neural information processing systems.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In Proceedings of the 2013 conference on empirical methods in natural language processing, pages 1631–1642.
- Jianlin Su, Jiarun Cao, Weijie Liu, and Yangyiwen Ou. 2021. Whitening sentence representations for better semantics and faster retrieval. arXiv preprint arXiv:2103.15316.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. Beir: A heterogeneous benchmark for zero-shot evaluation of information retrieval models. arXiv preprint arXiv:2104.08663.
- Ellen M Voorhees and Dawn M Tice. 2000. Building a question answering test collection. In Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval, pages 200–207.
- Alexander Wettig, Tianyu Gao, Zexuan Zhong, and Danqi Chen. 2022. Should you mask 15% in masked language modeling? arXiv preprint arXiv:2202.08005.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. Language resources and evaluation, 39(2):165–210.
- Xing Wu, Chaochen Gao, Liangjun Zang, Jizhong Han, Zhongyuan Wang, and Songlin Hu. 2021a. Esimcse: Enhanced sample building method for contrastive learning of unsupervised sentence embedding. arXiv preprint arXiv:2109.04380.
- Xing Wu, Chaochen Gao, Liangjun Zang, Jizhong Han, Zhongyuan Wang, and Songlin Hu. 2021b. Smoothed contrastive learning for unsupervised sentence embedding. arXiv preprint arXiv:2109.04321.
- Xing Wu, Guangyuan Ma, Meng Lin, Zijia Lin, Zhongyuan Wang, and Songlin Hu. 2022. Contextual mask auto-encoder for dense passage retrieval. arXiv preprint arXiv:2208.07670.
- Zhuofeng Wu, Sinong Wang, Jiatao Gu, Madian Khabsa, Fei Sun, and Hao Ma. 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. arXiv preprint arXiv:2105.11741.

Yan Zhang, Ruidan He, Zuozhu Liu, Kwan Hui Lim, and Lidong Bing. An unsupervised sentence embedding method by mutual information maximization. arXiv preprint arXiv:2009.12061.

Kun Zhou, Beichen Zhang, Wayne Xin Zhao, and Ji-Rong Wen. 2022. Debiased contrastive learning of unsupervised sentence representations. arXiv preprint arXiv:2205.00656.

Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In Proceedings of the IEEE international conference on computer vision, pages 19–27.