

Relational Sentence Embedding for Flexible Semantic Matching

Bin Wang[†], Haizhou Li^{‡,§,†}

[†]National University of Singapore, Singapore

[‡]School of Data Science, The Chinese University of Hong Kong, Shenzhen, China

[§]Shenzhen Research Institute of Big Data

bwang28c@gmail.com

Abstract

We present Relational Sentence Embedding (RSE), a new paradigm to further discover the potential of sentence embeddings. Prior work mainly models the similarity between sentences based on their embedding distance. Because of the complex semantic meanings conveyed, sentence pairs can have various relation types, including but not limited to entailment, paraphrasing, and question-answer. It poses challenges to existing embedding methods to capture such relational information. We handle the problem by learning associated relational embeddings. Specifically, a relation-wise translation operation is applied to the source sentence to infer the corresponding target sentence with a pre-trained Siamese-based encoder. Later, the fine-grained relational similarity scores can be estimated with learned relational embeddings. We benchmark our method on 19 datasets covering a wide range of tasks, including semantic textual similarity, transfer, and domain-specific tasks. Experimental results show that our method is effective and flexible in modelling sentence relations and outperforms a series of state-of-the-art sentence embedding methods.¹

1 Introduction

Sentence representation learning is a long-standing topic in natural language understanding and has broad applications in many tasks, including retrieval, clustering, and classification (Cer et al., 2018). It has been studied extensively in existing literature (Kiros et al., 2015; Conneau et al., 2017; Reimers and Gurevych, 2019; Gao et al., 2021b) and the problem is defined as learning vector representation (e_i) for sentence (s_i) and the similarity between sentences (s_i, s_j) is inferred by their pairwise cosine similarity: $\cos(e_i, e_j)$. Prior work found that the entailment relation in natural language inference (NLI) data aligns well with the

¹Our code is available <https://github.com/BinWang28/RSE>

Dataset	Relation	STS_avg
SNLI+MNLI	*Entailment	81.33
SNLI+MNLI	Entailment	77.67
QQP (149K)	Duplicate Question	75.96
Flicker (318K)	Same Caption	74.55
ParaNMT (5M)	Paraphrase	75.32
QNLI (55K)	Question Answer	72.21
All-Above	All 5 relations	74.62
RSE	<i>Mixed</i>	82.17
WIKI (1M)	Dropout	75.45

Table 1: Comparison of different supervised signals as positive sentence pairs. Averaged results on STS datasets are reported with $BERT_{base}$ model. * indicates the model is trained with hard negatives using contradiction pairs (Gao et al., 2021b). Experimental details can be found in Appendix A.

semantic similarity concept defined in STS datasets (Conneau et al., 2017; Reimers and Gurevych, 2019; Gao et al., 2021b). However, the concept of similarity is vaguely defined, and the relationship between sentences is indeed multifaceted (Wang et al., 2022). One can only rely on experimental verification to determine what relation between sentences aligns well with human perception of “semantic similarity”. The similarity intensity between sentences varies greatly based on the perspective of personal understanding.

In this work, we do not focus on discussing the detailed definition of semantic relations but focus on the explicit modelling of sentence relations. Instead of treating sentence pairs as either positive or negative samples in contrastive learning (Gao et al., 2021b; Yan et al., 2021; Wu et al., 2022), we treat sentence pairs as triples: ($s_i, relationship, s_j$), with relational information embedded. We learn a translational operation between sentence embeddings to contrastively model the relationship between sentences. Despite its simplicity, the re-

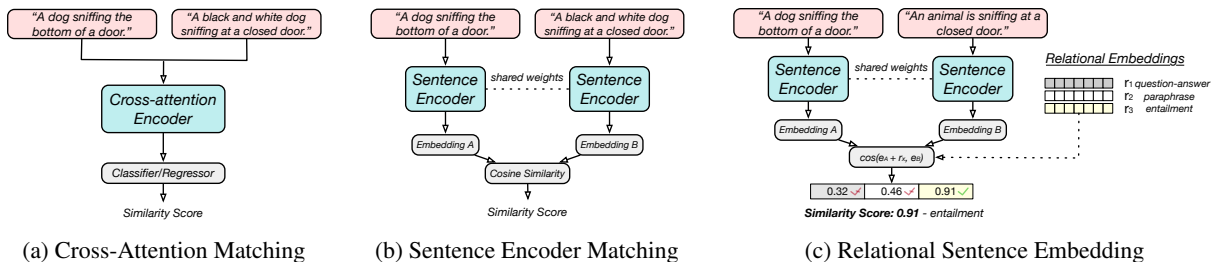


Figure 1: Comparison of different sentence-pair matching architectures. Relational sentence embedding adopts relation-specific embeddings for more flexible semantic matching with interpretable results. Meantime, the Siamese network allows much faster inference speech compared with cross-attention matching.

lational modelling component largely expands the concept of existing sentence embedding and introduces unique properties for more flexible semantic matching capabilities. Our method is called *Relational Sentence Embedding* (RSE).

First, with relation modelling, RSE can leverage multi-source relational data for improved generalizability. Earlier work mainly focuses on using NLI data for supervised sentence embedding learning, as the best performance is witnessed. It remains an under-explored problem to learn from multi-source relational data jointly for sentence representation learning. Table 1 shows our experimental study to combine multi-source relational data as the positive samples for contrastive sentence embedding learning (Gao et al., 2021b). Simply merging all relational data leads to performance down-gradation (ref. 'All-Above'). In contrast, with explicit relation modelling by translation operations, RSE can effectively leverage multi-source relational information into a shared sentence encoder while providing exceptional performance.

Second, unlike conventional SE, RSE can provide fine-grained relational similarity scores of sentence pairs. As shown in Figure 1b, conventional sentence embedding computes one unified similarity score for sentence pairs. It is less flexible and interpretable as no detailed information is provided about what aspects of sentences are similar. In contrast, as shown in Figure 1c, RSE can infer fine-grained sentence similarity scores according to the relation of interest, which are critical features for heterogeneous applications (e.g., retrieval and clustering). RSE can vastly extend the sentence embedding concept and has excellent potential in methodology development and applications.

We benchmark our approach on 19 datasets to verify the effectiveness and generalizability of our approach. For semantic textual similarity, we con-

duct experiment on seven standard STS (Agirre et al., 2012, 2013, 2014, 2015, 2016; Cer et al., 2017; Marelli et al., 2014) datasets and STR dataset (Abdalla et al., 2021). RSE achieves an averaged Spearman’s correlation of 82.18% with BERT_{base} model on STS datasets. We also achieve SOTA performance on 7 transfer tasks (Conneau and Kiela, 2018), and 4 domain-specific tasks (Wang et al., 2021b). We also present the case study and in-depth discussions.

2 Related Work

Sentence embedding can be divided into unsupervised and supervised approaches. Even though unsupervised methods have greatly improved with effective data augmentation techniques for contrastive learning (Gao et al., 2021b; Wu et al., 2022; Chen et al., 2022), supervised methods still outperform their unsupervised counterparts. The success of supervised sentence embedding is mainly because of the availability of high-quality natural language inference (NLI) data (Bowman et al., 2015; Williams et al., 2018). InferSent (Conneau et al., 2017), SBERT (Reimers and Gurevych, 2019), and SimCSE (Gao et al., 2021b) are three milestones on sentence embedding that effectively leverage NLI data using LSTM, Pre-trained BERT and contrastive learning, respectively. The NLI data covers three sentence relations: entailment, contradiction and neutral. The relationships between sentences are manifold. However, prior work found that training with other supervised signals (Young et al., 2014; Wieting and Gimpel, 2018) does not achieve comparable performance as NLI data (Gao et al., 2021b). In this work, we focus on learning from multi-source relational signals for better sentence representation learning.

Knowledge graph embeddings are designed for relational data modelling on triples: (h, r, t) . The




Figure 2: RSE takes the corresponding relation transformation as positive triples and treats other in-batch samples as negatives.

methods can be classified into structural-based models and semantic-enhanced models. The structural-based methods learn relational embedding and entity embedding jointly through scoring functions like TransE (Bordes et al., 2013), DistMult (Yang et al., 2015) and RotatE (Sun et al., 2019). Semantic-enhanced models incorporate additional information for entities, including entity types and textual descriptions, when learning embeddings. TEKE (Wang et al., 2016), InductiveE (Wang et al., 2021a) and KG-BERT (Yao et al., 2019) belong to this type. For sentence representation learning from multi-source supervised signals, the sentence pairs can be viewed as triples: $(s_i, relationship, s_j)$. Therefore, we optimize the semantic representation of sentences and perform translation operations between them to model their respective relations. An analogy to the word-level relational graph as Word-Net (Miller, 1995), and Relational Word Embedding (Camacho-Collados et al., 2019), our work can also be viewed as a pioneering study on the sentence level.

3 Relational Sentence Embedding

Given n pre-defined relations $R = \{r_1, r_2, \dots, r_n\}$, the multi-source relation sentence pair datasets are composed of triples: $D = \{(s_i, r_k, s_j)\}$, where s_i and s_j corresponds to the i_{th} and j_{th} sentence. The goal of relation sentence embedding is to learn sentence representations and relational embeddings jointly. The relational similarity intensity between sentences can be inferred from learned embeddings through either simple scoring functions or complex transformations. In this work, as a prototype of RSE, we use translation operation (Bordes et al., 2013) for sentence relation modelling.

3.1 Relational Contrastive Learning

Contrastive learning approaches have demonstrated their effectiveness in sentence embedding learning (Gao et al., 2021b; Janson et al., 2021). The core concept is learning representations that pull together close concepts and push away non-related ones. To build contrastive samples, we leverage the relational signal from sentence triples of supervised data. Assume we have relational triples (s_i, r_k, s_j) , let \mathbf{h}_i and \mathbf{h}_j denote the representation of s_i and s_j obtained from sentence encoder $\mathbf{h}_i = f_\theta(s_i)$ and \mathbf{h}_k^r denote the relational embedding of the k_{th} relation. We then treat $\mathbf{h}_i + \mathbf{h}_k^r$ and \mathbf{h}_j as positive pair. Following the previous contrastive learning approach (Chen et al., 2020), other in-batch tail sentences are treated as negatives. Assuming we have N sentence triples in one mini-batch, the training objective can be written as follow:

$$\mathcal{L}_i = -\log \frac{e^{sim(\mathbf{h}_i + \mathbf{h}_k^r, \mathbf{h}_j)/\tau}}{\sum_{m=1}^N e^{sim(\mathbf{h}_i + \mathbf{h}_k^r, \mathbf{h}_m)/\tau}} \quad (1)$$

where $sim(\cdot, \cdot)$ is the cosine similarity and τ is the temperature hyperparameter. The pre-trained BERT (Devlin et al., 2019) or RoBERTa (Liu et al., 2019) model $f_\theta(\cdot)$ is used to obtain sentence embedding. The relation embedding \mathbf{h}_k^r is randomly initialized and jointly optimized with the pre-trained language model with the above relational contrastive objective.

Figure 2 is a graphical illustration of Eq. 1. The relational information is learned by relational embedding through translation operation. That is, for triple (s_i, r_k, s_j) , we expect $\mathbf{h}_i + \mathbf{h}_k^r \approx \mathbf{h}_j$. The embedding property can further infer the relation intensity between sentences during inference.

In knowledge graph embedding, there are more diverse ways to model relations, including sim-

ple scoring functions (TransE (Bordes et al., 2013), TransR (Lin et al., 2015) and RotatE (Sun et al., 2019)), semantic matching methods (DistMult (Yang et al., 2015), TUCKER (Balazevic et al., 2019)) and neural network methods (ConvE (Dettmers et al., 2018)). In this work, we experiment with the simple translational operation for relation modeling for simplicity and interpretability. Other complex relation modeling methods can be explored in future work as they are more expressive in certain relation types, including symmetry, anti-symmetry, inversion, and composition (Sun et al., 2019).

3.2 Supervised Signals

The key to contrastive sentence embedding is building positive and negative samples. ConSERT (Yan et al., 2021) proposed to use adversarial attack, token shuffling, feature cutoff, and dropout as positive pair augmentation for unsupervised sentence embedding. SimCSE (Gao et al., 2021b) further explores the potential of dropout as positive-view augmentation. In this work, we collect multi-source relational data as positive samples and view the rest in-batch samples as negatives. In general, multi-source relational data can be either supervised or unsupervised. For unsupervised learning, we can view each data augmentation method as one relation and jointly learn from multiple data augmentation methods. In this work, we focus on artificial sentence triples and explore how to combine multiple human-labelled datasets with different relation types.

We deploy the following datasets for sentence triples (s_i, r_k, s_j) . 1) SNLI (Bowman et al., 2015) + MNLI (Williams et al., 2018): contain natural language inference data. The relations are *entailment* (r_1), *contradiction* and *neutral*. We use *entailment* relation with *contradiction* relation as hard negatives. 2) QQP²: The Quora Question dataset. The label for each question pair is either *duplicate-question* (r_2) or *non-duplicate-question*. We include the triples with *duplicate-question* relation. 3) ParaNMT (Wieting and Gimpel, 2018): *paraphrase* (r_3) dataset generated by back-translation with filtering. 4) Flicker (Young et al., 2014): image captioning dataset. Each image is with 5 human-written captions. We treat the captions for the same image as relation *same-caption* (r_4).

²<https://data.quora.com/First-Quora-Dataset-Release-Question-Pairs>

5) QNLI (Wang et al., 2018): contains question-answer pairs. The label is *qa-entailment* (r_5) or *qa-not-entailment* based on whether the sentences contain the true answer to the question. We only use *qa-entailment* triples. 6) Dropout (Gao et al., 2021b): is constructed in an unsupervised way. We use the Wiki-1M dataset and treat the sentence with different random dropouts as one relation. We call this relation *same-sent* (r_6). We leverage all samples from the NLI dataset and at most 150K samples for other relations as the size of QQP for easy comparison.

3.3 Mining Hard Negatives

Contrastive learning commonly uses in-batch sentences as negative samples (Janson et al., 2021). However, the in-batch negatives are limited in relational sentence embeddings. It is because the in-batch samples are coming from different relation triples. The sentences from other relations are less likely to be confused with the current anchor triple. Therefore, we extend the (s_i, r_k, s_j) triple to (s_i, r_k, s_j, s_j^-) with an extra in-relation negative sample s_j^- . The s_j^- is one random sampled sentence from the triple of relation r_k . To better understand, the in-relation negative for relation *duplicate-question* will be a random question, a ‘harder’ negative than in-batch sentences. For the *entailment* relation, we use *contradiction* sentence as the hard negative. With the hard negative, the training objective can be written as follows (Ref. to Eq.1):

$$-\log \frac{e^{\text{sim}(\mathbf{h}_i + \mathbf{h}_k^r, \mathbf{h}_j)/\tau}}{\sum_{m=1}^N (e^{\text{sim}(\mathbf{h}_i + \mathbf{h}_k^r, \mathbf{h}_m)/\tau} + e^{\text{sim}(\mathbf{h}_i + \mathbf{h}_k^r, \mathbf{h}_m^-)/\tau})} \quad (2)$$

As multiple relation triples are mixed in each batch during training, the batch size for each relation becomes smaller. Contrastive learning favors a relatively large batch size to reach good performance. Wu et al. (2022) incorporate momentum contrast (He et al., 2020) for better contrastive learning. However, we did not find its contribution to the supervised setting. As the GPU memory size constrains the batch size, we scale up the batch size of contrastive learning by decoupling back-propagation between contrastive learning and decoder using Gradient Cache (Gao et al., 2021a). As a result, we can train BERT_{base} model with batch size 512 within an 11G GPU memory limit.

How can we use? As the relational information is

Model	STS12	STS13	STS14	STS15	STS16	STS-B	SICK-R	Avg.
<i>Unsupervised methods</i>								
GloVe (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
CT - BERT _{base}	61.63	76.80	68.47	77.50	76.48	74.31	69.19	72.05
SimCSE - BERT _{base}	68.40	82.41	74.38	80.91	78.56	76.85	72.23	76.25
SimCSE - RoBERTa _{base}	70.16	81.77	73.24	81.36	80.65	80.22	68.56	76.57
<i>Supervised methods</i>								
InferSent - GloVe	52.86	66.75	62.15	72.77	66.87	68.03	65.65	65.01
USE	64.49	67.80	64.61	76.83	73.18	74.92	76.69	71.22
SBERT _{base}	70.97	76.53	73.19	79.09	74.30	77.03	72.91	74.89
SimCSE - BERT _{base}	75.30	84.67	80.19	85.40	80.82	84.25	80.39	81.57
RSE - BERT _{base}	76.27 _{±.34}	84.43 _{±.23}	80.60 _{±.12}	86.03 _{±.09}	81.86 _{±.12}	84.34 _{±.13}	81.73 _{±.07}	82.18 _{±.06}
RSE - BERT _{large}	76.94 _{±.40}	86.63 _{±.43}	81.81 _{±.28}	87.02 _{±.22}	83.25 _{±.22}	85.63 _{±.23}	82.68 _{±.24}	83.42 _{±.19}
RSE - RoBERTa _{base}	76.71 _{±.47}	84.38 _{±.34}	80.87 _{±.56}	87.12 _{±.42}	83.33 _{±.19}	84.62 _{±.18}	81.97 _{±.32}	82.71 _{±.20}
RSE - RoBERTa _{large}	77.81 _{±.37}	87.04 _{±.20}	82.73 _{±.22}	87.92 _{±.25}	85.44 _{±.15}	86.79 _{±.08}	83.13 _{±.07}	84.41 _{±.08}

Table 2: Performance on STS tasks. Spearman’s correlation is reported. For RSE approach, we also report the standard derivation of performance based on five runs with different random seeds. The best performance among all models and the best performance with BERT_{base} backbone is displayed in **bold**.

Model	STR
Unsup. SimCSE - BERT _{base}	73.98
Sup. SimCSE - BERT _{base}	80.72
RSE - BERT _{base}	81.23 _{±.13}
RSE - BERT _{large}	81.98 _{±.18}
RSE - RoBERTa _{base}	82.57 _{±.30}
RSE - RoBERTa _{large}	84.10 _{±.18}

Table 3: Comparison of SimCSE and RSE performance on STR dataset.

explicitly embedded in RSE, it becomes very flexible for different applications. Given two sentence (s_i, s_j), we can infer their similarity score under relation r_k by $f(s_i, s_j, r_k) = \text{sim}(\mathbf{h}_i + \mathbf{h}_k^r, \mathbf{h}_j)$. For retrieval and semantic similarity, we can use any relational similarity score or weighted sum of different relational scores as the final measure. We apply this in STS and USEB benchmarking tasks. As an additional classifier is trained, we use sentence embedding as the input for transfer tasks. We expect that with relational modelling capability, the result sentence embedding carries richer information.

4 Experiments on STS Tasks

We conduct experiment on 7 semantic textual similarity datasets including STS 12-16 (Agirre et al., 2012, 2013, 2014, 2015, 2016), STS-Benchmark (Cer et al., 2017) and SICK-

Relatedness (Marelli et al., 2014). We also experiment with STR-2022 (Abdalla et al., 2021). A new STS dataset with improved label quality using comparative ranking. For a fair comparison, we report the result with the fully unsupervised setting, Spearman’s correlation and ‘all’ for result aggregation following prior works (Reimers and Gurevych, 2019; Gao et al., 2021b).

4.1 Settings

We initialize our training checkpoint with pre-trained BERT_{base, large} (Devlin et al., 2019) and RoBERTa_{base, large} (Liu et al., 2019). The [CLS] token representation is taken as sentence representation (\mathbf{h}_i). We compare with both unsupervised and supervised sentence embedding methods, including 1) average GloVe embedding (Pennington et al., 2014), 2) average BERT embedding (Su et al., 2021), 3) CT-BERT (Janson et al., 2021), 4) InferSent (Conneau et al., 2017), 5) USE (Cer et al., 2018), 6) SBERT (Reimers and Gurevych, 2019) and 7) SimCSE (Gao et al., 2021b). For a fair comparison, we did not include T5-based methods to avoid the influence of pre-trained checkpoint (Ni et al., 2022; Chen et al., 2022). We use r_1 and r_2 as our labeled data with relational weights 1.0 and 0.5, respectively. We train our model for 3 epochs and evaluate every 125 training steps on the development set of STS-B to keep the best-performed checkpoint.

τ	0.001	0.01	0.05	0.1	1.0
Avg. STS	80.74	81.73	82.18	77.00	54.88

Table 4: Ablation study on temperature τ .




Figure 3: Performance comparison using different weights for ‘duplicate-question’ relation for STS tasks.

4.2 Main Results

Table 2 shows the evaluation results on 7 standard STS datasets. RSE can outperform all previous models in the supervised setting and improve the averaged STS result from 81.57% to 82.18% for Spearman’s correlation with $BERT_{base}$ model. A further performance boost is witnessed with better pre-training checkpoints and larger model sizes. The averaged STS result with $RoBERTa_{large}$ model reaches 84.41%. We then experiment with five random seeds to test training stability and report the standard derivation. Results show that our method is robust against randomness. The performance on the STR-2022 dataset is shown in Table 3. We observe the same result patterns, and RSE can outperform the supervised SimCSE model.

Effect of Temperature. We study the influence of different τ values, and the result is shown in Table 4. $\tau = 0.05$ provides the best result, and we use the same setting for all experiments.

Effect of Relational Weight. We examine the performance by adjusting the weight of the relation ‘duplication-question’ and the results are shown in Figure 3. We witness that the avg. STS improved using the similarity score of both ‘duplicate question’ and ‘entailment’ relations. This phenomenon generalizes to different pre-trained models. For $BERT_{base}$ and $RoBERTa_{large}$ model, RSE can outperform others by only using the ‘entailment’ relation. The model can perform better when in-




Figure 4: The similarity between relational embeddings. The relations $r_{1\sim 6}$ are *entailment*, *duplicate-question*, *paraphrase*, *same-caption*, *qa-entailment*, *same-sent*.

volving more relational scores, which we leave for future exploration.

5 Experiments on USEB Tasks

As there are some limitations in domain coverage of STS tasks (Wang et al., 2021b, 2022), we also experiment with Universal Sentence Embedding Benchmark (USEB) proposed in Wang et al. (2021b). It contains four tasks from heterogeneous domains, including Re-Ranking (RR), Information Retrieval (IR) and Paraphrase Identification (PI). The datasets are AskUbuntu (RR) (Lei et al., 2016), CQADupStack (IR) (Hoogveen et al., 2015), TwitterPara (PI) (Xu et al., 2015; Lan et al., 2017), and SciDocs (RR) (Cohan et al., 2020).

5.1 Settings

Similarly, we experiment with four different pre-trained checkpoints $BERT_{base,large}$ and $RoBERTa_{base,large}$. We compare with different benchmarking including unsupervised methods: GloVe, Sent2Vec (Pagliardini et al., 2018), BM25 (Robertson et al., 1995), BERT, unsupervised SimCSE, and supervised methods: SBERT (Reimers and Gurevych, 2019), and SimCSE (Gao et al., 2021b). We did not compare with other in-domain adaption-based methods for fair comparison because RSE is an off-the-shelf method. For training, we include all relational data ($r_{1\sim 6}$ in Sec. 3.2) and use the similarity score from one specific relation for each task. We train our model for 3 epochs and evaluate every 1000 training steps on the development set of USEB to keep the best-performed checkpoint.

Model	AskU.	CQADup.	TwitterP.			SciDocs					Avg.
			TURL	PIT	Avg.	Cite	CC	CR	CV	Avg.	
GloVe (avg.)	51.0	10.0	70.1	52.1	61.1	58.8	60.6	64.2	65.4	62.2	46.1
Sent2Vec	49.0	3.2	47.5	39.9	43.7	61.6	66.0	66.1	66.7	65.1	40.2
BM25	53.4	13.3	71.9	70.5	71.2	58.9	61.3	67.3	66.9	63.6	50.4
BERT _{base}	48.5	6.5	69.1	61.7	65.4	59.4	65.1	65.4	68.6	64.6	46.3
*SimCSE - BERT _{base}	55.9	12.4	74.5	62.5	68.5	62.5	65.1	67.7	67.6	65.7	50.6
SBERT _{base}	53.4	11.8	75.4	69.9	72.7	66.8	70.0	70.7	72.8	70.1	52.0
SimCSE - BERT _{base}	53.5	12.2	76.1	69.7	72.9	66.8	69.5	72.1	70.7	69.8	52.1
RSE - BERT _{base}	54.8	13.7	77.4	73.0	75.2	67.6	71.0	71.8	73.6	71.0	53.7
RSE - BERT _{large}	56.2	13.9	76.7	76.8	76.7	68.6	71.9	72.5	74.5	71.9	54.7
RSE - RoBERTa _{base}	56.2	13.3	74.8	74.8	74.8	66.9	69.4	70.2	72.5	69.7	53.5
RSE - RoBERTa _{large}	58.0	15.2	76.2	78.8	77.5	69.3	72.4	72.9	74.8	72.4	55.8

Table 5: Results on domain-specific tasks. The out-of-the-box results of unsupervised and supervised approaches are shown. The first and second blocks show unsupervised and supervised methods, respectively. * refers to unsupervised SimCSE with dropout augmentation. Our proposed RSE outperforms all benchmarking approaches.

5.2 Main Results

Table 5 shows the evaluation result on 4 USEB tasks. Among the benchmarking methods, we found that unsupervised approaches outperform supervised approaches in AskUbuntu and CQADupStack tasks. The unsupervised SimCSE and BM25 achieve the best result on AskUbuntu and CQADupStack, respectively. It shows the heterogeneous character of USEB task and the previous supervised methods do not generalize well to various domains.

In contrast, RSE performs the best on 3 of the 4 tasks and achieves the best overall performance. The average result improves from 52.1 to 53.7 with BERT_{base} model. For AskUbuntu and CQADupStack tasks, we find that the unsupervised SimCSE method achieves the best performance, even outperforming the supervised ones. RSE framework aligns with this finding that the best re-ranking and retrieval results are obtained by the similarity inferred from the ‘same-sent’ (r_6) relation. The TwitterPara dataset is for identifying whether two sentences are paraphrased or not. In RSE, we find the ‘paraphrase’ relation (r_3) trained on synthetic ParaNMT dataset works the best, which aligns with our intuition. Similarly, better performance is achieved with a larger pre-trained model as weight initialization. The USEB experiments verify the RSE’s capability in modelling multi-source relational sentence-pair data. The relational-aware similarity scores can be flexibly applied to heterogeneous domain tasks.

Relation Similarity. We obtained 6 relation em-

beddings in our experiments after training with multi-source relational data. To study the similarity between relations, we compute the cosine similarity between pairwise relational embeddings, and the result is shown in Figure 4. From the results, we can witness the more similar relation pairs are *entailment & same-sent*, *duplicate-question & paraphrase*, *duplicate-question & same-caption*, *duplicate-question & same-sent*, and *paraphrase & same-sent*. The similar relations are more or less corresponding to sentence similarity, and we can see that the learned relational embeddings capture the relational information well. The *qa-entailment* relation models the relationship between question and answer pairs. Therefore, we can see it is indeed less similar to all other relations.

6 Experiments on Transfer Tasks

Besides the semantic similarity, ranking, and retrieval tasks, we also experiment with the following transfer tasks: MR (Pang and Lee, 2005), CR (Hu and Liu, 2004), SUBJ (Pang and Lee, 2004), MPQA (Wiebe et al., 2005), SST (Socher et al., 2013), TREC (Voorhees and Tice, 2000) and MRPC (Dolan and Brockett, 2005). Because the RSE model captures richer information by learning from multi-source relational data, we expect it generalizes better to different classification tasks.

6.1 Settings

Similarly, we experiment with four different initialization checkpoints. The averaged [CLS] token representation from the last five layers is taken as

Model	MR	CR	SUBJ	MPQA	SST	TREC	MRPC	Avg.
<i>Unsupervised methods</i>								
GloVe (avg.)	55.14	70.66	59.73	68.25	63.66	58.02	53.76	61.32
BERT _{base} (avg.)	78.66	86.25	94.37	88.66	84.40	92.80	69.54	84.94
SimCSE - BERT _{base}	81.18	86.46	94.45	88.88	85.50	89.80	74.43	85.81
SimCSE - RoBERTa _{base}	81.04	87.74	93.28	86.94	86.60	84.60	73.68	84.84
<i>Supervised methods</i>								
InferSent - GloVe	81.57	86.54	92.50	90.38	84.18	88.20	75.77	85.59
USE	80.09	85.19	93.98	86.70	86.38	93.20	70.14	85.10
SBERT _{base}	84.64	89.43	94.39	89.86	88.96	89.60	76.00	87.41
SimCSE - BERT _{base}	82.69	89.25	94.81	89.59	87.31	88.40	73.51	86.51
w/ MLM	82.68	88.88	94.52	89.82	88.41	87.60	76.12	86.86
RSE - BERT _{base}	82.35	89.01	95.33	90.56	88.36	93.00	77.39	88.00
RSE - BERT _{large}	84.25	90.49	95.70	90.51	90.17	95.20	76.70	89.04
RSE - RoBERTa _{base}	84.52	91.21	94.39	90.21	91.49	91.40	77.10	88.62
RSE - RoBERTa _{large}	86.29	91.21	95.17	91.16	91.38	95.40	78.20	89.83

Table 6: Results on 7 transfer tasks. Both unsupervised and supervised methods are compared. Two versions of SimCSE are compared. The one with an additional mask language model (MLM) objective shows better performance. RSE outperforms other approaches on average.

the sentence representation (\mathbf{h}_i) (Wang and Kuo, 2020). The sentence encoder is used as a feature extractor. An MLP-based classifier is trained with sentence embeddings as the input. The default setting from the SentEval package is deployed (Conneau and Kiela, 2018).

We leverage the relations of *entailment* and *paraphrase* for training. As sentence embedding is used as the input to classifiers, relation embeddings are not incorporated in transfer task experiments. All models are trained for 3 epochs. We evaluate the model on the development set every 500 training steps and keep the best-performed checkpoint for final evaluation. We compare RSE with different competitive methods including InferSent (Conneau et al., 2017), USE (Cer et al., 2018), SBERT (Reimers and Gurevych, 2019), and SimCSE (Gao et al., 2021b).

6.2 Results and Analysis

Table 6 shows the evaluation results on 7 transfer tasks. The RSE model with BERT_{base} architecture achieves an average score of 88.0, outperforming all benchmarked methods. SBERT outperforms SimCSE in 6 of 7 tasks. We infer that the classification-based training objective makes SBERT more transferable to transfer tasks. Even

though the contrastive learning approach outperforms in semantic similarity tasks, it does not generalize well with additional trainable classifiers. In contrast, RSE is a contrastive-based approach, which already showed superior in semantic similarity tasks. Due to its learning capability from multi-source relational data, the learned sentence embeddings have better generalizability than single-source supervised learning. It leads to better transferability of the sentence embedding obtained from RSE models. Similar to previous findings, the performance of RSE improves with larger and better pre-trained checkpoints as initialization. Even though RSE is mainly designed for flexible semantic matching with multiple relational similarity scores, the learned sentence embedding captures richer semantic meanings for various transfer tasks.

7 Conclusion and Future Work

In this work, we propose RSE, a new sentence embedding paradigm with explicit relation modelling by translation operation. RSE can learn from multi-source relational data and provide interpretable relational similarity scores. Experiments on semantic similarity, domain-specific, and transfer tasks show that RSE can outperform SOTA approaches, verifying its flexibility and effectiveness. Relational

sentence embedding can broaden the concept of conventional similarity-based sentence embedding and can be extended for different applications.

For future work, first, with more sentence relational data available, the relational sentence embedding can become more competitive and generalizable to different domains. The collection of sentence-level relational data has great potential. Second, more complex relation modelling methods can be incorporated into handling heterogeneous relations. Besides the translation operation in this work, other relation modelling functions, including rotation and semantic matching, can be further explored. Third, this work mainly focuses on supervised relation signals, and unsupervised relational sentence embedding still needs to be explored.

Limitations

The first limitation of RSE comes from data sparsity. Unlike word-level relational data, the variety of sentences is much larger than words. Therefore, even though we can collect more and more relational data on sentence pairs, it is generally hard to become densely connected between sentences. To alleviate this issue and generate sustainable sentence relational data, we should design automatic tools for sentence relation labelling with human-in-the-loop supervision.

Second, unsupervised sentence embedding also shows decent performance on semantic textual similarity tasks, and the focus is mainly on the design of unsupervised contrastive samples. Therefore, exploring different unsupervised view-augmentation techniques in relational sentence embedding remains an open question. In this work, we only study incorporate dropout augmentation without much discussion on other techniques like pseudo-labelling (Chen et al., 2022), word drop, word swap, and word repetition (Yan et al., 2021).

For relational data modelling, complex operations (Sun et al., 2019; Trouillon et al., 2016) for relation modelling are introduced and proven effective in knowledge graph completion. However, it is not straightforward to incorporate complex operations with pre-trained semantic representations. The study of RSE in complex space modelling remains unexplored.

Acknowledgement

This work is supported by the Agency for Science, Technology and Research (A*STAR) under

its AME Programmatic Funding Scheme (Project No. A18A2b0046), Singapore, and by National Natural Science Foundation of China (Grant No. 62271432).

References

- Mohamed Abdalla, Krishnapriya Vishnubhotla, and Saif M Mohammad. 2021. [What makes sentences semantically related: A textual relatedness dataset and empirical study](#). *arXiv preprint arXiv:2110.04845*.
- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Iñigo Lopez-Gazpio, Montse Maritxalar, Rada Mihalcea, German Rigau, Larraitz Uria, 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, Denver, Colorado. Association for Computational Linguistics.
- 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, Dublin, Ireland. Association for Computational Linguistics.
- Eneko Agirre, Carmen Banea, Daniel Cer, Mona 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, San Diego, California. 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, Montréal, Canada. Association for Computational Linguistics.
- 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, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Ivana Balazevic, Carl Allen, and Timothy Hospedales. 2019. [TuckER: Tensor factorization for knowledge](#)

- graph completion**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5185–5194, Hong Kong, China. Association for Computational Linguistics.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. **A large annotated corpus for learning natural language inference**. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Jose Camacho-Collados, Luis Espinosa Anke, and Steven Schockaert. 2019. **Relational word embeddings**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3286–3296, Florence, Italy. Association for Computational Linguistics.
- Daniel Cer, Mona 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, Vancouver, Canada. Association for Computational Linguistics.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. **Universal sentence encoder for English**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, Brussels, Belgium. Association for Computational Linguistics.
- 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.
- Yiming Chen, Yan Zhang, Bin Wang, Zuozhu Liu, and Haizhou Li. 2022. **Generate, discriminate and contrast: A semi-supervised sentence representation learning framework**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8150–8161, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel Weld. 2020. **SPECTER: Document-level representation learning using citation-informed transformers**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2270–2282, Online. Association for Computational Linguistics.
- Alexis Conneau and Douwe Kiela. 2018. **SentEval: An evaluation toolkit for universal sentence representations**. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. **Supervised learning of universal sentence representations from natural language inference data**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **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, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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, Yunyi Zhang, Jiawei Han, and Jamie Callan. 2021a. **Scaling deep contrastive learning batch size under memory limited setup**. In *Proceedings of the 6th Workshop on Representation Learning for NLP (RepL4NLP-2021)*, pages 316–321, Online. Association for Computational Linguistics.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021b. **SimCSE: Simple contrastive learning of sentence embeddings**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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.
- Doris Hoogeveen, Karin M. Verspoor, and Timothy Baldwin. 2015. **Cquadupstack: A benchmark data set for community question-answering research**. In

- Proceedings of the 20th Australasian Document Computing Symposium, ADCS '15*, New York, NY, USA. Association for Computing Machinery.
- 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.
- Sverker Janson, Evangelina Gogoulou, Erik Ylipää, Amaru Cuba Gyllensten, and Magnus Sahlgren. 2021. Semantic re-tuning with contrastive tension. In *International Conference on Learning Representations, 2021*.
- Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. *Advances in neural information processing systems*, 28.
- Wuwei Lan, Siyu Qiu, Hua He, and Wei Xu. 2017. A continuously growing dataset of sentential phrases. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1224–1234, Copenhagen, Denmark. Association for Computational Linguistics.
- Tao Lei, Hrishikesh Joshi, Regina Barzilay, Tommi Jaakkola, Kateryna Tymoshenko, Alessandro Moschitti, and Lluís Màrquez. 2016. Semi-supervised question retrieval with gated convolutions. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1279–1289, San Diego, California. Association for Computational Linguistics.
- Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning entity and relation embeddings for knowledge graph completion. In *Twenty-ninth AAAI conference on artificial intelligence*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. A SICK cure for the evaluation of compositional distributional semantic models. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 216–223, Reykjavik, Iceland. European Language Resources Association (ELRA).
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Jianmo Ni, Gustavo Hernandez Abrego, Noah Constant, Ji Ma, Keith Hall, Daniel Cer, and Yinfei Yang. 2022. Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1864–1874, Dublin, Ireland. Association for Computational Linguistics.
- Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. 2018. Unsupervised learning of sentence embeddings using compositional n-gram features. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 528–540, New Orleans, Louisiana. Association for Computational Linguistics.
- Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, pages 271–278, Barcelona, Spain.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 115–124, Ann Arbor, Michigan. Association for Computational Linguistics.
- 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.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. 1995. Okapi at trec-3. *Nist Special Publication Sp*, 109:109.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew 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, Seattle, Washington, USA. 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*.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. Rotate: Knowledge graph embedding by relational rotation in complex space. In *International Conference on Learning Representations*.

- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. [Complex embeddings for simple link prediction](#). In *International conference on machine learning*, pages 2071–2080. PMLR.
- 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.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Bin Wang and C.-C. Jay Kuo. 2020. SBERT-WK: A sentence embedding method by dissecting bert-based word models. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2146–2157.
- Bin Wang, C.-C. Jay Kuo, and Haizhou Li. 2022. [Just rank: Rethinking evaluation with word and sentence similarities](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6060–6077, Dublin, Ireland. Association for Computational Linguistics.
- Bin Wang, Guangtao Wang, Jing Huang, Jiakuan You, Jure Leskovec, and C.-C. Jay Kuo. 2021a. [Inductive learning on commonsense knowledge graph completion](#). In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Kexin Wang, Nils Reimers, and Iryna Gurevych. 2021b. [TSDAE: Using transformer-based sequential denoising auto-encoder for unsupervised sentence embedding learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 671–688, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zhigang Wang, Juanzi Li, Zhiyuan Liu, and Jie Tang. 2016. [Text-enhanced representation learning for knowledge graph](#). In *Proceedings of International joint conference on artificial intelligent (IJCAI)*, pages 4–17.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. [Annotating expressions of opinions and emotions in language](#). *Language resources and evaluation*, 39(2):165–210.
- John Wieting and Kevin Gimpel. 2018. [ParaNMT-50M: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Melbourne, Australia. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Xing Wu, Chaochen Gao, Liangjun Zang, Jizhong Han, Zhongyuan Wang, and Songlin Hu. 2022. [ESimCSE: Enhanced sample building method for contrastive learning of unsupervised sentence embedding](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3898–3907, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Wei Xu, Chris Callison-Burch, and Bill Dolan. 2015. [SemEval-2015 task 1: Paraphrase and semantic similarity in Twitter \(PIT\)](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 1–11, Denver, Colorado. Association for Computational Linguistics.
- Yuanmeng Yan, Rumei Li, Sirui Wang, Fuzheng Zhang, Wei Wu, and Weiran Xu. 2021. [ConSERT: A contrastive framework for self-supervised sentence representation transfer](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5065–5075, Online. Association for Computational Linguistics.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. [Embedding entities and relations for learning and inference in knowledge bases](#). *International Conference on Learning Representations*.
- Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. [KG-BERT: BERT for knowledge graph completion](#). *arXiv preprint arXiv:1909.03193*.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. [From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions](#). *Transactions of the Association for Computational Linguistics*, 2:67–78.

A Joint-Learning with Multiple Datasets

SimCSE (Gao et al., 2021b) found that the positive sentence pairs extracted from *entailment* relation are the most effective supervision signals in contrastive sentence representation learning. We first experiment with different supervision as positive sentence pairs. The averaged STS result is reported in Table 1. The result shows that using *entailment* for positive and *contrastive* for negatives provides the best performance. It outperforms all other supervision, including QQP, Flicker, ParaNMT, and

QNLI datasets. This explains why only NLI dataset is widely used in previous work for supervised sentence representation learning (Conneau et al., 2017; Reimers and Gurevych, 2019; Gao et al., 2021b).

Another question comes to the stage: Can we simultaneously leverage the supervision from the above-mentioned datasets for better-supervised representation learning? An intuitive way to learn from multi-source relational data is to treat all supervised signals as positive samples in contrastive learning. Therefore, we experiment with merging all supervision data and performing contrastive sentence representation learning following Gao et al. (2021b). The result is shown in Table 1. We can tell that merging all sentence pairs for joint training will hurt the performance (77.67 \rightarrow 74.62) of STS datasets compared with a single source for positive sentence pairs.

Based on the above experimental results, we conclude it is not a trivial task to leverage multi-source supervision signals in supervised sentence representation learning. With explicit relation modeling techniques proposed in RSE, we can effectively learn from multi-source relational data and achieve better flexibility and performance in various tasks.

Relation	MRR	Hits@1	Hits@3	Hits@10
Entailment	0.80	0.74	0.85	0.93
Duplicate Question	0.83	0.76	0.88	0.97
Paraphrase	0.97	0.95	0.98	1.00
QA Entailment	0.85	0.80	0.90	0.95
Same Caption	0.56	0.45	0.63	0.79
All	0.81	0.74	0.84	0.92

Table 7: Performance on in-domain relation prediction. We follow the evaluation protocol of link prediction in knowledge graph completion literature.

B In-domain Relation Prediction

As the relation sentence embedding is trained in a contrastive manner, the trained model can also be used to predict the validity of new triples. Therefore, we did further experiments to test the model’s performance on link prediction.

For dataset creation, we use data from five different relations, including *entailment*, *duplicate-question*, *paraphrase*, *same-caption*, and *qa-entailment*. We did not include the prediction for *same-sent* relation as the prediction for the same sentence is a trivial task. We use all available data for training. For validation and testing, we ran-

domly sampled 2,000 sentence triples each.

We perform link prediction experiments following the literature of knowledge graph completion (Sun et al., 2019). The task is to predict the tail entity (sentence) given the head entity (sentence) and relation. It is a retrieval-based task, and performance is measured by the rank of the valid positive samples. Our experimental result is shown in Table 7. From the result, we can see that the *paraphrase* relation is the easiest to predict. It is understandable because the sentences from the paraphrased data are very similar. In contrast, we see that the *same-caption* relation is the hardest to be retrieved. It is because the sentence pairs are collected from the descriptions for the same image. The captions for each image can be generated based on various aspects. The relationship between large number sentence pairs is not very strong. It poses challenges on *same-caption* relations in link prediction.

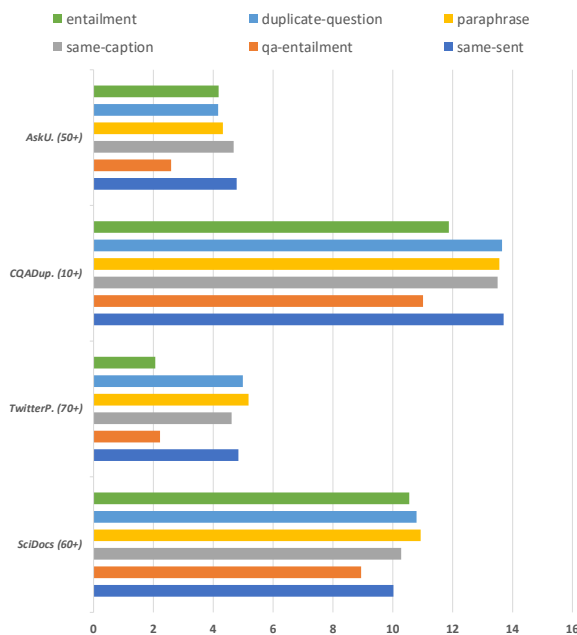


Figure 5: Performance on the USEB benchmark (4 datasets) using different relational similarity scores. The model is $BERT_{base}$.

C Comparison of Relational Score Selection on USEB Tasks

In Section 5, we report the best result on four tasks with the best relational similarity score. We found that different dataset favors different relational similarity score. We look into the detailed performance with all relational similarity scores to further study

the effect. The result is shown in Figure 5.

For AskU. dataset, the unsupervised SimCSE shows the best performance. Similarly, we also found that the unsupervised relation (*same-sent*) demonstrated to be the best relation similarity score to be used in AskU. and CQADup. datasets. TwitterP. dataset is a paraphrase identification task and the similarity score under *paraphrase* relation in RSE shows the best result. In contrast, as no dataset in USEB focuses on question-answer pair retrieval or ranking, the similarity score from *qa-entailment* relation does not lead to a good performance. From the above discussion, we can see that different relational similarity scores will apply to different applications. A wise choice of relational score can lead to better performance, proving that our proposed RSE model can be flexible to various tasks with explicit relation modeling capability.

D Hyperparameter Settings

Table 8 shows the datasets, relations, examples and their applied tasks in our experiments. We train our model for 3 epochs in all experiments and evaluate every 125, 1000, and 500 steps for best checkpoint selection on STS, USEB, and Transfer tasks, respectively. We select the best model based on its performance on the development set. Due to resource limitations, we perform a small-scale grid search for the model hyperparameter settings. The grid search is on batch size $\in \{128, 256, 512\}$ and learning rate $\in \{1e-5, 3e-5, 5e-5\}$. The learning rate for relation embedding is set to $1e-2$ in all experiments because they are randomly initialized parameters without pre-training. All experiments take the representation for [CLS] token as sentence representation. The temperature hyperparameter is set to 0.05, and the maximum sentence length is set to 32. We did not use warm-up or weight decay in all experiments.

E Case Study

To better understand the property of RSE, we present the case study of four sentence pairs and their inferred relational scores in Table 9. The results show that the RSE model can estimate the most likely relationship between sentence pairs. Because the QA entailment relation is very distinct from other relation types, the question-answering pairs are easily distinguished. The QA sentences show the highest score in QA entailment relation, as indicated by the first two cases. Meantime, be-

cause of the similar property between relations of ‘*duplicate question*’, ‘*paraphrase*’ and ‘*same sent*’, from the third case, we can see highly similar sentence pairs have high relational scores in all three relations. It also aligns with our relation study in Figure 4. The last case shows two entailment sentence pairs, and we can see that RSE successfully presents the highest score as an entailment relation.

Dataset	Relationship	Example Sentence Pair	Tasks
MNLI & SNLI	Entailment (r_1)	(s_1) He mostly hangs out with a group of older, Southern black men, who call him Jumper and Black Cat. (s_2) The group of guys he tends to hang out with gave him the nickname Jumper.	STS, USEB, Transfer
QQP	Duplicate Question (r_2)	(s_1) How can I be a good geologist? (s_2) What should I do to be a great geologist?	STS, USEB
ParaNMT	Paraphrase (r_3)	(s_1) he loved that little man , by the way . (s_2) he liked the little boy .	USEB, Transfer
Flicker	Same Caption (r_4)	(s_1) Two men in green shirts are standing in a yard. (s_2) Two friends enjoy time spent together.	USEB
QNLI	QA entailment (r_5)	(s_1) How many alumni does Olin Business School have worldwide? (s_2) Olin has a network of more than 16,000 alumni worldwide.	USEB
WIKI	Same Sent (r_6)	(s_1) A meeting of promoters was also held at Presbyterian Church. (s_2) A meeting of promoters was also held at Presbyterian Church.	USEB

Table 8: Example sentence triples for different relations and their applied tasks.

Sentences	Entailment	Duplicate question	Paraphrase	Same caption	QA entailment	Same sent
“Where is ACL 2023 hold?” “The 61st Annual Meeting of the Association for Computational Linguistics (ACL23) will take place in Toronto, Canada from July 9th to July 14th, 2023.”	0.1639	0.3204	0.3037	0.3468	0.4470	0.3078
“Is Singapore a city or state?” “Singapore is a sunny, tropical island in South-east Asia, off the southern tip of the Malay Peninsula.”	0.4220	0.5640	0.5605	0.5774	0.6599	0.5451
“Giraffes can consume 75 pounds of food a day.” “Giraffes can eat up to 75 pounds of food in a day.”	0.8622	0.9596	0.9546	0.9434	0.9021	0.9569
“Young man breakdancing on a sidewalk in front of many people.” “A man dancing outside.”	0.6680	0.6126	0.6176	0.5606	0.5062	0.6396

Table 9: Case study on sentence pairs and the relation scores inferred by our RSE model.