

# PAUSE: Positive and Annealed Unlabeled Sentence Embedding

Lele Cao<sup>1</sup> Emil Larsson<sup>1,2</sup> Vilhelm von Ehrenheim<sup>1</sup>  
Dhiana Deva Cavalcanti Rocha<sup>1</sup> Anna Martin<sup>1</sup> Sonja Horn<sup>1</sup>

<sup>1</sup>Motherbrain, EQT Group, Stockholm, Sweden

<sup>2</sup>Modulai, Stockholm, Sweden

{lele.cao, vilhelm.vonehrenheim}@eqtpartners.com

{dhiana.deva, anna.martin}@eqtpartners.com

emil@modulai.io sonja@eqtventures.com

## Abstract

Sentence embedding refers to a set of effective and versatile techniques for converting raw text into numerical vector representations that can be used in a wide range of natural language processing (NLP) applications. The majority of these techniques are either supervised or unsupervised. Compared to the unsupervised methods, the supervised ones make less assumptions about optimization objectives and usually achieve better results. However, the training requires a large amount of labeled sentence pairs, which is not available in many industrial scenarios. To that end, we propose a generic and end-to-end approach – PAUSE (Positive and Annealed Unlabeled Sentence Embedding), capable of learning high-quality sentence embeddings from a partially labeled dataset. We experimentally show that PAUSE achieves, and sometimes surpasses, state-of-the-art results using only a small fraction of labeled sentence pairs on various benchmark tasks. When applied to a real industrial use case where labeled samples are scarce, PAUSE encourages us to extend our dataset without the burden of extensive manual annotation work.

## 1 Introduction

A sentence embedding is a numerical representation used to describe the meaning of an entire sentence. Embeddings of this type are becoming increasingly important for many downstream tasks in the language understanding domain, such as similarity or sentiment analysis. Some earlier methods, like GloVe (Pennington et al., 2014), BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) pool directly from underlying token-level embeddings to create a sentence representation. Recently, these pooling strategies have been challenged by various parameterized policies that can be optimized on domain specific tasks. The majority of these are either *unsupervised* or *supervised*. While unsupervised methods only utilize unlabeled sentences,

supervised methods can quickly customize the embeddings by using domain specific labels. As a consequence, supervised methods make less assumptions about optimization objectives and usually achieve better results. However, supervised training requires a large amount of labeled sentence pairs, which is usually unavailable. In many real scenarios, the dataset turns out to be *positive-unlabeled* (i.e. *PU dataset*), where the majority is unlabeled and the rest of the samples are labeled as positive. The methods that enable learning binary classifiers on PU datasets are called *PU learning*. To bridge the gap between supervised and unsupervised approaches, we incorporate state-of-the-art PU learning with the general supervised sentence embedding approaches, proposing a novel method – PAUSE (Positive and Annealed Unlabeled Sentence Embedding)<sup>1</sup>. The main highlights of PAUSE include:

- (1) good sentence embeddings can be learned from datasets with only a few positive labels;
- (2) it can be trained in an end-to-end fashion;
- (3) it can be directly applied to any dual-encoder model architecture;
- (4) it is extended to scenarios with an arbitrary number of classes;
- (5) polynomial annealing of the PU loss is proposed to stabilize the training;
- (6) our experiments show that PAUSE constantly outperforms baseline methods.

## 2 Related Work

Among *unsupervised* sentence embedding methods, some are capable of exploring the relations among

<sup>1</sup>The source code, pre-processed data, and trained models are accessible publicly from: <https://github.com/EQTPartners/pause>

sub-sentences, such as skip-thoughts (Kiros et al., 2015), FastSent (Hill et al., 2016), quick-thoughts (Logeswaran and Lee, 2018) and DiscSent (Jernite et al., 2017). These methods assume that adjacent sentences always have similar semantics. However, not every corpus is long enough, perfectly ordered or coherent enough to fulfill that assumption, which limits their applicable domains. Other unsupervised methods merely focus on the internal structures within each sentence, such as paragraph-vectors (Le and Mikolov, 2014), Doc2VecC (Chen, 2017), Sent2Vec (Pagliardini et al., 2018; Gupta et al., 2019), WMD (Wu et al., 2018), GEM (Yang et al., 2019) and IS-BERT (Zhang et al., 2020). In general, those unsupervised approaches optimize objectives based on assumptions, which limits their embeddings from being adapted towards different applications. Recently, several concurrently proposed methods, such as Yan et al. 2021; Kim et al. 2021; Carlsson et al. 2021; Giorgi et al. 2021, adopt contrastive objectives by constructing different views from the same sentence. Gao et al. 2021 achieved superior results by simply using dropout to create different views.

The *supervised* approaches, on the other hand, are usually (1) trained in an end-to-end manner, (2) following a dual encoder architecture and (3) finetuned from a model pretrained on SNLI (Stanford Natural Language Inference) (Bowman et al., 2015) and Multi-Genre NLI (Williams et al., 2018) datasets. NLI is the task of determining whether a *hypothesis* is true (entailment), false (contradiction), or undetermined (neutral) given a *premise*<sup>2</sup>. The recent representative methods include InferSent (Conneau et al., 2017), USE (Universal Sentence Encoder) variants (Cer et al., 2018; Chidambaram et al., 2019), SBERT (SentenceBERT) (Reimers and Gurevych, 2019) and LaBSE (Language-agnostic BERT Sentence Embedding) (Feng et al., 2020). Built upon pretrained models, they can effectively learn good embeddings from the labeled sentence pairs. However, this approach is not feasible in scenarios where the quantity of annotations is limited.

Rather than purely labeled or unlabeled, in many real scenarios, the dataset turns out to be *positive-unlabeled* (PU), where a small portion of sentence pairs are labeled as positive samples and the rest are unlabeled. To address this type of problem, the gap between supervised and unsupervised methods

<sup>2</sup><http://nlpprogress.com>

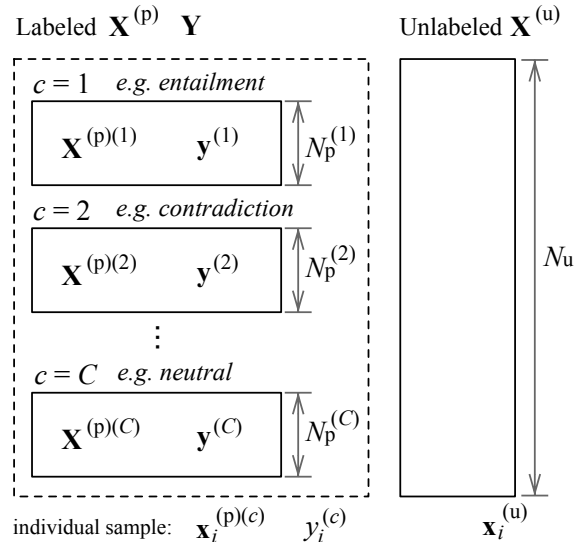


Figure 1: Problem setting and notations.

has to be filled. Levi experimented with incorporating an unsupervised regularization criteria in the supervised loss. Although (Levi, 2018) reported better generalization capability, all of the samples still have to be labeled. Jiang et al. made an early attempt to apply PU learning – particularly for matrix factorization (Yu et al., 2017) – to obtain word embeddings for low-resource languages.

Existing PU learning methods can be divided into three categories based on how unlabeled data is handled. The first category, with methods like (Li and Liu, 2003; Yang et al., 2017), tries to assign labels to unlabeled data in a heuristic-driven and iterative manner which makes the training scattered in steps/phases and hard to implement in practice. The second includes methods like (Liu et al., 2003; Lee and Liu, 2003) that treat unlabeled data as negative with lower confidence. This can be more computationally expensive to tune. The third category, with methods such as uPU (Du Plessis et al., 2014), nnPU (Kiryo et al., 2017), PUbN (Hsieh et al., 2019) and Self-PU (Chen et al., 2020), regards each unlabeled sample as a weighted mixture of being positive and negative. This third category optimizes a so called PU loss, and has recently become dominant due to its wide applicability and end-to-end nature. However, a major limitation is that these algorithms are only applicable to binary classification problems. In this work, we show how to adapt PU learning to effectively learn sentence embeddings from multi-class PU datasets.

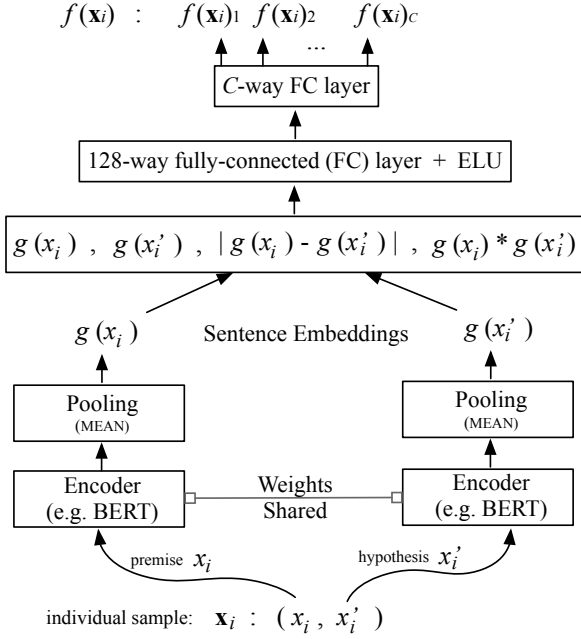


Figure 2: The dual encoder model architecture.

### 3 The Proposed Method

We propose a generic and end-to-end approach to obtain sentence embeddings in a setup that is a generalized natural language inference (GNLI) task. This approach can have (1) any number of classes and (2) the majority of the sentence pairs unlabeled. Let  $\mathbf{X}$  be the set of sentence pairs in the entire dataset; as illustrated in Figure 1,  $\mathbf{X} = \cup_{c=1}^C \mathbf{X}^{(p)(c)} \cup \mathbf{X}^{(u)}$ , where  $C (\geq 2)$  is the total number of entailment classes,  $\mathbf{X}^{(p)(c)}$  denotes the  $N_p^{(c)}$  sentence pairs labeled as the  $c$ -th class, and  $\mathbf{X}^{(u)}$  represents the  $N_u$  unlabeled pairs. For the  $N_p = \sum_{c=1}^C N_p^{(c)}$  labeled pairs, we use  $\mathbf{Y} \in \mathbb{R}^{N_p \times C}$  to denote their mutually exclusive and one-hot encoded labels, hence the binary label for the  $c$ -th entailment class should have the form  $\mathbf{y}^{(c)} \in \mathbb{R}^{N_p^{(c)}}$ . On the individual sample level, we use  $\mathbf{x}_i^{(p)(c)}$ ,  $y_i^{(c)}$ , and  $\mathbf{x}_i^{(u)}$  to denote the  $i$ -th sentence pair that is labeled as class  $c$ , the binary label towards the  $c$ -th class for the  $i$ -th pair, and the  $i$ -th unlabeled sample respectively.

#### 3.1 Dual encoder model architecture

The model architecture of PAUSE follows a dual encoder schema (Figure 2) that is widely adopted in supervised sentence embedding training. Each individual sample  $\mathbf{x}_i$  contains a pair of hypothesis and premise sentences  $(x_i, x'_i)$ , each of which is fed into a pretrained encoder (e.g. BERT). As shown in Figure 2, the two encoders are identical during

the training by sharing their weights. We add a pooling operation to the output of both encoders to obtain the fixed sized sentence embeddings  $g(x_i)$  and  $g(x'_i)$ , and following the empirical suggestion of (Reimers and Gurevych, 2019), we apply the MEAN-strategy (i.e. calculating the average of the encoder output vectors). Once the sentence embeddings are generated, three matching methods are applied to extract relations between  $g(x_i)$  and  $g(x'_i)$ : (1) concatenation of the two vectors, (2) absolute element-wise difference  $|g(x_i) - g(x'_i)|$ , and (3) element-wise product  $g(x_i) * g(x'_i)$ . The results of the three matching methods are then concatenated into a vector, which captures information from both the premise and the hypothesis. This vector is fed into a 128-way fully-connected (FC) layer with ELU (Exponential Linear Unit) activation (Clevert et al., 2016), the output of which is transformed by a  $C$ -way linear FC layer, obtaining the final output  $\mathbf{f}(\mathbf{x}_i) = [f(\mathbf{x}_i)_1, f(\mathbf{x}_i)_2, \dots, f(\mathbf{x}_i)_C]$ .

#### 3.2 Supervised loss

For multi-class and mono-label problems, we calculate cross entropy (CE) loss using the labeled samples:

$$\mathcal{L}_{\text{CE}} = \frac{-1}{CN_p} \sum_{i=1}^{N_p} \sum_{c=1}^C |y_i^{(c)}| \log \left( \frac{e^{f(\mathbf{x}_i)_c}}{\sum_{j=1}^C e^{f(\mathbf{x}_i)_j}} \right). \quad (1)$$

For multi-class multi-label problems, the supervised loss can be binary CE:

$$\mathcal{L}_{\text{CE}} = \frac{-1}{CN_p} \sum_{i=1}^{N_p} \sum_{c=1}^C |y_i^{(c)}| \log \left( \frac{1}{1 + e^{-f(\mathbf{x}_i)_c}} \right). \quad (2)$$

When there is absolutely no negative label for binary classification problems, the supervised loss can safely be ignored.

#### 3.3 Positive unlabeled loss

To leverage unlabeled data in obtaining better sentence embeddings, we turn to state-of-the-art PU learning methods, among which we largely follow (Du Plessis et al., 2014, 2015; Kiryo et al., 2017) due to their effectiveness and simplicity. Recently, Chen et al. proposed an updated version which achieved marginal improvement but with the cost of greatly increased training complexity.

To facilitate computing PU loss, we address each class separately as a binary classification problem, therefore  $y_i^{(c)} \in \{\pm 1\}$ . For each class

$C$ , we define  $p(\mathbf{x}, y^{(c)})$  as the joint density of  $(\mathbf{X}, \mathbf{y}^{(c)})$ ,  $p_p^{(c)}(\mathbf{x}) = p(\mathbf{x}|y^{(c)} = +1)$  as the positive marginal,  $p_n^{(c)}(\mathbf{x}) = p(\mathbf{x}|y^{(c)} = -1)$  as the negative marginal,  $p^{(c)}(\mathbf{x})$  as the unlabeled marginal,  $\pi_p^{(c)} = p(y^{(c)} = +1)$  as the positive prior, and  $\pi_n^{(c)} = 1 - \pi_p^{(c)}$  as the negative prior.

Assuming that we have all samples labeled for the  $c$ -th class, we can easily estimate the error risk  $\mathcal{R}(f)_c$  as

$$\mathcal{R}(f)_c = \mathbb{E}_{(\mathbf{x}, y^{(c)}) \sim p(\mathbf{x}, y^{(c)})}[\ell(f(\mathbf{x})_c, y^{(c)})], \quad (3)$$

where function  $f$  is approximated by the model depicted in Figure 2, and  $\ell : \mathbb{R} \times \{\pm 1\} \rightarrow \mathbb{R}$  is the loss function, such that the value  $\ell(a, b)$  means the loss incurred by predicting an output  $a$  when the ground truth is  $b$ . The feasible  $\ell$  functions can be referred to in (Kiryo et al., 2017). Noticing that  $\mathcal{R}(f)_c$  can be equivalently calculated by summing the error from positive and negative samples (denoted as  $\mathcal{R}_p^+(f)_c$  and  $\mathcal{R}_n^-(f)_c$  respectively):

$$\begin{aligned} \mathcal{R}(f)_c &= \pi_p^{(c)} \mathcal{R}_p^+(f)_c + \pi_n^{(c)} \mathcal{R}_n^-(f)_c \\ &= \pi_p^{(c)} \mathbb{E}_{\mathbf{x} \sim p_p^{(c)}(\mathbf{x})}[\ell^+] \\ &\quad + \pi_n^{(c)} \mathbb{E}_{\mathbf{x} \sim p_n^{(c)}(\mathbf{x})}[\ell^-], \end{aligned} \quad (4)$$

where  $\ell^- = \ell(f(\mathbf{x})_c, -1)$  and  $\ell^+ = \ell(f(\mathbf{x})_c, +1)$ . Since  $p^{(c)}(\mathbf{x}) = \pi_p^{(c)} p_p^{(c)}(\mathbf{x}) + \pi_n^{(c)} p_n^{(c)}(\mathbf{x})$ , we have:

$$\begin{aligned} \mathbb{E}_{\mathbf{x} \sim p^{(c)}(\mathbf{x})}[\ell^-] &= \pi_p^{(c)} \mathbb{E}_{\mathbf{x} \sim p_p^{(c)}(\mathbf{x})}[\ell^-] \\ &\quad + \pi_n^{(c)} \mathbb{E}_{\mathbf{x} \sim p_n^{(c)}(\mathbf{x})}[\ell^-]. \end{aligned} \quad (5)$$

For the sake of simplicity, we denote the terms  $\mathbb{E}_{\mathbf{x} \sim p^{(c)}(\mathbf{x})}[\ell^-]$  and  $\mathbb{E}_{\mathbf{x} \sim p_p^{(c)}(\mathbf{x})}[\ell^-]$  as  $\mathcal{R}_u^-(f)_c$  and  $\mathcal{R}_p^-(f)_c$  respectively. As a result, (5) becomes

$$\mathcal{R}_u^-(f)_c = \pi_p^{(c)} \mathcal{R}_p^-(f)_c + \pi_n^{(c)} \mathcal{R}_n^-(f)_c. \quad (6)$$

By solving (4) and (6), we can eliminate the term  $\mathcal{R}_n^-(f)_c$  in (4):

$$\begin{aligned} \mathcal{R}(f)_c &= \pi_p^{(c)} \mathcal{R}_p^+(f)_c \\ &\quad + \left[ \mathcal{R}_u^-(f)_c - \pi_p^{(c)} \mathcal{R}_p^-(f)_c \right], \end{aligned} \quad (7)$$

where the two terms are called positive risk and negative risk respectively. Kiryo et al. argue that when the value of the negative risk becomes less than zero, it often indicates overfitting. In that circumstance, we empirically choose to drop the

positive risk and optimize reversely in respect to the negative risk term. Hence, in implementation, the error risk for the  $c$ -th class has the form of

$$\begin{aligned} \mathcal{R}(f)_c &= \mathbb{1}\{\mathcal{R}_u^-(f)_c - \pi_p^{(c)} \mathcal{R}_p^-(f)_c \geq 0\} \\ &\quad \times \pi_p^{(c)} \mathcal{R}_p^+(f)_c \\ &\quad + \max\{\mathcal{R}_u^-(f)_c - \pi_p^{(c)} \mathcal{R}_p^-(f)_c, \\ &\quad \pi_p^{(c)} \mathcal{R}_p^-(f)_c - \mathcal{R}_u^-(f)_c\}. \end{aligned} \quad (8)$$

For  $\ell$ , we choose to use sigmoid loss, i.e.  $\ell(a, b) = (1 + e^{ab})^{-1}$ , and we can conveniently calculate  $\mathcal{R}(f)_c$  by plugging in the flowing equations:

$$\begin{aligned} \mathcal{R}_p^+(f)_c &= \frac{1}{N_p^{(c)}} \sum_{i=1}^{N_p^{(c)}} \left(1 + e^{f(\mathbf{x}_i^{(p)})_c}\right)^{-1}, \\ \mathcal{R}_p^-(f)_c &= \frac{1}{N_p^{(c)}} \sum_{i=1}^{N_p^{(c)}} \left(1 + e^{-f(\mathbf{x}_i^{(p)})_c}\right)^{-1}, \quad (9) \\ \mathcal{R}_u^-(f)_c &= \frac{1}{N_u} \sum_{i=1}^{N_u} \left(1 + e^{-f(\mathbf{x}_i^{(u)})_c}\right)^{-1}. \end{aligned}$$

The overall PU loss  $\mathcal{L}_{\text{PU}}$  can be constructed using (8) and (9):

$$\mathcal{L}_{\text{PU}} = \frac{1}{C} \sum_{c=1}^C \mathcal{R}(f)_c \quad (10)$$

### 3.4 Annealed joint optimization

During training, the model is optimized in an end-to-end manner on mini-batches. Every mini-batch is sampled from all subsets of the dataset (cf. Figure 1) according to the relative subset sizes, so that the mini-batches reflect the composition of the entire dataset. Because the initial estimations of positive/negative risks in (7) tend to be inaccurate, simply optimizing both losses (i.e.  $\mathcal{L}_{\text{CE}}$  and  $\mathcal{L}_{\text{PU}}$ ) jointly with the same weights often leads to sub-optimal and unstable solution. This problem is particularly prominent when the dataset is large or the model is highly flexible. As a result, we propose to apply an annealing strategy to the PU loss component when constructing the overall loss:

$$\mathcal{L} = \mathcal{L}_{\text{CE}} + \left(\frac{t}{T}\right)^\alpha \mathcal{L}_{\text{PU}}, \quad (11)$$

where  $T$  denotes the total number of training steps and  $1 \leq t \leq T$  is the elapsed number of steps. The hyper-parameter  $\alpha \geq 2$  controls the annealing speed. We empirically discover that  $\alpha = 3$  usually

offers optimal and stable performance. For binary classification problems ( $C = 1$ ), the overall loss will fall back to  $\mathcal{L}_{PU}$  when there is no negative labels available.

## 4 Experiments

Inspired by previous work (Hill et al., 2016; Reimers and Gurevych, 2019; Zhang et al., 2020), we evaluate PAUSE on STS (Semantic Textual Similarity) and SentEval<sup>3</sup> in both supervised and unsupervised settings. We also show the robustness of PAUSE in a real industrial use case. In our experiments, we test two versions of PAUSE: PAUSE-base (110M parameters) and PAUSE-small (4.4M parameters), which use uncased BERT-base (Devlin et al., 2019) and BERT-small (Turc et al., 2019)<sup>4</sup> as their encoder model, respectively. PAUSE is trained by minimizing (11) with  $\alpha = 3$  (searched 3, 4, and 5) using the Adam optimizer with learning rate 7.5e-5 (searched 1e-3, 1e-4, 7.5e-5, 5e-5, 2.5e-5, 1e-5 and 1e-6). The experiments are carried out on a machine (managed virtual machine instance<sup>5</sup>) with four virtual CPUs (Intel Xeon 2.30GHz), 15GB RAM, and four GPUs (NVIDIA Tesla P100). Since PAUSE requires a large batch size to ensure enough labeled samples from each class in every mini-batch, we use a batch size of 128 and 1,024 for PAUSE-base and PAUSE-small respectively, fully utilizing the GPU capacity.

### 4.1 Unsupervised STS

We first evaluate PAUSE on STS tasks without using any STS data for training. Specifically, we choose the datasets of STS 2012-2016 (Agirre et al., 2012, 2013, 2014, 2015, 2016), STS benchmark (STSb) (Cer et al., 2017), and SICK-Relatedness (SICK-R) (Marelli et al., 2014). These datasets have labels between 0 and 5 indicating the semantic relatedness of sentence pairs. We compare PAUSE with two groups of baselines. The first group is the unsupervised methods, which includes FastSent (Hill et al., 2016), IS-BERT-NLI (Zhang et al., 2020), the average of GloVe embeddings, the average of the last layer representations of BERT, and the [CLS] embedding of BERT. The second group consists of supervised approaches: InferSent-Glove

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

<sup>4</sup><https://github.com/google-research/bert>

<sup>5</sup><https://cloud.google.com/ai-platform/docs/technical-overview#notebooks>

(Conneau et al., 2017), USE (Cer et al., 2018), SBERT, and SRoBERTa (Reimers and Gurevych, 2019). All models are trained on the combination of the SNLI (Bowman et al., 2015) and Multi-Genre NLI (Williams et al., 2018) datasets, which contains one million sentence pairs annotated with three labels ( $C = 3$ ): *entailment*, *contradiction* and *neutral*. PAUSE is trained for 2 epochs with a linear learning rate warm-up over the first 10% of the training steps.

As suggested in (Reimers et al., 2016; Reimers and Gurevych, 2019; Zhang et al., 2020), we calculate the Spearman’s rank correlation between the cosine-similarity of the sentence embeddings and the labels, which is presented in Table 1. The results show that most of the supervised methods achieve superior performances compared to unsupervised ones, which has been previously evidenced by (Hill et al., 2016; Cer et al., 2018; Zhang et al., 2020). PAUSE using the BERT-base encoder (PAUSE-NLI-base) performs much better than the versions using the BERT-small encoder (PAUSE-NLI-small). PAUSE-NLI-base takes on average 220 minutes to complete one epoch of training, while PAUSE-NLI-small takes merely 9 minutes. The post-fix (i.e. 1%, 5%, ..., 100%) in the names of the PAUSE variants indicate the percentage of NLI labels that are used during training. Although the performance monotonically drops when using less labels, this drop remains marginal even when only 50% of the labels are used. PAUSE-NLI-base-100% obtained slightly better results than SBERT-NLI-base. PAUSE-NLI-base-100% and SBERT-NLI-base are trained on the same amount of labeled samples, yet the former achieves slightly better results, probably due to differences in (1) the BERT-base versions<sup>6</sup> and/or (2) the layers following the encoding step. Observing the average results of each PAUSE-NLI-base variant, the model trained on merely 10% of the labels results in a performance about 2% lower than the one relying on all labels. This demonstrates that PAUSE is a label-efficient sentence embedding approach applicable to situations where only a small number of samples are labeled.

In an attempt to train the SBERT-NLI-base model using only 1%, 5%, 10%, and 30% of the labels for 2 epochs, we found that all trials suffered from overfitting in varying degrees. While this

<sup>6</sup>PAUSE uses [https://tfhub.dev/tensorflow/bert\\_en\\_uncased\\_L-12\\_H-768\\_A-12/3](https://tfhub.dev/tensorflow/bert_en_uncased_L-12_H-768_A-12/3)

Model		STS12	STS13	STS14	STS15	STS16	STSb	SICK-R	Avg.
unsupervised	FastSent <sup>*</sup>	-	-	63.00	-	-	-	61.00	-
	Avg. GloVe embeddings <sup>†</sup>	55.14	70.66	59.73	68.25	63.66	58.02	53.76	61.32
	Avg. BERT embeddings <sup>†</sup>	38.78	57.98	57.98	63.15	61.06	46.35	58.40	54.81
	BERT [CLS]-vector <sup>†</sup>	20.16	30.01	20.09	36.88	38.08	16.50	42.63	29.19
	IS-BERT-NLI <sup>‡</sup>	56.77	69.24	61.21	75.23	70.16	69.21	64.25	66.58
supervised	InferSent-Glove <sup>†</sup>	52.86	66.75	62.15	72.77	66.87	68.03	65.65	65.01
	USE <sup>†</sup>	64.49	67.80	64.61	76.83	73.18	74.92	<b>76.69</b>	71.22
	SBERT-NLI-base <sup>†</sup>	70.97	<b>76.53</b>	<b>73.19</b>	<b>79.09</b>	74.30	77.03	72.91	74.89
	SRoBERTa-NLI-base <sup>†</sup>	71.54	72.49	70.80	<b>78.74</b>	73.69	77.77	74.46	74.21
PAUSE variants	PAUSE-NLI-small-100%	67.40	68.47	62.34	67.73	69.08	71.24	71.65	68.27
	PAUSE-NLI-small-70%	65.58	67.60	61.05	67.25	67.85	70.42	70.85	67.23
	PAUSE-NLI-small-50%	65.60	66.95	61.52	66.87	66.48	69.35	70.83	66.80
	PAUSE-NLI-small-30%	64.50	65.91	60.42	67.28	66.42	69.33	70.00	66.27
	PAUSE-NLI-small-10%	60.85	65.47	60.90	66.56	65.99	68.12	67.24	65.02
	PAUSE-NLI-small-5%	59.21	65.32	59.56	66.14	65.93	67.58	65.81	64.22
	PAUSE-NLI-small-1%	57.06	64.00	59.47	64.28	63.25	63.14	61.09	61.76
	PAUSE-NLI-base-100%	<b>73.68</b>	<b>77.08</b>	<b>74.10</b>	<b>78.58</b>	<b>76.58</b>	<b>80.34</b>	<b>74.69</b>	<b>76.44</b>
	PAUSE-NLI-base-70%	<b>73.96</b>	<b>76.26</b>	73.13	77.57	<b>75.74</b>	<b>78.98</b>	74.53	<b>75.74</b>
	PAUSE-NLI-base-50%	<b>73.72</b>	75.89	72.80	77.29	<b>75.40</b>	<b>78.62</b>	<b>74.55</b>	<b>75.47</b>
	PAUSE-NLI-base-30%	73.54	75.15	72.44	77.09	74.43	78.39	73.55	74.94
	PAUSE-NLI-base-10%	71.51	74.46	<b>73.23</b>	77.16	74.67	78.22	73.49	74.68
	PAUSE-NLI-base-5%	67.59	73.76	68.68	73.24	73.47	75.77	72.28	72.11
	PAUSE-NLI-base-1%	62.44	72.01	63.88	69.39	66.32	70.22	69.64	67.70

<sup>\*</sup> The results are extracted from (Hill et al., 2016).

<sup>†</sup> The results are extracted from (Reimers and Gurevych, 2019).

<sup>‡</sup> The results are extracted from (Zhang et al., 2020).

Table 1: Unsupervised evaluation results: Spearman’s rank correlation  $\rho$  between the cosine similarity of sentence embeddings and the labels for STS tasks.  $\rho \times 100$  is reported here. The PAUSE results are averaged from three runs with random seeds. Bold font indicates the top-3 results on each dataset.

problem could be addressed by hyper-parameter optimization and regularization, such alterations would compromise the fairness of the comparison.

## 4.2 Supervised STS benchmark

Similar to (Reimers and Gurevych, 2019; Zhang et al., 2020), we use the STS benchmark (STSb) dataset (Cer et al., 2017) to evaluate the models’ performance on the supervised STS task. STSb includes 8,628 sentence pairs from the categories of *captions*, *news*, and *forums*. The dataset is split into train (5,749), dev (1,500), and test (1,379) subsets. We use the training set to finetune PAUSE (pretrained using partially labeled NLI) using a regression objective function. On the test set, we compute the cosine similarity between each pair of sentences. Since PAUSE obtains better results using BERT-base compared to BERT-small (cf. Table 1), we only report the results for PAUSE-NLI-STSb-base models, which are trained with five random seeds and four epochs.

In Table 2, we compare PAUSE to three categories of baselines: (1) not trained on STSb at all, (2) only trained on STSb, and (3) first trained on the fully labeled NLI, then finetuned on STSb. It is clear that finetuning on STSb greatly improves the model performance and pretraining on NLI further uplifts the performance slightly. Using merely 10% to 70% of NLI labels, PAUSE manages to achieve results comparable to the baselines that use all NLI labels. Another interesting finding is that when pretraining PAUSE on less than 5% of the labels, the performance becomes inferior to directly finetuning on STSb. This suggests that pretraining PAUSE with too few labeled samples may result in embeddings that are hard to finetune in downstream regression tasks. In addition, we observe a clear trend in Table 2: the standard deviation increases when less labels are used, which is also observed in the unsupervised experiments.

Model	$\rho \times 100$
<i>Not trained on STSb</i>	
Avg. BERT embeddings	46.35
SBERT-NLI-base	77.03
SROBERTa-NLI-base	<b>77.77</b>
<i>Only trained on STSb</i>	
IS-BERT-STSB (ft)	74.25 $\pm$ 0.94
IS-BERT-STSB (ssl+ft)	84.84 $\pm$ 0.43
SBERT-STSB-base	84.67 $\pm$ 0.19
SROBERTa-STSB-base	<b>84.92</b> $\pm$ 0.34
<i>Trained on NLI, then on STSb</i>	
SBERT-NLI-STSB-base	<b>85.35</b> $\pm$ 0.17
SROBERTa-NLI-STSB-base	84.79 $\pm$ 0.38
PAUSE-NLI-STSB-base-100%	84.83 $\pm$ 0.30
<i>Trained on partially labeled NLI, then on STSb</i>	
PAUSE-NLI-STSB-base-70%	<b>84.02</b> $\pm$ 0.29
PAUSE-NLI-STSB-base-50%	83.86 $\pm$ 0.35
PAUSE-NLI-STSB-base-30%	83.53 $\pm$ 0.37
PAUSE-NLI-STSB-base-10%	83.39 $\pm$ 0.44
PAUSE-NLI-STSB-base-5%	81.77 $\pm$ 0.41
PAUSE-NLI-STSB-base-1%	73.04 $\pm$ 0.58

Table 2: The supervised evaluation results (Spearman’s rank correlation  $\rho$ ) on the STS benchmark (STSb) test set. The results of the non-PAUSE baselines are extracted from (Reimers and Gurevych, 2019) and (Zhang et al., 2020).

### 4.3 SentEval: domain specific tasks

In order to give an impression of the quality of our sentence embeddings for various domain specific tasks, we choose to evaluate PAUSE on seven SentEval tasks (Conneau and Kiela, 2018): (1) **TREC** - fine grained question type classification (Li and Roth, 2002), (2) **CR** - sentiment prediction of customer product reviews (Hu and Liu, 2004), (3) **MRPC** - Microsoft Research Paraphrase Corpus from parallel news sources (Dolan et al., 2004), (4) **SUBJ** - Subjectivity prediction of sentences from movie reviews and plot summaries (Pang and Lee, 2004), (5) **MR** - sentiment prediction for movie reviews (Pang and Lee, 2005), (6) **MPQA** - opinion polarity classification (Wiebe et al., 2005) and (7) **SST** - binary sentiment analysis (Socher et al., 2013). Unlike (Devlin et al., 2019) and (Zhang et al., 2020) that finetune the encoder on these tasks, we directly use the sentence embeddings from PAUSE-NLI-base models (cf. Section 4.1) as features for a logistic regression classifier that is trained in a 10-fold cross-validation setup, where the prediction accuracy is computed for the test

fold.

The results can be found in Table 3, where the top-3 results on each task are presented in bold face. Largely speaking, the sentence embeddings from SBERT and PAUSE successfully capture domain specific information with the exception of the TREC task where pretraining on question-answering data (e.g. USE) seems to be beneficial. In Table 1, we have observed poor results from Avg. BERT embeddings, BERT [CLS]-vector, PAUSE-NLI-base-5% and PAUSE-NLI-base-1%. However, on the selected SentEval tasks, they all achieve decent results and the performance of PAUSE does not even degrade as we use significantly less labels. This inconsistency can be explained by how we measure the model performance: on STS datasets we calculated the cosine-similarity between sentence embeddings, which treats all dimensions indifferently, while SentEval fits a logistic regression classifier to the embeddings allowing different dimensions to have different impact on the classifier’s output. As a result, cosine-similarity can only be relied on when the sentence embeddings are finetuned on related datasets with a large number of labeled samples. When the sentence embeddings are directly used as input features for training discriminative models on downstream tasks, finetuning on NLI only gives approximately 1~2% performance uplift. Moreover, PAUSE appears to be unaffected by a drastic decrease in labeled samples, which is consistent with the results of IS-BERT-task.

We also notice that unsupervised PAUSE-NLI-base-100% performs better than SBERT-NLI-base in Table 1, yet this is not the case for supervised fine-tuning (Table 2 and 3). This might be a consequence of several differences: (1) PAUSE uses a newer version of the pretrained BERT-base model<sup>6</sup> compared to (Reimers and Gurevych, 2019), (2) PAUSE has an extra term  $g(x_i) * g(x'_i)$  when extracting relations between two sentences, as seen in Figure 2, and (3) PAUSE treats the NLI sentence pairs with ambiguous/conflicting labels as unlabeled samples that are utilized by the PU loss term during model optimization.

### 4.4 Use case: finding similar companies

In this section, we will discuss the potential of PAUSE on a real industrial use case from EQT Group<sup>7</sup>. The EQT investment professionals use

<sup>7</sup>EQT Group is a global investment organization (<https://eqtgroup.com>).

Model	MR	CR	SUBJ	MPQA	SST	TREC	MRPC	Avg.
Unigram-TFIDF*	73.7	79.2	90.3	82.4	-	85.0	73.6	-
SDAE*	74.6	78.0	90.8	86.9	-	78.4	73.7	-
ParagraphVec DBOW*	60.2	66.9	76.3	70.7	-	59.4	72.9	-
FastSent*	70.8	78.4	88.7	80.6	-	76.8	72.2	-
SkipThought*	76.5	80.1	93.6	87.1	82.0	<b>92.2</b>	73.0	83.50
Avg. GloVe embeddings†	77.25	78.30	91.17	87.85	80.18	83.0	72.87	81.52
Avg. BERT embeddings†	78.66	86.25	<b>94.37</b>	88.66	84.40	<b>92.8</b>	69.54	84.94
BERT [CLS]-vector†	78.68	84.85	94.21	88.23	84.13	91.4	71.13	84.66
IS-BERT-task‡	81.09	87.18	<b>94.96</b>	88.75	85.96	88.64	74.24	<b>85.91</b>
InferSent-Glove†	<b>81.57</b>	86.54	92.50	<b>90.38</b>	84.18	88.2	75.77	85.59
USE†	80.09	85.19	93.98	86.70	86.38	<b>93.2</b>	70.14	85.10
SBERT-NLI-base†	<b>83.64</b>	<b>89.43</b>	<b>94.39</b>	<b>89.86</b>	<b>88.96</b>	89.6	<b>76.00</b>	<b>87.41</b>
PAUSE-NLI-base-100%	<b>81.40</b>	<b>87.87</b>	93.64	88.85	<b>88.41</b>	84.4	<b>75.83</b>	<b>85.77</b>
PAUSE-NLI-base-70%	80.04	87.74	92.65	<b>89.42</b>	<b>87.70</b>	84.2	75.71	85.35
PAUSE-NLI-base-50%	80.27	<b>87.52</b>	93.09	89.14	<b>87.70</b>	81.6	74.61	84.85
PAUSE-NLI-base-30%	80.42	87.31	92.99	89.04	87.26	84.8	74.09	85.13
PAUSE-NLI-base-10%	79.93	86.49	93.43	88.60	86.44	86.2	73.68	84.97
PAUSE-NLI-base-5%	80.06	87.05	93.37	88.86	85.61	86.0	75.01	85.14
PAUSE-NLI-base-1%	79.99	85.83	93.33	88.16	86.66	86.4	<b>77.10</b>	85.35

\* The results are extracted from (Hill et al., 2016).

† The results are extracted from (Reimers and Gurevych, 2019).

‡ The results are extracted from (Zhang et al., 2020); IS-BERT-task is finetuned on each of the task-specific dataset (without label) to produce sentence embeddings, which are then used for training downstream classifiers.

Table 3: Accuracy of test fold (10 fold cross validation) on domain specific SentEval tasks. Sentence embeddings from different models are used as features to train a logistic regression classifier.

Motherbrain<sup>8</sup> to accomplish tasks such as deal sourcing, market analysis and metrics benchmarking to name a few. One concrete use case is competitor mapping, where the Motherbrain users track competitors for companies they are studying (a.k.a. anchor companies). Competitors are defined as the companies running a business similar to that of the anchor company. In our database, there are more than eight million companies with textual descriptions and only 12,326 of these have annotations indicating their relatedness to other companies, forming 166,832 company pairs with noisy binary labels indicating whether they are similar (55,139 pairs) or dissimilar (111,693 pairs). Table 4 demonstrates our similar company (SC) dataset containing these annotated company pairs, which is divided into train (156,925), dev (4,922) and test (4,985) sets. A BERT-base encoder is finetuned using the PAUSE approach on the training set. The finetuned encoder model is served via Tensorflow Serving<sup>9</sup>,

which is called by an Apache Beam<sup>10</sup> job encoding the incoming company descriptions into company embeddings in a streaming fashion. The company embeddings are then indexed by Elasticsearch<sup>11</sup> to support fast similarity search initiated by platform users.

To benchmark PAUSE in this setting, we trained models using different percentages of labeled samples (100%, 50%, 10%, 5%). Table 5 shows that it is sufficient to have 10% of the samples labeled and still reach high accuracy, precision, and recall. When all samples are labeled, the accuracy is only increased by around 3%. In practice, these results encourage us to extend our dataset without the burden of manually labeling all samples. We also speculate that increasing the size of the dataset while maintaining the balance between labeled and unlabeled samples could improve performance further. Essentially, this implies that we can achieve results close to that of a fully labeled dataset with a fraction of the manual annotation work.

<sup>8</sup>Motherbrain is EQT’s proprietary investment platform driven by diversified big data and cutting-edge algorithms. (<https://eqtgroup.com/motherbrain>, <https://eqtventures.com/motherbrain>)

<sup>9</sup><https://github.com/tensorflow/serving>

<sup>10</sup><https://beam.apache.org>

<sup>11</sup><https://www.elastic.co>



ID	Anchor Company	Candidate Company	label
1	Owner and operator of data centers in UK ... distribute data in data centers and the global digital economy.	Independent co-location / data center provider in Slovenia.	1
2	Provider of human-computer interaction technology designed to ... documentation is common to both CX and ES.	Developer of a mobile stock-trading application designed to make ... enabling traders to discover and invest in markets without a hassle.	0
3	Provide a reliable and fast veterinary diagnostic service ... We intend to be your partner in the daily medical diagnosis!	Provider of laboratory services. We care about ... Biochemical and haematological examinations are available 24 hours a day.	1
4	Manufacturer of smart electric scooters designed to offer ... noise-free scooters that run on electricity.	A peer-to-peer rental marketplace which allows people to rent spare items ... can be borrowed within minutes.	0
... ... 166,832 rows (company pairs) in total			

Table 4: The first four samples (shortened and anonymized) of the similar company (SC) dataset used to train company embeddings for identifying similar companies. Label=0 and 1 represent dissimilarity and similarity, respectively.

Model	Acc.	Precision	Recall
PAUSE-SC-100%	76.13	77.55	61.87
PAUSE-SC-50%	75.27	74.30	64.41
PAUSE-SC-10%	73.35	73.29	60.58
PAUSE-SC-5%	64.40	63.80	54.07

Table 5: The performance of PAUSE (using the BERT-base encoder) trained on the similar company (SC) dataset using different percentages of labeled samples.

## 5 Conclusions and Future Work

In this work, we attempt to bridge the gap between supervised and unsupervised sentence embedding techniques, proposing PAUSE – a generic and end-to-end sentence embedding approach that exploits the labels and explores the unlabeled sentence pairs simultaneously. PAUSE trained on NLI datasets achieves state-of-the-art results on unsupervised STS tasks, and also performs well on many downstream domain-specific tasks. In all of our experiments, we observe that PAUSE keeps performing well with a reduced number of labeled samples, as long as more than 5-10% of the dataset is labeled. This indicates that PAUSE is a label-efficient sentence embedding approach that can be effectively applied to datasets where only a small part is labeled while the rest remains unlabeled. We also demonstrate that PAUSE helps lower the labeling requirement for an industrial use case aimed at encoding company descriptions. In that sense, PAUSE pushes the application boundary of sentence embeddings to include many more real-world scenarios where labeled samples are scarce. The possible extensions of this work include (1) augmenting the the labels with dropout, (2) experimenting with contrastive supervised loss, and (3) exploring how PAUSE can be extended with con-

textual sentence embeddings.

## Acknowledgements

EQT Group and the Motherbrain team have provided great support along the journey of accomplishing this work; particularly, we would like to appreciate the insights/support of all sorts from (alphabetically ordered) Alex Patow, Andjela Kusmuk, Andreas Beccau, Andrey Melentyev, Anton Andersson Andrejic, Anton Ask Åström Daniel Ström, Daniel Wroblewski, Elin Bäcklund, Emil Broman, Emma Sjöström, Erik Ferm, Filip Byrén, Guillermo Rodas, Hannes Ingelhart, Henrik Landgren, Joar Wandborg, Love Larsson, Lucas Magnum, Niklas Skaar, Peter Finnman, Sarah Bernelind, Olof Hernell, Pietro Casella, Richard Stahl, Sebastian Lindblom, Sven Törnkvist, Ylva Lundegård. Additionally, the first author would also like to thank Xiaolong Liu (Intel Labs) and Xiaoxue Li (Shanghai University of Finance and Economics) for the initial discussion around PU learning methodologies and their applications.

## 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)*, pages 252–263.
- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2014. [SemEval-2014 task 10: Multilingual semantic textual similarity](#). In *Proceedings of the*

- 8th International Workshop on Semantic Evaluation (SemEval)*, pages 81–91.
- Eneko Agirre, Carmen Banea, Daniel Cer, Mona Diab, Aitor Gonzalez Agirre, Rada Mihalcea, German Rigau Claramunt, and Janyce Wiebe. 2016. [SemEval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval)*, pages 497–511.
- Eneko Agirre, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. 2012. [Semeval-2012 task 6: A pilot on semantic textual similarity](#). In *Proceedings of the 1st Joint Conference on Lexical and Computational Semantics (SEM) and the 6th International Workshop on Semantic Evaluation (SemEval)*, pages 385–393.
- Eneko Agirre, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and Weiwei Guo. 2013. [SEM 2013 shared task: Semantic textual similarity](#). In *Proceedings of the 2nd Joint Conference on Lexical and Computational Semantics (SEM)*, pages 32–43.
- 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 (EMNLP)*, pages 632–642.
- Fredrik Carlsson, Amaru Cuba Gyllensten, Evangelia Gogoulou, Erik Ylipää Hellqvist, and Magnus Sahlgren. 2021. [Semantic re-tuning with contrastive tension](#). In *Proceedings of the 10th International Conference on Learning Representations*.
- Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and cross-lingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval)*, pages 1–14.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. [Universal sentence encoder for english](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP): System Demonstrations*, pages 169–174.
- Minmin Chen. 2017. [Efficient vector representation for documents through corruption](#). In *Proceedings of the 6th International Conference on Learning Representations (ICLR)*.
- Xuxi Chen, Wuyang Chen, Tianlong Chen, Ye Yuan, Chen Gong, Kewei Chen, and Zhangyang Wang. 2020. [Self-PU: Self boosted and calibrated positive-unlabeled training](#). In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 1510–1519.
- Muthu Chidambaram, Yinfei Yang, Daniel Cer, Steve Yuan, Yunhsuan Sung, Brian Strope, and Ray Kurzweil. 2019. [Learning cross-lingual sentence representations via a multi-task dual-encoder model](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP)*, pages 250–259.
- Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. 2016. [Fast and accurate deep network learning by exponential linear units \(ELUs\)](#). In *Proceedings of the 4th International Conference on Learning Representations (ICLR)*.
- Alexis Conneau and Douwe Kiela. 2018. [SentEval: An evaluation toolkit for universal sentence representations](#). In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC)*.
- 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 (EMNLP)*, pages 670–680.
- 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 (NAACL-HLT)*, pages 4171–4186.
- William B Dolan, Chris Quirk, and Chris Brockett. 2004. [Unsupervised construction of large paraphrase corpora: Exploiting massively parallel news sources](#). In *Proceedings of the 20th International Conference on Computational Linguistics (COLING)*, pages 350–356.
- Marthinus Christoffel Du Plessis, Gang Niu, and Masashi Sugiyama. 2014. [Analysis of learning from positive and unlabeled data](#). In *Proceedings of the 28th International Conference on Neural Information Processing Systems (NeurIPS)*, pages 703–711.
- Marthinus Christoffel Du Plessis, Gang Niu, and Masashi Sugiyama. 2015. [Convex formulation for learning from positive and unlabeled data](#). In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pages 1386–1394.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. [Language-agnostic BERT sentence embedding](#). *CoRR*, abs/2007.01852.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

- John Giorgi, Osvald Nitski, Bo Wang, and Gary Bader. 2021. [DeCLUTR: Deep contrastive learning for unsupervised textual representations](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, pages 879–895.
- Prakhar Gupta, Matteo Pagliardini, and Martin Jaggi. 2019. [Better word embeddings by disentangling contextual n-gram information](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 933–939.
- Felix Hill, Kyunghyun Cho, and Anna Korhonen. 2016. [Learning distributed representations of sentences from unlabelled data](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 1367–1377.
- Yu-Guan Hsieh, Gang Niu, and Masashi Sugiyama. 2019. [Classification from positive, unlabeled and biased negative data](#). In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 2820–2829.
- Minqing Hu and Bing Liu. 2004. [Mining and summarizing customer reviews](#). In *Proceedings of the 10th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177.
- Yacine Jernite, Samuel R. Bowman, and David A. Sonntag. 2017. [Discourse-based objectives for fast unsupervised sentence representation learning](#). *CoRR*, abs/1705.00557.
- Chao Jiang, Hsiang-Fu Yu, Cho-Jui Hsieh, and Kai-Wei Chang. 2018. [Learning word embeddings for low-resource languages by pu learning](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 1024–1034.
- Taeuk Kim, Kang Min Yoo, and Sang-goo Lee. 2021. [Self-guided contrastive learning for BERT sentence representations](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, pages 2528–2540.
- Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2015. [Skip-thought vectors](#). In *Proceedings of the 28th International Conference on Neural Information Processing Systems (NeurIPS)*, pages 3294–3302.
- Ryuichi Kiryo, Gang Niu, Marthinus C du Plessis, and Masashi Sugiyama. 2017. [Positive-unlabeled learning with non-negative risk estimator](#). In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS)*, pages 1674–1684.
- Quoc Le and Tomas Mikolov. 2014. [Distributed representations of sentences and documents](#). In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 1188–1196.
- Wee Sun Lee and Bing Liu. 2003. [Learning with positive and unlabeled examples using weighted logistic regression](#). In *Proceedings of the 20th International Conference on Machine Learning (ICML)*, pages 448–455.
- Gil Levi. 2018. [Connecting supervised and unsupervised sentence embeddings](#). In *Proceedings of The 3rd Workshop on Representation Learning for NLP (RepL4NLP)*, pages 79–83.
- Xiaoli Li and Bing Liu. 2003. [Learning to classify texts using positive and unlabeled data](#). In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 587–592.
- Xin Li and Dan Roth. 2002. [Learning question classifiers](#). In *Proceedings of the 19th International Conference on Computational Linguistics (COLING)*.
- Bing Liu, Yang Dai, Xiaoli Li, Wee Sun Lee, and Philip S Yu. 2003. [Building text classifiers using positive and unlabeled examples](#). In *the 3rd IEEE International Conference on Data Mining (ICDM)*, pages 179–186.
- 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](#). *CoRR*, abs/1907.11692.
- Lajanugen Logeswaran and Honglak Lee. 2018. [An efficient framework for learning sentence representations](#). In *Proceedings of the 6th International Conference on Learning Representations (ICLR)*.
- 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 9th International Conference on Language Resources and Evaluation (LREC)*, pages 216–223.
- 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 (NAACL-HLT)*, pages 528–540.
- 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)*, pages 271–278.

- 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)*, pages 115–124.
- 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.
- Nils Reimers, Philip Beyer, and Iryna Gurevych. 2016. [Task-oriented intrinsic evaluation of semantic textual similarity](#). In *Proceedings of the 26th International Conference on Computational Linguistics (COLING): Technical Papers*, pages 87–96.
- 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 3973–3983.
- 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 (EMNLP)*, pages 1631–1642.
- Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Well-read students learn better: On the importance of pre-training compact models](#). *CoRR*, abs/1908.08962.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. [Annotating expressions of opinions and emotions in language](#). *Language resources and evaluation*, 39(2):165–210.
- 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 (NAACL-HLT)*, pages 1112–1122.
- Lingfei Wu, Ian En-Hsu Yen, Kun Xu, Fangli Xu, Avinash Balakrishnan, Pin-Yu Chen, Pradeep Ravikumar, and Michael J Witbrock. 2018. [Word mover’s embedding: From Word2Vec to document embedding](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4524–4534.
- 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 (ACL-IJCNLP)*, pages 5065–5075.
- Pengyi Yang, Wei Liu, and Jean Yang. 2017. [Positive unlabeled learning via wrapper-based adaptive sampling](#). In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3273–3279.
- Ziyi Yang, Chenguang Zhu, and Weizhu Chen. 2019. [Parameter-free sentence embedding via orthogonal basis](#). 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 638–648.
- Hsiang-Fu Yu, Mikhail Bilenko, and Chih-Jen Lin. 2017. [Selection of negative samples for one-class matrix factorization](#). In *Proceedings of the 2017 SIAM International Conference on Data Mining (SDM)*, pages 363–371.
- Yan Zhang, Ruidan He, Zuozhu Liu, Kwan Hui Lim, and Lidong Bing. 2020. [An unsupervised sentence embedding method by mutual information maximization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1601–1610.