

DeCLUTR: Deep Contrastive Learning for Unsupervised Textual Representations

John Giorgi^{1,5,6} Osvald Nitski^{2,7} Bo Wang^{1,4,6,7,†} Gary Bader^{1,3,5,†}

¹Department of Computer Science, University of Toronto

²Faculty of Applied Science and Engineering, University of Toronto

³Department of Molecular Genetics, University of Toronto

⁴Department of Laboratory Medicine and Pathobiology, University of Toronto

⁵Terrence Donnelly Centre for Cellular & Biomolecular Research

⁶Vector Institute for Artificial Intelligence

⁷Peter Munk Cardiac Center, University Health Network

†Co-senior authors

{john.giorgi, osvald.nitski, gary.bader}@mail.utoronto.ca
bowang@vectorinstitute.ai

Abstract

Sentence embeddings are an important component of many natural language processing (NLP) systems. Like word embeddings, sentence embeddings are typically learned on large text corpora and then transferred to various downstream tasks, such as clustering and retrieval. Unlike word embeddings, the highest performing solutions for learning sentence embeddings require labelled data, limiting their usefulness to languages and domains where labelled data is abundant. In this paper, we present DeCLUTR: Deep Contrastive Learning for Unsupervised Textual Representations. Inspired by recent advances in deep metric learning (DML), we carefully design a self-supervised objective for learning universal sentence embeddings that does not require labelled training data. When used to extend the pretraining of transformer-based language models, our approach closes the performance gap between unsupervised and supervised pretraining for universal sentence encoders. Importantly, our experiments suggest that the quality of the learned embeddings scale with both the number of trainable parameters and the amount of unlabelled training data. Our code and pretrained models are publicly available and can be easily adapted to new domains or used to embed unseen text.¹

1 Introduction

Due to the limited amount of labelled training data available for many natural language processing (NLP) tasks, transfer learning has become ubiquitous (Ruder et al., 2019). For some time, transfer learning in NLP was limited to pretrained word embeddings (Mikolov et al., 2013; Pennington et al.,

2014). Recent work has demonstrated strong transfer task performance using pretrained *sentence* embeddings. These fixed-length vectors, often referred to as “universal” sentence embeddings, are typically learned on large corpora and then transferred to various downstream tasks, such as clustering (e.g. topic modelling) and retrieval (e.g. semantic search). Indeed, sentence embeddings have become an area of focus, and many supervised (Conneau et al., 2017), semi-supervised (Subramanian et al., 2018; Phang et al., 2018; Cer et al., 2018; Reimers and Gurevych, 2019) and unsupervised (Le and Mikolov, 2014; Jernite et al., 2017; Kiros et al., 2015; Hill et al., 2016; Logeswaran and Lee, 2018) approaches have been proposed. However, the highest performing solutions require labelled data, limiting their usefulness to languages and domains where labelled data is abundant. Therefore, closing the performance gap between unsupervised and supervised universal sentence embedding methods is an important goal.

Pretraining transformer-based language models has become the primary method for learning textual representations from unlabelled corpora (Radford et al., 2018; Devlin et al., 2019; Dai et al., 2019; Yang et al., 2019; Liu et al., 2019; Clark et al., 2020). This success has primarily been driven by masked language modelling (MLM). This self-supervised *token*-level objective requires the model to predict the identity of some randomly masked tokens from the input sequence. In addition to MLM, some of these models have mechanisms for learning *sentence*-level embeddings via self-supervision. In BERT (Devlin et al., 2019), a special classification token is prepended to every input sequence, and its representation is used in a binary classifi-

¹<https://github.com/JohnGiorgi/DeCLUTR>

cation task to predict whether one textual segment follows another in the training corpus, denoted Next Sentence Prediction (NSP). However, recent work has called into question the effectiveness of NSP (Conneau and Lample, 2019; You et al., 1904; Joshi et al., 2020). In RoBERTa (Liu et al., 2019), the authors demonstrated that removing NSP during pretraining leads to unchanged or even slightly improved performance on downstream sentence-level tasks (including semantic text similarity and natural language inference). In ALBERT (Lan et al., 2020), the authors hypothesize that NSP conflates *topic* prediction and *coherence* prediction, and instead propose a Sentence-Order Prediction objective (SOP), suggesting that it better models inter-sentence coherence. In preliminary evaluations, we found that neither objective produces good universal sentence embeddings (see Appendix A). Thus, we propose a simple but effective self-supervised, sentence-level objective inspired by recent advances in metric learning.

Metric learning is a type of representation learning that aims to learn an embedding space where the vector representations of similar data are mapped close together, and vice versa (Lowe, 1995; Mika et al., 1999; Xing et al., 2002). In computer vision (CV), deep metric learning (DML) has been widely used for learning visual representations (Wohlgart and Lepetit, 2015; Wen et al., 2016; Zhang and Saligrama, 2016; Bucher et al., 2016; Leal-Taixé et al., 2016; Tao et al., 2016; Yuan et al., 2020; He et al., 2018; Grabner et al., 2018; Yelamathi et al., 2018; Yu et al., 2018). Generally speaking, DML is approached as follows: a “pretext” task (often self-supervised, e.g. colourization or inpainting) is carefully designed and used to train deep neural networks to generate useful feature representations. Here, “useful” means a representation that is easily adaptable to other downstream tasks, unknown at training time. Downstream tasks (e.g. object recognition) are then used to evaluate the quality of the learned features (independent of the model that produced them), often by training a linear classifier on the task using these features as input. The most successful approach to date has been to design a pretext task for learning with a pair-based contrastive loss function. For a given *anchor* data point, contrastive losses attempt to make the distance between the anchor and some *positive* data points (those that are similar) smaller than the distance between the anchor and some *neg-*

ative data points (those that are dissimilar) (Hadsell et al., 2006). The highest-performing methods generate anchor-positive pairs by randomly augmenting the same image (e.g. using crops, flips and colour distortions); anchor-negative pairs are randomly chosen, augmented views of different images (Bachman et al., 2019; Tian et al., 2020; He et al., 2020; Chen et al., 2020). In fact, Kong et al., 2020 demonstrate that the MLM and NSP objectives are also instances of contrastive learning.

Inspired by this approach, we propose a self-supervised, contrastive objective that can be used to pretrain a sentence encoder. Our objective learns universal sentence embeddings by training an encoder to minimize the distance between the embeddings of textual segments randomly sampled from nearby in the same document. We demonstrate our objective’s effectiveness by using it to extend the pretraining of a transformer-based language model and obtain state-of-the-art results on SentEval (Conneau and Kiela, 2018) – a benchmark of 28 tasks designed to evaluate universal sentence embeddings. Our primary contributions are:

- We propose a self-supervised sentence-level objective that can be used alongside MLM to pretrain transformer-based language models, inducing generalized embeddings for sentence- and paragraph-length text without any labelled data (subsection 5.1).
- We perform extensive ablations to determine which factors are important for learning high-quality embeddings (subsection 5.2).
- We demonstrate that the quality of the learned embeddings scale with model and data size. Therefore, performance can likely be improved simply by collecting more unlabelled text or using a larger encoder (subsection 5.3).
- We open-source our solution and provide detailed instructions for training it on new data or embedding unseen text.²

2 Related Work

Previous works on universal sentence embeddings can be broadly grouped by whether or not they use labelled data in their pretraining step(s), which we refer to simply as *supervised* or *semi-supervised* and *unsupervised*, respectively.

²<https://github.com/JohnGiorgi/DeCLUTR>

Supervised or semi-supervised The highest performing universal sentence encoders are pretrained on the human-labelled natural language inference (NLI) datasets Stanford NLI (SNLI) (Bowman et al., 2015) and MultiNLI (Williams et al., 2018). NLI is the task of classifying a pair of sentences (denoted the “hypothesis” and the “premise”) into one of three relationships: *entailment*, *contradiction* or *neutral*. The effectiveness of NLI for training universal sentence encoders was demonstrated by the supervised method InferSent (Conneau et al., 2017). Universal Sentence Encoder (USE) (Cer et al., 2018) is semi-supervised, augmenting an unsupervised, Skip-Thoughts-like task (Kiros et al. 2015, see section 2) with supervised training on the SNLI corpus. The recently published Sentence Transformers (Reimers and Gurevych, 2019) method fine-tunes pretrained, transformer-based language models like BERT (Devlin et al., 2019) using labelled NLI datasets.

Unsupervised Skip-Thoughts (Kiros et al., 2015) and FastSent (Hill et al., 2016) are popular unsupervised techniques that learn sentence embeddings by using an encoding of a sentence to predict words in neighbouring sentences. However, in addition to being computationally expensive, this generative objective forces the model to reconstruct the *surface* form of a sentence, which may capture information irrelevant to the *meaning* of a sentence. QuickThoughts (Logeswaran and Lee, 2018) addresses these shortcomings with a simple discriminative objective; given a sentence and its context (adjacent sentences), it learns sentence representations by training a classifier to distinguish context sentences from non-context sentences. The unifying theme of unsupervised approaches is that they exploit the “distributional hypothesis”, namely that the meaning of a word (and by extension, a sentence) is characterized by the word context in which it appears.

Our overall approach is most similar to Sentence Transformers – we extend the pretraining of a transformer-based language model to produce useful sentence embeddings – but our proposed objective is self-supervised. Removing the dependence on labelled data allows us to exploit the vast amount of unlabelled text on the web without being restricted to languages or domains where labelled data is plentiful (e.g. English Wikipedia). Our objective most closely resembles QuickThoughts; some distinctions include: we relax our sampling to




Figure 1: Overview of the self-supervised contrastive objective. (A) For each document d in a minibatch of size N , we sample A anchor spans per document and P positive spans per anchor. For simplicity, we illustrate the case where $A = P = 1$ and denote the anchor-positive span pair as s_i, s_j . Both spans are fed through the same encoder $f(\cdot)$ and pooler $g(\cdot)$ to produce the corresponding embeddings $e_i = g(f(s_i)), e_j = g(f(s_j))$. The encoder and pooler are trained to minimize the distance between embeddings via a contrastive prediction task, where the other embeddings in a minibatch are treated as negatives (omitted here for simplicity). (B) Positive spans can overlap with, be adjacent to or be subsumed by the sampled anchor span. (C) The length of anchors and positives are randomly sampled from beta distributions, skewed toward longer and shorter spans, respectively.

textual *segments* of up to paragraph length (rather than natural *sentences*), we sample one or more positive segments per anchor (rather than strictly one), and we allow these segments to be adjacent, overlapping or subsuming (rather than strictly adjacent; see Figure 1, B).

3 Model

3.1 Self-supervised contrastive loss

Our method learns textual representations via a contrastive loss by maximizing agreement between textual segments (referred to as “spans” in the rest of the paper) sampled from nearby in the same document. Illustrated in Figure 1, this approach comprises the following components:

- A data loading step randomly samples paired anchor-positive spans from each document in a minibatch of size N . Let A be the number of anchor spans sampled per document, P be the number of positive spans sampled per anchor and $i \in \{1 \dots AN\}$ be the index of an arbitrary anchor span. We denote an anchor span

and its corresponding $p \in \{1 \dots P\}$ positive spans as \mathbf{s}_i and \mathbf{s}_{i+pAN} respectively. This procedure is designed to maximize the chance of sampling semantically similar anchor-positive pairs (see subsection 3.2).

- An encoder $f(\cdot)$ maps each token in the input spans to an embedding. Although our method places no constraints on the choice of encoder, we chose $f(\cdot)$ to be a transformer-based language model, as this represents the state-of-the-art for text encoders (see subsection 3.3).
- A pooler $g(\cdot)$ maps the encoded spans $f(\mathbf{s}_i)$ and $f(\mathbf{s}_{i+pAN})$ to fixed-length embeddings $\mathbf{e}_i = g(f(\mathbf{s}_i))$ and its corresponding mean positive embedding

$$\mathbf{e}_{i+AN} = \frac{1}{P} \sum_{p=1}^P g(f(\mathbf{s}_{i+pAN}))$$

Similar to Reimers and Gurevych 2019, we found that choosing $g(\cdot)$ to be the mean of the token-level embeddings (referred to as “mean pooling” in the rest of the paper) performs well (see Appendix, Table 4). We pair each anchor embedding with the mean of multiple positive embeddings. This strategy was proposed by Saunshi et al. 2019, who demonstrated theoretical and empirical improvements compared to using a single positive example for each anchor.

- A contrastive loss function defined for a contrastive prediction task. Given a set of embedded spans $\{\mathbf{e}_k\}$ including a positive pair of examples \mathbf{e}_i and \mathbf{e}_{i+AN} , the contrastive prediction task aims to identify \mathbf{e}_{i+AN} in $\{\mathbf{e}_k\}_{k \neq i}$ for a given \mathbf{e}_i

$$\ell(i, j) = -\log \frac{\exp(\text{sim}(\mathbf{e}_i, \mathbf{e}_j)/\tau)}{\sum_{k=1}^{2AN} \mathbb{1}_{[i \neq k]} \cdot \exp(\text{sim}(\mathbf{e}_i, \mathbf{e}_k)/\tau)}$$

where $\text{sim}(\mathbf{u}, \mathbf{v}) = \mathbf{u}^T \mathbf{v} / (\|\mathbf{u}\|_2 \|\mathbf{v}\|_2)$ denotes the cosine similarity of two vectors \mathbf{u} and \mathbf{v} , $\mathbb{1}_{[i \neq k]} \in \{0, 1\}$ is an indicator function evaluating to 1 if $i \neq k$, and $\tau > 0$ denotes the temperature hyperparameter.

During training, we randomly sample minibatches of N documents from the train set and define the contrastive prediction task on anchor-positive pairs $\mathbf{e}_i, \mathbf{e}_{i+AN}$ derived from the N documents, resulting in $2AN$ data points. As proposed in (Sohn, 2016), we treat the other $2(AN - 1)$ instances within a minibatch as negative examples. The cost function takes the following form

$$\mathcal{L}_{\text{contrastive}} = \sum_{i=1}^{AN} \ell(i, i + AN) + \ell(i + AN, i)$$

This is the InfoNCE loss used in previous works (Sohn, 2016; Wu et al., 2018; Oord et al., 2018) and denoted normalized temperature-scale cross-entropy loss or “NT-Xent” in (Chen et al., 2020). To embed text with a trained model, we simply pass batches of tokenized text through the model, without sampling spans. Therefore, the computational cost of our method at test time is the cost of the encoder, $f(\cdot)$, plus the cost of the pooler, $g(\cdot)$, which is negligible when using mean pooling.

3.2 Span sampling

We start by choosing a minimum and maximum span length; in this paper, $\ell_{\min} = 32$ and $\ell_{\max} = 512$, the maximum input size for many pretrained transformers. Next, a document d is tokenized to produce a sequence of n tokens $\mathbf{x}^d = (x_1, x_2 \dots x_n)$. To sample an anchor span \mathbf{s}_i from \mathbf{x}^d , we first sample its length ℓ_{anchor} from a beta distribution and then randomly (uniformly) sample its starting position s_i^{start}

$$\begin{aligned} \ell_{\text{anchor}} &= \lfloor p_{\text{anchor}} \times (\ell_{\max} - \ell_{\min}) + \ell_{\min} \rfloor \\ s_i^{\text{start}} &\sim \{0, \dots, n - \ell_{\text{anchor}}\} \\ s_i^{\text{end}} &= s_i^{\text{start}} + \ell_{\text{anchor}} \\ \mathbf{s}_i &= \mathbf{x}_{s_i^{\text{start}}:s_i^{\text{end}}}^d \end{aligned}$$

We then sample $p \in \{1 \dots P\}$ corresponding positive spans \mathbf{s}_{i+pAN} independently following a similar procedure

$$\begin{aligned} \ell_{\text{positive}} &= \lfloor p_{\text{positive}} \times (\ell_{\max} - \ell_{\min}) + \ell_{\min} \rfloor \\ s_{i+pAN}^{\text{start}} &\sim \{s_i^{\text{start}} - \ell_{\text{positive}}, \dots, s_i^{\text{end}}\} \\ s_{i+pAN}^{\text{end}} &= s_{i+pAN}^{\text{start}} + \ell_{\text{positive}} \\ \mathbf{s}_{i+pAN} &= \mathbf{x}_{s_{i+pAN}^{\text{start}}:s_{i+pAN}^{\text{end}}}^d \end{aligned}$$

where $p_{\text{anchor}} \sim \text{Beta}(\alpha = 4, \beta = 2)$, which skews anchor sampling towards longer spans, and $p_{\text{positive}} \sim \text{Beta}(\alpha = 2, \beta = 4)$, which skews positive sampling towards shorter spans (Figure 1, C). In practice, we restrict the sampling of anchor spans from the same document such that they are a minimum of $2 * \ell_{\text{max}}$ tokens apart. In Appendix B, we show examples of text that has been sampled by our method. We note several carefully considered decisions in the design of our sampling procedure:

- Sampling span lengths from a distribution clipped at $\ell_{\text{min}} = 32$ and $\ell_{\text{max}} = 512$ encourages the model to produce good embeddings for text ranging from sentence- to paragraph-length. At test time, we expect our model to be able to embed up-to paragraph-length texts.
- We found that sampling longer lengths for the anchor span than the positive spans improves performance in downstream tasks (we did not find performance to be sensitive to the specific choice of α and β). The rationale for this is twofold. First, it enables the model to learn global-to-local view prediction as in (Hjelm et al., 2019; Bachman et al., 2019; Chen et al., 2020) (referred to as “subsumed view” in Figure 1, B). Second, when $P > 1$, it encourages diversity among positives spans by lowering the amount of repeated text.
- Sampling positives nearby to the anchor exploits the distributional hypothesis and increases the chances of sampling valid (i.e. semantically similar) anchor-positive pairs.
- By sampling multiple anchors per document, each anchor-positive pair is contrasted against both *easy* negatives (anchors and positives sampled from *other* documents in a minibatch) and *hard* negatives (anchors and positives sampled from the *same* document).

In conclusion, the sampling procedure produces three types of positives: positives that partially overlap with the anchor, positives adjacent to the anchor, and positives subsumed by the anchor (Figure 1, B) and two types of negatives: easy negatives sampled from a different document than the anchor, and hard negatives sampled from the same document as the anchor. Thus, our stochastically generated training set and contrastive loss implicitly define a family of predictive tasks which can be

used to train a model, independent of any specific encoder architecture.

3.3 Continued MLM pretraining

We use our objective to extend the pretraining of a transformer-based language model (Vaswani et al., 2017), as this represents the state-of-the-art encoder in NLP. We implement the MLM objective as described in (Devlin et al., 2019) on each anchor span in a minibatch and sum the losses from the MLM and contrastive objectives before backpropagating

$$\mathcal{L} = \mathcal{L}_{\text{contrastive}} + \mathcal{L}_{\text{MLM}}$$

This is similar to existing pretraining strategies, where an MLM loss is paired with a sentence-level loss such as NSP (Devlin et al., 2019) or SOP (Lan et al., 2020). To make the computational requirements feasible, we do not train from scratch, but rather we continue training a model that has been pretrained with the MLM objective. Specifically, we use both RoBERTa-base (Liu et al., 2019) and DistilRoBERTa (Sanh et al., 2019) (a distilled version of RoBERTa-base) in our experiments. In the rest of the paper, we refer to our method as DeCLUTR-small (when extending DistilRoBERTa pretraining) and DeCLUTR-base (when extending RoBERTa-base pretraining).

4 Experimental setup

4.1 Dataset, training, and implementation

Dataset We collected all documents with a minimum token length of 2048 from OpenWebText (Gokaslan and Cohen, 2019) an open-access subset of the WebText corpus (Radford et al., 2019), yielding 497,868 documents in total. For reference, Google’s USE was trained on 570,000 human-labelled sentence pairs from the SNLI dataset (among other unlabelled datasets). InferSent and Sentence Transformer models were trained on both SNLI and MultiNLI, a total of 1 million human-labelled sentence pairs.

Implementation We implemented our model in PyTorch (Paszke et al., 2017) using AllenNLP (Gardner et al., 2018). We used the NT-Xent loss function implemented by the PyTorch Metric Learning library (Musgrave et al., 2019) and the pretrained transformer architecture and weights from the Transformers library (Wolf et al., 2020). All models were trained on up to four NVIDIA Tesla V100 16 or 32GB GPUs.

Table 1: Trainable model parameter counts and sentence embedding dimensions. DeCLUTR-small and DeCLUTR-base are pretrained DistilRoBERTa and RoBERTa-base models respectively after continued pretraining with our method.

Model	Parameters	Embedding dim.
<i>Bag-of-words (BoW) baselines</i>		
GloVe	–	300
fastText	–	300
<i>Supervised and semi-supervised</i>		
InferSent	38M	4096
Universal Sentence Encoder	147M	512
Sentence Transformers	125M	768
<i>Unsupervised</i>		
QuickThoughts	73M	4800
DeCLUTR-small	82M	768
DeCLUTR-base	125M	768

Training Unless specified otherwise, we train for one to three epochs over the 497,868 documents with a minibatch size of 16 and a temperature $\tau = 5 \times 10^{-2}$ using the AdamW optimizer (Loshchilov and Hutter, 2019) with a learning rate (LR) of 5×10^{-5} and a weight decay of 0.1. For every document in a minibatch, we sample two anchor spans ($A = 2$) and two positive spans per anchor ($P = 2$). We use the Slanted Triangular LR scheduler (Howard and Ruder, 2018) with a number of train steps equal to training instances and a cut fraction of 0.1. The remaining hyperparameters of the underlying pretrained transformer (i.e. DistilRoBERTa or RoBERTa-base) are left at their defaults. All gradients are scaled to a vector norm of 1.0 before backpropagating. Hyperparameters were tuned on the SentEval validation sets.

4.2 Evaluation

We evaluate all methods on the SentEval benchmark, a widely-used toolkit for evaluating general-purpose, fixed-length sentence representations. SentEval is divided into 18 *downstream* tasks – representative NLP tasks such as sentiment analysis, natural language inference, paraphrase detection and image-caption retrieval – and ten *probing* tasks, which are designed to evaluate what linguistic properties are encoded in a sentence representation. We report scores obtained by our model and the relevant baselines on the downstream and probing tasks using the SentEval toolkit³ with default parameters (see Appendix C for details). Note that

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

all the supervised approaches we compare to are trained on the SNLI corpus, which is included as a downstream task in SentEval. To avoid train-test contamination, we compute average downstream scores without considering SNLI when comparing to these approaches in Table 2.

4.2.1 Baselines

We compare to the highest performing, most popular sentence embedding methods: InferSent, Google’s USE and Sentence Transformers. For InferSent, we compare to the latest model.⁴ We use the latest “large” USE model⁵, as it is most similar in terms of architecture and number of parameters to DeCLUTR-base. For Sentence Transformers, we compare to “roberta-base-nli-mean-tokens”⁶, which, like DeCLUTR-base, uses the RoBERTa-base architecture and pretrained weights. The only difference is each method’s extended pretraining strategy. We include the performance of averaged GloVe⁷ and fastText⁸ word vectors as weak baselines. Trainable model parameter counts and sentence embedding dimensions are listed in Table 1. Despite our best efforts, we could not evaluate the pretrained QuickThought models against the full SentEval benchmark. We cite the scores from the paper directly. Finally, we evaluate the pretrained transformer model’s performance before it is subjected to training with our contrastive objective, denoted “Transformer-*”. We use mean pooling on the pretrained transformers token-level output to produce sentence embeddings – the same pooling strategy used in our method.

5 Results

In subsection 5.1, we compare the performance of our model against the relevant baselines. In the remaining sections, we explore which components contribute to the quality of the learned embeddings.

5.1 Comparison to baselines

Downstream task performance Compared to the underlying pretrained models DistilRoBERTa

⁴<https://dl.fbaipublicfiles.com/infersent/infersent2.pkl>

⁵<https://tfhub.dev/google/universal-sentence-encoder-large/5>

⁶https://www.sbert.net/docs/pretrained_models.html

⁷<http://nlp.stanford.edu/data/glove.840B.300d.zip>

⁸<https://dl.fbaipublicfiles.com/fasttext/vectors-english/crawl-300d-2M.vec.zip>

Table 2: Results on the downstream tasks from the test set of SentEval. QuickThoughts scores are taken directly from (Logeswaran and Lee, 2018). USE: Google’s Universal Sentence Encoder. Transformer-small and Transformer-base are pretrained DistilRoBERTa and RoBERTa-base models respectively, using mean pooling. DeCLUTR-small and DeCLUTR-base are pretrained DistilRoBERTa and RoBERTa-base models respectively after continued pretraining with our method. Average scores across all tasks, excluding SNLI, are shown in the top half of the table. Bold: best scores. Δ : difference to DeCLUTR-base average score. \uparrow and \downarrow denote increased or decreased performance with respect to the underlying pretrained model. *: Unsupervised evaluations.

Model	CR	MR	MPQA	SUBJ	SST2	SST5	TREC	MRPC	SNLI	Avg.	Δ
<i>Bag-of-words (BoW) weak baselines</i>											
GloVe	78.78	77.70	87.76	91.25	80.29	44.48	83.00	73.39/81.45	65.85	65.47	-13.63
fastText	79.18	78.45	87.88	91.53	82.15	45.16	83.60	74.49/82.44	68.79	68.56	-10.54
<i>Supervised and semi-supervised</i>											
InferSent	84.37	79.42	89.04	93.03	84.24	45.34	90.80	76.35/83.48	84.16	76.00	-3.10
USE	85.70	79.38	88.89	93.11	84.90	46.11	95.00	72.41/82.01	83.25	78.89	-0.21
Sent. Transformers	90.78	84.98	88.72	92.67	90.55	52.76	87.40	76.64/82.99	84.18	77.19	-1.91
<i>Unsupervised</i>											
QuickThoughts	86.00	82.40	90.20	94.80	87.60	–	92.40	76.90/84.00	–	–	–
Transformer-small	86.60	82.12	87.04	94.77	88.03	49.50	91.60	74.55/81.75	71.88	72.58	-6.52
Transformer-base	88.19	84.35	86.49	95.28	89.46	51.27	93.20	74.20/81.44	72.19	72.70	-6.40
DeCLUTR-small	87.52 \uparrow	82.79 \uparrow	87.87 \uparrow	94.96 \uparrow	87.64 \downarrow	48.42 \downarrow	90.80 \downarrow	75.36/82.70 \uparrow	73.59 \uparrow	77.50 \uparrow	-1.60
DeCLUTR-base	90.68 \uparrow	85.16 \uparrow	88.52 \uparrow	95.78 \uparrow	90.01 \uparrow	51.18 \downarrow	93.20 \uparrow	74.61/82.65 \uparrow	74.74 \uparrow	79.10 \uparrow	–
Model	SICK-E	SICK-R	STS-B	COCO	STS12*	STS13*	STS14*	STS15*	STS16*		
GloVe	78.89	72.30	62.86	0.40	53.44	51.24	55.71	59.62	57.93	–	–
fastText	79.01	72.98	68.26	0.40	58.85	58.83	63.42	69.05	68.24	–	–
InferSent	86.30	83.06	78.48	65.84	62.90	56.08	66.36	74.01	72.89	–	–
USE	85.37	81.53	81.50	62.42	68.87	71.70	72.76	83.88	82.78	–	–
Sent. Transformers	82.97	79.17	74.28	60.96	64.10	65.63	69.80	74.71	72.85	–	–
QuickThoughts	–	–	–	60.55	–	–	–	–	–	–	–
Transformer-small	81.96	77.51	70.31	60.48	53.99	45.53	57.23	65.57	63.51	–	–
Transformer-base	80.29	76.84	69.62	60.14	53.28	46.10	56.17	64.69	62.79	–	–
DeCLUTR-small	83.46 \uparrow	77.66 \uparrow	77.51 \uparrow	60.85 \uparrow	63.66 \uparrow	68.93 \uparrow	70.40 \uparrow	78.25 \uparrow	77.74 \uparrow	–	–
DeCLUTR-base	83.84 \uparrow	78.62 \uparrow	79.39 \uparrow	62.35 \uparrow	63.56 \uparrow	72.58 \uparrow	71.70 \uparrow	79.95 \uparrow	79.59 \uparrow	–	–

Table 3: Results on the probing tasks from the test set of SentEval. USE: Google’s Universal Sentence Encoder. Transformer-small and Transformer-base are pretrained DistilRoBERTa and RoBERTa-base models respectively, using mean pooling. DeCLUTR-small and DeCLUTR-base are pretrained DistilRoBERTa and RoBERTa-base models respectively after continued pretraining with our method. Bold: best scores. \uparrow and \downarrow denote increased or decreased performance with respect to the underlying pretrained model.

Model	SentLen	WC	TreeDepth	TopConst	BShift	Tense	SubjNum	ObjNum	SOMO	CoordInv	Avg.
<i>Bag-of-words (BoW) weak baselines</i>											
GloVe	57.82	81.10	31.41	62.70	49.74	83.58	78.39	76.31	49.55	53.62	62.42
fastText	55.46	82.10	32.74	63.32	50.16	86.68	79.75	79.81	50.21	51.41	63.16
<i>Supervised and semi-supervised</i>											
InferSent	78.76	89.50	37.72	80.16	61.41	88.56	86.83	83.91	52.11	66.88	72.58
USE	73.14	69.44	30.87	73.27	58.88	83.81	80.34	79.14	56.97	61.13	66.70
Sent. Transformers	69.21	51.79	30.08	50.38	69.70	83.02	79.74	77.85	60.10	60.33	63.22
<i>Unsupervised</i>											
Transformer-small	88.62	65.00	40.87	75.38	88.63	87.84	86.68	84.17	63.75	64.78	74.57
Transformer-base	81.96	59.67	38.84	74.02	90.08	88.59	85.51	83.33	68.54	71.32	74.19
DeCLUTR-small (ours)	88.85 \uparrow	74.87 \uparrow	38.48 \downarrow	75.17 \downarrow	86.12 \downarrow	88.71 \uparrow	86.31 \downarrow	84.30 \uparrow	61.27 \downarrow	62.98 \downarrow	74.71 \uparrow
DeCLUTR-base (ours)	84.62 \uparrow	68.98 \uparrow	38.35 \downarrow	74.78 \uparrow	87.85 \downarrow	88.82 \uparrow	86.56 \uparrow	83.88 \uparrow	65.08 \downarrow	67.54 \downarrow	74.65 \uparrow

and RoBERTa-base, DeCLUTR-small and DeCLUTR-base obtain large boosts in average downstream performance, +4% and +6% respectively (Table 2). DeCLUTR-base leads to improved or equivalent performance for every downstream task but one (SST5) and DeCLUTR-small for all but three (SST2, SST5 and TREC). Compared

to existing methods, DeCLUTR-base matches or even outperforms average performance without using any hand-labelled training data. Surprisingly, we also find that DeCLUTR-small outperforms Sentence Transformers while using \sim 34% less trainable parameters.

Probing task performance With the exception of InferSent, existing methods perform poorly on the probing tasks of SentEval (Table 3). Sentence Transformers, which begins with a pretrained transformer model and fine-tunes it on NLI datasets, scores approximately 10% lower on the probing tasks than the model it fine-tunes. In contrast, both DeCLUTR-small and DeCLUTR-base perform comparably to the underlying pretrained model in terms of average performance. We note that the purpose of the probing tasks is *not* the development of ad-hoc models that attain top performance on them (Conneau et al., 2018). However, it is still interesting to note that high downstream task performance can be obtained without sacrificing probing task performance. Furthermore, these results suggest that fine-tuning transformer-based language models on NLI datasets may discard some of the linguistic information captured by the pretrained model’s weights. We suspect that the inclusion of MLM in our training objective is responsible for DeCLUTR’s relatively high performance on the probing tasks.

Supervised vs. unsupervised downstream tasks

The downstream evaluation of SentEval includes supervised and unsupervised tasks. In the unsupervised tasks, the embeddings of the method to evaluate are used as-is without any further training (see Appendix C for details). Interestingly, we find that USE performs particularly well across the unsupervised evaluations in SentEval (tasks marked with a * in Table 2). Given the similarity of the USE architecture to Sentence Transformers and DeCLUTR and the similarity of its supervised NLI training objective to InferSent and Sentence Transformers, we suspect the most likely cause is one or more of its additional training objectives. These include a conversational response prediction task (Henderson et al., 2017) and a Skip-Thoughts (Kiros et al., 2015) like task.

5.2 Ablation of the sampling procedure

We ablate several components of the sampling procedure, including the number of anchors sampled per document A , the number of positives sampled per anchor P , and the sampling strategy for those positives (Figure 2). We note that when $A = 2$, the model is trained on twice the number of spans and twice the effective batch size ($2AN$, where N is the number of documents in a minibatch) as compared to when $A = 1$. To control for this, all experi-




Figure 2: Effect of the number of anchor spans sampled per document (a), the number of positive spans sampled per anchor (b), and the sampling strategy (c). Averaged downstream task scores are reported from the validation set of SentEval. Performance is computed over a grid of hyperparameters and plotted as a distribution. The grid is defined by all permutations of number of anchors $A = \{1, 2\}$, number of positives $P = \{1, 2, 4\}$, temperatures $\tau = \{5 \times 10^{-3}, 1 \times 10^{-2}, 5 \times 10^{-2}\}$ and learning rates $\alpha = \{5 \times 10^{-5}, 1 \times 10^{-4}\}$. $P = 4$ is omitted for DeCLUTR-base as these experiments did not fit into GPU memory.

ments where $A = 1$ are trained for two epochs (twice the number of epochs as when $A = 2$) and for two times the minibatch size ($2N$). Thus, both sets of experiments are trained on the same number of spans and the same effective batch size ($4N$), and the only difference is the number of anchors sampled per document (A).

We find that sampling multiple anchors per document has a large positive impact on the quality of learned embeddings. We hypothesize this is because the difficulty of the contrastive objective increases when $A > 1$. Recall that a minibatch is composed of random documents, and each anchor-positive pair sampled from a document is contrasted against all other anchor-positive pairs in the minibatch. When $A > 1$, anchor-positive pairs will be contrasted against other anchors and positives from the same document, increasing the difficulty of the contrastive objective, thus leading to better representations. We also find that a positive sampling strategy that allows positives to be adjacent to *and* subsumed by the anchor outperforms a strategy that only allows adjacent *or* subsuming views, suggesting that the information captured by these views is complementary. Finally, we note that sampling multiple positives per anchor ($P > 1$) has minimal impact on performance. This is in contrast to (Saunshi et al., 2019), who found both theoretical and empirical improvements when




Figure 3: Effect of training objective, train set size and model capacity on SentEval performance. DeCLUTR-small has 6 layers and ~ 82 M parameters. DeCLUTR-base has 12 layers and ~ 125 M parameters. Averaged downstream task scores are reported from the validation set of SentEval. 100% corresponds to 1 epoch of training with all 497,868 documents from our OpenWebText subset.

multiple positives are averaged and paired with a given anchor.

5.3 Training objective, train set size and model capacity

To determine the importance of the training objectives, train set size, and model capacity, we trained two sizes of the model with 10% to 100% (1 full epoch) of the train set (Figure 3). Pretraining the model with both the MLM and contrastive objectives improves performance over training with either objective alone. Including MLM alongside the contrastive objective leads to monotonic improvement as the train set size is increased. We hypothesize that including the MLM loss acts as a form of regularization, preventing the weights of the pretrained model (which itself was trained with an MLM loss) from diverging too dramatically, a phenomenon known as “catastrophic forgetting” (McCloskey and Cohen, 1989; Ratcliff, 1990). These results suggest that the quality of embeddings learned by our approach scale in terms of model capacity and train set size; because the training method is completely self-supervised, scaling the train set would simply involve collecting more unlabelled text.

6 Discussion and conclusion

In this paper, we proposed a self-supervised objective for learning universal sentence embeddings. Our objective does not require labelled training data and is applicable to any text encoder. We demonstrated the effectiveness of our objective by evaluating the learned embeddings on the SentE-

val benchmark, which contains a total of 28 tasks designed to evaluate the transferability and linguistic properties of sentence representations. When used to extend the pretraining of a transformer-based language model, our self-supervised objective closes the performance gap with existing methods that require human-labelled training data. Our experiments suggest that the learned embeddings’ quality can be further improved by increasing the model and train set size. Together, these results demonstrate the effectiveness and feasibility of replacing hand-labelled data with carefully designed self-supervised objectives for learning universal sentence embeddings. We release our model and code publicly in the hopes that it will be extended to new domains and non-English languages.

Acknowledgments

This research was enabled in part by support provided by Compute Ontario (<https://computeontario.ca/>), Compute Canada (www.computecanada.ca) and the CIFAR AI Chairs Program and partially funded by the US National Institutes of Health (NIH) [U41 HG006623, U41 HG003751].

References

- 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.
- Philip Bachman, R. Devon Hjelm, and William Buchwalter. 2019. [Learning representations by maximizing mutual information across views](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 15509–15519.
- 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.
- Maxime Bucher, Stéphane Herbin, and Frédéric Jurie. 2016. [Improving semantic embedding consistency by metric learning for zero-shot classification](#). In *European Conference on Computer Vision*, pages 730–746. Springer.
- 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 E. Hinton. 2020. [A simple framework for contrastive learning of visual representations](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: pre-training text encoders as discriminators rather than generators](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- 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.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. [What you can cram into a single \$\mathbb{R}^d\$ vector: Probing sentence embeddings for linguistic properties](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136, Melbourne, Australia. Association for Computational Linguistics.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 7057–7067.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- 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.
- Bill Dolan, Chris Quirk, and Chris Brockett. 2004. [Unsupervised construction of large paraphrase corpora: Exploiting massively parallel news sources](#). In *COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics*, pages 350–356, Geneva, Switzerland. COLING.

- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [AllenNLP: A deep semantic natural language processing platform](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. Association for Computational Linguistics.
- Aaron Gokaslan and Vanya Cohen. 2019. Openweb-text corpus. <http://Skylion007.github.io/OpenWebTextCorpus>.
- Alexander Grabner, Peter M. Roth, and Vincent Lepetit. 2018. [3d pose estimation and 3d model retrieval for objects in the wild](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3022–3031. IEEE Computer Society.
- Raia Hadsell, Sumit Chopra, and Yann LeCun. 2006. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742. Ieee.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2020. [Momentum contrast for unsupervised visual representation learning](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 9726–9735. IEEE.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society.
- Xinwei He, Yang Zhou, Zhichao Zhou, Song Bai, and Xiang Bai. 2018. [Triplet-center loss for multi-view 3d object retrieval](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 1945–1954. IEEE Computer Society.
- Matthew Henderson, Rami Al-Rfou, Brian Strope, Yun-Hsuan Sung, László Lukács, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. 2017. Efficient natural language response suggestion for smart reply. *arXiv preprint arXiv:1705.00652*.
- 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*, pages 1367–1377, San Diego, California. Association for Computational Linguistics.
- R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. 2019. [Learning deep representations by mutual information estimation and maximization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- 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.
- Yacine Jernite, Samuel R. Bowman, and David A. Sontag. 2017. [Discourse-based objectives for fast unsupervised sentence representation learning](#). *CoRR*, abs/1705.00557.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Skip-thought vectors](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 3294–3302.
- Lingpeng Kong, Cyprien de Masson d’Autume, Lei Yu, Wang Ling, Zihang Dai, and Dani Yogatama. 2020. [A mutual information maximization perspective of language representation learning](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [ALBERT: A lite BERT for self-supervised learning of language representations](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Quoc V. Le and Tomas Mikolov. 2014. [Distributed representations of sentences and documents](#). In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, volume 32 of *JMLR Workshop and Conference Proceedings*, pages 1188–1196. JMLR.org.
- Laura Leal-Taixe, Cristian Canton-Ferrer, and Konrad Schindler. 2016. Learning by tracking: Siamese cnn for robust target association. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 33–40.

- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- 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*.
- Lajanugen Logeswaran and Honglak Lee. 2018. An efficient framework for learning sentence representations. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- David G Lowe. 1995. Similarity metric learning for a variable-kernel classifier. *Neural computation*, 7(1):72–85.
- 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).
- Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- Sebastian Mika, Gunnar Ratsch, Jason Weston, Bernhard Scholkopf, and Klaus-Robert Mullers. 1999. Fisher discriminant analysis with kernels. In *Neural networks for signal processing IX: Proceedings of the 1999 IEEE signal processing society workshop (cat. no. 98th8468)*, pages 41–48. Ieee.
- Tomás Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3111–3119.
- Kevin Musgrave, Ser-Nam Lim, and Serge Belongie. 2019. Pytorch metric learning. <https://github.com/KevinMusgrave/pytorch-metric-learning>.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- 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.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*.
- 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.
- Jason Phang, Thibault Févry, and Samuel R. Bowman. 2018. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *CoRR*, abs/1811.01088.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *URL https://s3-us-west-2.amazonaws.com/openai-assets/researchcovers/languageunsupervised/language_understanding_paper.pdf*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Roger Ratcliff. 1990. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review*, 97(2):285.
- 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.
- Sebastian Ruder, Matthew E. Peters, Swabha Swayamdipta, and Thomas Wolf. 2019. Transfer learning in natural language processing. In

- Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorials*, pages 15–18, Minneapolis, Minnesota. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. 2019. [A theoretical analysis of contrastive unsupervised representation learning](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 5628–5637. PMLR.
- 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.
- Kihyuk Sohn. 2016. [Improved deep metric learning with multi-class n-pair loss objective](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1849–1857.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J. Pal. 2018. [Learning general purpose distributed sentence representations via large scale multi-task learning](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Ran Tao, Efstratios Gavves, and Arnold W. M. Smeulders. 2016. [Siamese instance search for tracking](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 1420–1429. IEEE Computer Society.
- Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2020. Contrastive multiview coding. In *ECCV*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- 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 R. Bowman. 2019. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. 2016. A discriminative feature learning approach for deep face recognition. In *European conference on computer vision*, pages 499–515. Springer.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language resources and evaluation*, 39(2-3):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, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Paul Wohlhart and Vincent Lepetit. 2015. [Learning descriptors for object recognition and 3d pose estimation](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 3109–3118. IEEE Computer Society.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. 2018. [Unsupervised feature learning via non-parametric instance discrimination](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3733–3742. IEEE Computer Society.
- Eric P. Xing, Andrew Y. Ng, Michael I. Jordan, and Stuart J. Russell. 2002. [Distance metric learning with application to clustering with side-information](#). In *Advances in Neural Information Processing Systems*

15 [Neural Information Processing Systems, NIPS 2002, December 9-14, 2002, Vancouver, British Columbia, Canada], pages 505–512. MIT Press.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5754–5764.

Sasi Kiran Yelamarthi, Shiva Krishna Reddy, Ashish Mishra, and Anurag Mittal. 2018. A zero-shot framework for sketch based image retrieval. In *European Conference on Computer Vision*, pages 316–333. Springer.

Yang You, Jing Li, Jonathan Hseu, Xiaodan Song, James Demmel, and C Hsieh. 1904. Reducing bert pre-training time from 3 days to 76 minutes. [corr abs/1904.00962](#) (2019).

Rui Yu, Zhiyong Dou, Song Bai, Zhaoxiang Zhang, Yongchao Xu, and Xiang Bai. 2018. Hard-aware point-to-set deep metric for person re-identification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 188–204.

Ye Yuan, Wuyang Chen, Yang Yang, and Zhangyang Wang. 2020. In defense of the triplet loss again: Learning robust person re-identification with fast approximated triplet loss and label distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.

Ziming Zhang and Venkatesh Saligrama. 2016. [Zero-shot learning via joint latent similarity embedding](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 6034–6042. IEEE Computer Society.

A Pretrained transformers make poor universal sentence encoders

Certain pretrained transformers, such as BERT and ALBERT, have mechanisms for learning sequence-level embeddings via self-supervision. These models prepend every input sequence with a special classification token (e.g. “[CLS]”), and its representation is learned using a simple classification task, such as Next Sentence Prediction (NSP) or Sentence-Order Prediction (SOP) (see [Devlin et al. 2019](#) and [Lan et al. 2020](#) respectively for details on these tasks). However, during preliminary experiments, we noticed that these models are not good universal sentence encoders, as measured by their performance on the SentEval benchmark ([Conneau and Kiela, 2018](#)). As a simple experiment, we

evaluated three pretrained transformer models on SentEval: one trained with the NSP loss (BERT), one trained with the SOP loss (ALBERT) and one trained with neither, RoBERTa ([Liu et al., 2019](#)). We did not find that the CLS embeddings produced by models trained against the NSP or SOP losses to outperform that of a model trained without either loss and sometimes failed to outperform a bag-of-words (BoW) baseline ([Table 4](#)). Furthermore, we find that pooling token embeddings via averaging (referred to as “mean pooling” in our paper) outperforms pooling via the CLS classification token. Our results are corroborated by [Liu et al. 2019](#), who find that removing NSP loss leads to the same or better results on downstream tasks and [Reimers and Gurevych 2019](#), who find that directly using the output of BERT as sentence embeddings leads to poor performances on the semantic similarity tasks of SentEval.

B Examples of sampled spans

In [Table 5](#), we present examples of anchor-positive and anchor-negative pairs generated by our sampling procedure. We show one example for each possible view of a sampled positive, e.g. positives adjacent to, overlapping with, or subsumed by the anchor. For each anchor-positive pair, we show examples of both a *hard* negative (derived from the same document) and an *easy* negative (derived from another document). Recall that a minibatch is composed of random documents, and each anchor-positive pair sampled from a document is contrasted against all other anchor-positive pairs in the minibatch. Thus, hard negatives, as we have described them here, are generated only when sampling multiple anchors per document ($A > 1$).

C SentEval evaluation details

SentEval is a benchmark for evaluating the quality of fixed-length sentence embeddings. It is divided into 18 *downstream* tasks, and 10 *probing* tasks. Sentence embedding methods are evaluated on these tasks via a simple interface⁹, which standardizes training, evaluation and hyperparameters. For most tasks, the method to evaluate is used to produce fix-length sentence embeddings, and a simple logistic regression (LR) or multi-layer perception (MLP) model is trained on the task using these embeddings as input. For other tasks (namely

⁹<https://github.com/facebookresearch/SentEval>

Table 4: Results on the downstream and probing tasks from the validation set of SentEval. We compare models trained with the Next Sentence Prediction (NSP) and Sentence-Order Prediction (SOP) losses to a model trained with neither, using two different pooling strategies: ”*-CLS”, where the special classification token is used as its sentence representation, and ”*-mean”, where each sentence is represented by the mean of its token embeddings.

Model	Parameters	Embed. dim.	SentEval	
			Downstream	Probing
<i>Bag-of-Words (BoW) weak baselines</i>				
GloVe	–	300	66.05	62.93
fastText	–	300	68.75	63.46
<i>Trained with Next Sentence Prediction (NSP) loss</i>				
BERT-base-CLS	110M	768	63.53	69.57
BERT-base-mean	110M	768	71.98	73.37
<i>Trained with Sentence-Order Prediction (SOP) loss</i>				
ALBERT-base-V2-CLS	11M	768	58.75	69.88
ALBERT-base-V2-mean	11M	768	69.39	74.83
<i>Trained with neither NSP or SOP losses</i>				
RoBERTa-base-CLS	125M	768	68.53	66.92
RoBERTa-base-mean	125M	768	72.84	74.59

several semantic text similarity tasks), the embeddings are used as-is without any further training. Note that this setup is different from evaluations on the popular GLUE benchmark (Wang et al., 2019), which typically use the task data to fine-tune the parameters of the sentence embedding model.

In subsection C.1, we present the individual tasks of the SentEval benchmark. In subsection C.2, we explain our method for computing the average downstream and average probing scores presented in our paper.

C.1 SentEval tasks

The downstream tasks of SentEval are representative NLP tasks used to evaluate the transferability of fixed-length sentence embeddings. We give a brief overview of the broad categories that divide the tasks below (see Conneau and Kiela 2018 for more details):

- **Binary and multi-class classification:** These tasks cover various types of sentence classification, including sentiment analysis (MR Pang and Lee 2005, SST2 and SST5 Socher et al. 2013), question-type (TREC) (Voorhees and Tice, 2000), product reviews (CR) (Hu and Liu, 2004), subjectivity/objectivity (SUBJ) (Pang and Lee, 2004) and opinion polarity (MPQA) (Wiebe et al., 2005).
- **Entailment and semantic relatedness:** These tasks cover multiple entailment

datasets (also known as natural language inference or NLI), including SICK-E (Marelli et al., 2014) and the Stanford NLI dataset (SNLI) (Bowman et al., 2015) as well as multiple semantic relatedness datasets including SICK-R and STS-B (Cer et al., 2017).

- **Semantic textual similarity** These tasks (STS12 Agirre et al. 2012, STS13 Agirre et al. 2013, STS14 Agirre et al. 2014, STS15 Agirre et al. 2015 and STS16 Agirre et al. 2016) are similar to the semantic relatedness tasks, except the embeddings produced by the encoder are used as-is in a cosine similarity to determine the semantic similarity of two sentences. No additional model is trained on top of the encoder’s output.
- **Paraphrase detection** Evaluated on the Microsoft Research Paraphrase Corpus (MRPC) (Dolan et al., 2004), this binary classification task is comprised of human-labelled sentence pairs, annotated according to whether they capture a paraphrase/semantic equivalence relationship.

Table 5: Examples of text spans generated by our method. During training, we randomly sample one or more anchors from every document in a minibatch. For each anchor, we randomly sample one or more positives adjacent to, overlapping with, or subsumed by the anchor. All anchor-positive pairs are contrasted with every other anchor-positive pair in the minibatch. This leads to *easy* negatives (anchors and positives sampled from *other* documents in a minibatch) and *hard* negatives (anchors and positives sampled from the *same* document). Here, examples are capped at a maximum length of 64 tokens. During training, we sample spans up to a length of 512 tokens.

Anchor	Positive	Hard negative	Easy negative
immigrant-rights advocates and law enforcement professionals were skeptical of the new program. Any effort by local cops to enforce immigration laws, they felt, would be bad for community policing, since immigrant victims or witnesses of crime wouldn't feel comfortable talking to police.	feel comfortable talking to police. Some were skeptical that ICE's intentions were really to protect public safety, rather than simply to deport unauthorized immigrants more easily.	liberal parts of the country with large immigrant populations, like Santa Clara County in California and Cook County in Illinois, agreed with the critics of Secure Communities. They worried that implementing the program would strain their relationships with immigrant residents.	that a new location is now available for exploration. A good area, in my view, feels like a natural progression of a game world it doesn't seem tacked on or arbitrary. That in turn needs it to relate
if the ash stops belching out of the volcano then, after a few days, the problem will have cleared, so that's one of the factors. "The other is the wind speed and direction." At the moment the weather patterns are very volatile which is what is making it quite difficult, unlike last year, to predict	where the ash will go. "The public can be absolutely confident that airlines are only able to operate when it is safe to do so." Ryanair said it could not see any ash cloud	A British Airways jumbo jet was grounded in Canada on Sunday following fears the engines had been contaminated with volcanic ash	events are processed in FIFO order. When this nextTickQueue is emptied, the event loop considers all operations to have been completed for the current phase and transitions to the next phase.
Far Cry Primal is an action-adventure video game developed by Ubisoft Montreal and published by Ubisoft. It was released worldwide for PlayStation 4 and Xbox One on February 23, 2016, and for Microsoft Windows on March 1, 2016. The game is a spin-off of the main Far Cry series. It is the first Far Cry game set in the Mesolithic Age.	Ubisoft. It was released worldwide for PlayStation 4 and Xbox One on February 23, 2016, and for Microsoft Windows on March 1, 2016. The game is a spin-off of the main Far Cry series.	Players take on the role of a Wenja tribesman named Takkar, who is stranded in Oros with no weapons after his hunting party is ambushed by a Saber-tooth Tiger.	to such feelings. Fawkes cried out and flew ahead, and Albus Dumbledore followed. Further along the Dementors' path, people were still alive to be fought for. And no matter how much he himself was hurting, while there were still people who needed him he would go on. For

- **Caption-Image retrieval** This task is comprised of two sub-tasks: ranking a large collection of images by their relevance for some given query text (Image Retrieval) and ranking captions by their relevance for some given query image (Caption Retrieval). Both tasks are evaluated on data from the COCO dataset (Lin et al., 2014). Each image is represented by a pretrained, 2048-dimensional embedding produced by a ResNet-101 (He et al., 2016).

The probing tasks are designed to evaluate what linguistic properties are encoded in a sentence representation. All tasks are binary or multi-class classification. We give a brief overview of each task below (see Conneau et al. 2018 for more details):

- **Sentence length (SentLen)**: A multi-class classification task where a model is trained to predict the length of a given input sentence, which is binned into six possible length ranges.
- **Word content (WC)**: A multi-class classification task where, given 1000 words as targets, the goal is to predict which of the target words appears in a given input sentence. Each sentence contains a single target word, and the word occurs exactly once in the sentence.
- **Tree depth (TreeDepth)**: A multi-class classification task where the goal is to predict the maximum depth (with values ranging from 5 to 12) of a given input sentence’s syntactic tree.
- **Bigram Shift (BShift)**: A multi-class classification task where the goal is to predict whether two consecutive tokens within a given sentence have been inverted.
- **Top Constituents (TopConst)**: A multi-class classification task where the goal is to predict the top constituents (from a choice of 19) immediately below the sentence (S) node of the sentence’s syntactic tree.
- **Tense**: A binary classification task where the goal is to predict the tense (past or present) of the main verb in a sentence.
- **Subject number (SubjNum)**: A binary classification task where the goal is to predict the number (singular or plural) of the subject of the main clause.

- **Object number (ObjNum)**: A binary classification task, analogous to SubjNum, where the goal is to predict the number (singular or plural) of the direct object of the main clause.
- **Semantic odd man out (SOMO)**: A binary classification task where the goal is to predict whether a sentence has had a single randomly picked noun or verb replaced with another word with the same part-of-speech.
- **Coordinate inversion (CoordInv)**: A binary classification task where the goal is to predict whether the order of two coordinate clauses in a sentence has been inverted.

C.2 Computing an average score

In our paper, we present averaged downstream and probing scores. Computing averaged probing scores was straightforward; each of the ten probing tasks reports a simple accuracy, which we averaged. To compute an averaged downstream score, we do the following:

- If a task reports Spearman correlation (i.e. SICK-R, STS-B), we use this score when computing the average downstream task score. If the task reports a mean Spearman correlation for multiple subtasks (i.e. STS12, STS13, STS14, STS15, STS16), we use this score.
- If a task reports both an accuracy and an F1-score (i.e. MRPC), we use the average of these two scores.
- For the Caption-Image Retrieval task, we report the average of the Recall@K, where $K \in \{1, 5, 10\}$ for the Image and Caption retrieval tasks (a total of six scores). This is the default behaviour of SentEval.
- Otherwise, we use the reported accuracy.