

Paraphrastic Representations at Scale

John Wieting¹, Kevin Gimpel², Graham Neubig³, and Taylor Berg-Kirkpatrick⁴

¹Google Research

²Toyota Technological Institute at Chicago, Chicago, IL, 60637, USA

³Carnegie Mellon University, Pittsburgh, PA, 15213, USA

⁴University of California San Diego, San Diego, CA, 92093, USA

jwieting@alumni.cmu.edu, kgimpel@ttic.edu, gneubig@cs.cmu.edu, tberg@eng.ucsd.edu

Abstract

We present a system that allows users to train their own state-of-the-art paraphrastic sentence representations in a variety of languages. We release trained models for English, Arabic, German, Spanish, French, Russian, Turkish, and Chinese. We train these models on large amounts of data, achieving significantly improved performance from our original papers on a suite of monolingual semantic similarity, cross-lingual semantic similarity, and bitext mining tasks. Moreover, the resulting models surpass all prior work on efficient unsupervised semantic textual similarity, even significantly outperforming supervised BERT-based models like Sentence-BERT (Reimers and Gurevych, 2019). Most importantly, our models are orders of magnitude faster than other strong similarity models and can be used on CPU with little difference in inference speed (even improved speed over GPU when using more CPU cores), making these models an attractive choice for users without access to GPUs or for use on embedded devices. Finally, we add significantly increased functionality to the code bases for training paraphrastic sentence models, easing their use for both inference and for training them for any desired language with parallel data. We also include code to automatically download and preprocess training data.¹

1 Introduction

Measuring sentence similarity (Agirre et al., 2012) is an important task in natural language processing, and has found many uses including paraphrase detection (Dolan et al., 2004), bitext mining (Schwenk and Douze, 2017), language modelling (Khandelwal et al., 2019), question-answering (Lewis et al., 2021), and as reward functions or evaluation metrics for language generation

¹Code, including an easy to install PyPi package, released models including Hugging Face implementations, demo, and data are available at <https://github.com/jwieting/paraphrastic-representations-at-scale>.

tasks (Wieting et al., 2019a). Within this context, fast and light-weight methods are particularly useful as they make it easy to compute similarity over the ever-increasing volumes of web text available. For instance, we may want to mine a hundred million parallel sentences (Schwenk et al., 2021) or use a semantic similarity reward when fine-tuning language generation models on tens of millions of training examples. These tasks are much more feasible when using approaches that are fast, can be run on CPU, and use little RAM, allowing for increased batch size.

This need for fast inference is one motivation for using sentence embeddings. Sentence embeddings allow the search for similar sentences to be linear in the number of sentences, or even sub-linear when using highly optimized tools like Faiss (Johnson et al., 2017) that allow for efficient nearest neighbor search. This is contrast to models, like cross-attention models, which are quadratic during inference as they require both of the texts being compared as inputs. As we show in this paper, our *simple* and *interpretable* word-averaging sentence embedding models (Wieting et al., 2016b; Wieting and Gimpel, 2018; Wieting et al., 2019b), are orders of magnitude faster to compute than prior embedding approaches while simultaneously possessing significantly stronger performance on monolingual and cross-lingual semantic similarity tasks. Since we are simply averaging embeddings and have no neural architecture, any models based on neural architectures, especially large pretrained neural architectures which are increasingly used, will not be as fast as the models described in this paper. Lastly, we also show that this approach is competitive with LASER (Artetxe and Schwenk, 2019), a state-of-the-art multilingual model, on mining bitext and has stronger performance on cross-lingual semantic similarity, while having inference speeds that are twice as fast on GPU and orders of magnitude faster on CPU.

We make several contributions in this paper that go beyond our prior work. Firstly, we reformat the code to support training models on tens of millions of sentence pairs efficiently and with low RAM usage. Secondly, we train an English model on 25.85 million paraphrase pairs from ParaNMT (Wieting and Gimpel, 2018), a paraphrase corpus we previously constructed automatically from bitext. We then train models directly on X-English bitext for Arabic, German, Spanish, French, Russian, Turkish, and Chinese, producing models that are able to distinguish both paraphrases in English and their respective languages as well as cross-lingual X-English paraphrases. Even though all models are able to model semantic similarity in English, we find that training on ParaNMT specifically leads to stronger models as it is easier to filter the data to remove noise and sentence pairs with little to no diversity. We refer to our models as PARAGRAM-SP, abbreviated as P-SP,² referring to how the models are based on averaging subword units generated by sentencepiece (Kudo and Richardson, 2018). We make all of these models available to the community for use on downstream tasks.

We also add functionality to our implementation. Besides the support for efficient, low-memory training on tens of million of sentence pairs described above, we add code to support (1) reading in a list of sentences and producing a saved numpy array of the sentence embeddings; (2) reading in a list of sentence pairs and producing cosine similarity scores; and (3) downloading and preprocessing evaluation data, bitext, and paraphrase data. For bitext and paraphrase data, we provide support for training using either text files or HDF5 files.

Lastly, this paper contains new experiments showcasing the limits of these scaled-up models and detailed comparisons with prior work on a suite of semantic similarity tasks in a variety of languages. We release our code and models to the community in the hope that they will be found useful for research and applications, as well as using them as a base to build stronger, faster models covering more of the languages of the world.

²Our English model is P-SP, and the cross-lingual models are P-SP-AR, P-SP-DE, P-SP-ES, P-SP-FR, P-SP-RU, P-SP-TR, and P-SP-ZH.

2 Related Work

2.1 English Semantic Similarity

Our learning and evaluation setting is the same as that of our earlier work that seeks to learn paraphrastic sentence embeddings that can be used for downstream tasks (Wieting et al., 2016b,a; Wieting and Gimpel, 2017; Wieting et al., 2017; Wieting and Gimpel, 2018). We trained models on noisy paraphrase pairs and evaluated them primarily on semantic textual similarity (STS) tasks. More recently, we made use of parallel bitext for training paraphrastic representations for other languages as well that are also able to model cross-lingual semantic similarity (Wieting et al., 2019a, 2020). Prior work in learning general sentence embeddings has used autoencoders (Socher et al., 2011; Hill et al., 2016), encoder-decoder architectures (Kiros et al., 2015; Gan et al., 2017), and other sources of supervision and learning frameworks (Le and Mikolov, 2014; Pham et al., 2015; Arora et al., 2017; Pagliardini et al., 2017).

For English semantic similarity, we compare to well known sentence embedding models such as InferSent (Conneau et al., 2017), GenSen (Subramanian et al., 2018), the Universal Sentence Encoder (USE) (Cer et al., 2018), as well as BERT (Devlin et al., 2019).³ We use the pretrained BERT model in two ways to create a sentence embedding. The first way is to concatenate the hidden states for the CLS token in the last four layers. The second way is to concatenate the hidden states of all word tokens in the last four layers and mean pool these representations. Both methods result in a 4096 dimension embedding. We also compare to a more recently released model called Sentence-BERT (Reimers and Gurevych, 2019). This model is similar to InferSent in that it is trained on natural language inference data (SNLI; Bowman et al., 2015). However, instead of using pretrained word embeddings, they fine-tune BERT in a way to induce sentence embeddings. Lastly, we also compare to the unsupervised version of SimCSE (Gao et al., 2021), which fine-tunes a pretrained encoder on contrastive pairs, where positive pairs are obtained by using dropout on a single input sentence.

³Note that in all experiments using BERT, including Sentence-BERT, the large, uncased version is used.

2.2 Cross-Lingual Semantic Similarity and Semantic Similarity in Non-English Languages

Most previous work for cross-lingual representations has focused on models based on encoders from neural machine translation (Espana-Bonet et al., 2017; Schwenk and Douze, 2017; Schwenk, 2018) or deep architectures using contrastive losses (Grégoire and Langlais, 2018; Guo et al., 2018; Chidambaram et al., 2019). Recently, other approaches using large Transformer (Vaswani et al., 2017) have been proposed, trained on vast quantities of text (Conneau et al., 2020; Liu et al., 2020; Tran et al., 2020). We primarily focus our comparison for these settings on LASER (Artetxe and Schwenk, 2019), a model trained for semantic similarity across more than 100 languages. Their model uses an LSTM encoder-decoder trained on hundreds of millions of parallel sentences. They achieve state-of-the-art performance on a variety of multilingual sentence embeddings tasks including bitext mining. We also compare to LaBSE (Feng et al., 2022), a contrastive model trained on six billion parallel pairs across languages and was also trained on monolingual text using a masked language modelling objective.

3 Methods

We first describe our objective function and then describe our encoder.

Training. The training data consists of a sequence of parallel sentence pairs (s_i, t_i) in source and target languages respectively. Note that for training our English model, the source and target languages are both English as we are able to make use of an existing paraphrase corpus. For each sentence pair, we randomly choose a *negative* target sentence t'_i during training that is not a translation or paraphrase of s_i . Our objective is to have source and target sentences be more similar than source and negative target examples by a margin δ :

$$\min_{\theta_{\text{src}}, \theta_{\text{tgt}}} \sum_i \left[\delta - f_{\theta}(s_i, t_i) + f_{\theta}(s_i, t'_i) \right]_+ \quad (1)$$

where the similarity function is defined as:

$$f_{\theta}(s, t) = \cos\left(g(s; \theta_{\text{src}}), g(t; \theta_{\text{tgt}})\right) \quad (2)$$

where g is the sentence encoder with parameters for each language $\theta = (\theta_{\text{src}}, \theta_{\text{tgt}})$. To select t'_i

we choose the most similar sentence in some set according to the current model parameters, i.e., the one with the highest cosine similarity. We found we could achieve the strongest performance by tying all parameters together for each language, more precisely, θ_{src} and θ_{tgt} are the same.

Negative Sampling. Negative examples are selected from the sentences in the batch from the opposing language when training with bitext and from any sentence in the batch when using paraphrase data. In all cases, we choose the negative example with the highest cosine similarity to the given sentence s , ensuring that the negative is not in fact paired with s in the batch. To select even more difficult negative examples that aid training, we use the *mega-batching* procedure of Wieting and Gimpel (2018), which aggregates M mini-batches to create one “mega-batch” and selects negative examples from this mega-batch. Once each pair in the mega-batch has a negative example, the mega-batch is split back up into M mini-batches for training. Additionally, we anneal the mega-batch size by slowly increasing it during training. This yields improved performance by a significant margin.

Encoder. Our sentence encoder g simply averages the embeddings of subword units generated by sentencepiece (Kudo and Richardson, 2018); we refer to our model as PARAGRAM-SP, abbreviated as P-SP. This means that the sentence piece embeddings themselves are the only learned parameters of this model.

4 Code and Usage

We added a number of features to the code base to improve performance and make it easier to use. First, we added code to support easier inference. Examples of using the code programmatically to embed sentences and score sentence pairs (using cosine similarity) are shown in Figure 1.

Our code base also supports functionality that allows one to read in a list of sentences and produce a saved numpy array of the sentence embeddings. We also included functionality that allows one to read in a list of sentence pairs and produce the sentence pairs along with their cosine similarity scores in an output file. These scripts allow our models to be used without any programming for the two most common use cases: embedding sentences and scoring sentence pairs. Examples of their usage with a trained model are shown in Figure 2.

```

1 from models import load_model
2
3 text1 = 'This is a test.'
4 text2 = 'This is another test.'
5
6 # Load English paraphrase model
7 model_name = 'paraphrase-at-scale/model.
   para.lc.100.pt'
8 sp_model = 'paraphrase-at-scale/paranmt.
   model'
9
10 model, _ = load_model(model_name=
   model_name, sp_model=sp_model)
11
12 # Obtain sentence embedding
13 embeddings = model.embed_raw_text([text1
   , text2]) # 2D numpy array of
   sentence embeddings
14 cosine_scores = model.score_raw_text([[
   text1, text2]]) # list of cosine
   scores

```

Figure 1: Usage example of programmatically loading one of our pretrained models and obtaining sentence embeddings and scores for two sentences.

```

1 python -u embed_sentences.py --sentence-
   file paraphrase-at-scale/example-
   sentences.txt --load-file paraphrase
   -at-scale/model.para.lc.100.pt --
   output-file sentence_embeds.np
2
3 python score_sentence_pairs.py --
   sentence-pair-file paraphrase-at-
   scale/example-sentences-pairs.txt --
   load-file paraphrase-at-scale/model.
   para.lc.100.pt

```

Figure 2: Usage examples to embed sentences and score sentence pairs. The first command is a usage example of scoring a list of sentence pairs. The file `example-sentences-pairs.txt` contains a list of sentences, one per line. The output of the script is a saved numpy array of sentence embeddings in the same order of the input sentences. The second command is a usage example of scoring a list of sentence pairs. The file `example-sentences-pairs.txt` contains pairs of tab-separated sentences, one per line. The output of the script is a text file containing the tab separated list of sentences along with their cosine scores in the same order of the input sentences.

Secondly, we added a training mode using HDF5⁴ format, allowing training data to remain on disk during training. This leads to a significant reduction in RAM usage during training, which is especially true when using more than 10 million training examples. Efficient training can now be done on CPU only using only a few gigabytes of RAM.

Lastly, we also added code for preprocessing

⁴<https://docs.h5py.org/en/stable/>

```

1 cd preprocess/bilingual && bash do_all.
   sh fr-es-de
2 cd ../../
3 cd preprocess/paranmt && bash do_all.sh
   0.4 1.0 0.7

```

Figure 3: Usage examples to download and preprocess bilingual and ParaNMT data. The first command downloads and preprocesses (filters, trains sentencepiece models, tokenizes if language is zh, converts files to hdf5 format) en-X bilingual data. The third command downloads and preprocesses ParaNMT data. The arguments are used to filter the data (semantic similarity scores between 0.4 and 1.0 and trigram overlap below 0.7, which have been used in prior papers when generating training data for paraphrase generation (Iyyer et al., 2018; Krishna et al., 2020)).

en	ar	de	es	fr	ru	tr	zh
25.85M	8.23M	6.47M	6.75M	6.46M	9.09M	5.12M	4.18M

Table 1: The number of sentence pairs used to train our models. For English, the data is ParaNMT, and for the other languages, the data is a collection of bitext detailed in Section 5.1.

data, including scripts to download and evaluate on the STS data (English, non-English, and cross-lingual), as well as code to download and process bitext and ParaNMT automatically. For bitext, our scripts download the data, filter the data by length,⁵ lowercase, remove duplicates, train a sentencepiece model, encode the data with the sentencepiece model, shuffle the data, and process the data into HDF5 format for efficient use. For ParaNMT, our scripts download the data, use a language classifier to filter out non-English sentences⁶ (Joulin et al., 2017), filter the data by paraphrase score, trigram overlap, and length,⁷ train a sentencepiece model, encode the data with the sentencepiece model, and process the data into HDF5 format. Examples are shown in Figure 3.

5 Experiments

5.1 Experimental Setup

Data. For our English model, we train on selected sentence pairs from ParaNMT (Wieting and Gimpel, 2018). We filter the corpus by only including sentence pairs where the paraphrase score for the two sentences is ≥ 0.4 . We additionally filtered

⁵We remove sentences with the number of tokens (untokenized) smaller than 3 or greater than 100.

⁶<https://fasttext.cc>

⁷We remove sentences with the number of tokens (untokenized) smaller than 5 or greater than 40.

Model	Semantic Textual Similarity (STS)					Avg.
	2012	2013	2014	2015	2016	
BERT (CLS)	33.2	29.6	34.3	45.1	48.4	38.1
BERT (Mean)	48.8	46.5	54.0	59.2	63.4	54.4
InferSent	61.1	51.4	68.1	70.9	70.7	64.4
GenSen	60.7	50.8	64.1	73.3	66.0	63.0
USE	61.4	59.0	70.6	74.3	73.9	67.8
Sentence-BERT	66.9	63.2	74.2	77.3	72.8	70.9
LASER	63.1	47.0	67.7	74.9	71.9	64.9
P-SP	68.7	64.7	78.1	81.4	80.0	74.6
Sentence-BERT	71.0	76.5	73.2	79.1	74.3	74.8
P-SP	71.2	76.5	74.6	83.0	79.1	76.9

Table 2: Results of our models and models from prior work on English STS. In the first part of the table, we show results, measured in Pearson’s $r \times 100$, for each year of the STS tasks 2012-2016 as well as the average performance across all years. In the second part, we evaluate based on the Spearman’s $\rho \times 100$ of the concatenation of the datasets of each year with the 2013 SMT dataset removed following (Reimers and Gurevych, 2019).

Model	Dim.	ar-ar	ar-en	es-es	es-en	tr-en
LASER	1024	69.3\68.8	65.5\66.5	79.7\79.7	59.7\58.0	72.0\72.1
LaBSE	768	68.6\69.1	72.2\74.5	79.5\80.8	65.5\65.7	72.9\72.1
Espana-Bonet et al. (2017)	2048	59	44	78	49	76
Chidambaram et al. (2019)	512	-	-	64.2	58.7	-
2017 STS 1st Place	-	75.4	74.9	85.6	83.0	77.1
2017 STS 2nd Place	-	75.4	71.3	85.0	81.3	74.2
2017 STS 3rd Place	-	74.6	70.0	84.9	79.1	73.6
P-SP	1024	76.2\76.7	78.3\78.4	85.8\85.6	78.4\77.8	79.2\79.5

Table 3: Comparison of our models with those in the literature on non-English and cross-lingual STS. We also include the top 3 systems for each dataset from the SemEval 2017 STS shared task. Performance is measured in Pearson’s $r \times 100$. We also include results in Spearman’s $\rho \times 100$ after a slash for LASER, LaBSE, and P-SP.

sentence pairs by their trigram overlap (Wieting et al., 2017), which is calculated by counting trigrams in the two sentences, and then dividing the number of shared trigrams by the total number in the sentence with fewer tokens. We only include sentence pairs where the trigram overlap score is ≤ 0.7 . The paraphrase score is calculated by averaging PARAGRAM-PHRASE embeddings (Wieting et al., 2016b) for the two sentences in each pair and then computing their cosine similarity. The purpose of the lower threshold is to remove noise while the higher threshold is meant to remove paraphrases that are too similar.

Our training data is a mixture of Open Subtitles 2018⁸ (Lison and Tiedemann, 2016), Tanzil corpus⁹ (Tiedemann, 2012), Europarl¹⁰ for Spanish, Global Voices¹¹ (Tiedemann, 2012), and the MultiUN corpus.¹² We follow the same distribution for our languages of interest across data sources as Artetxe and Schwenk (2019) for a fair comparison.

⁸<http://opus.nlpl.eu/OpenSubtitles.php>

⁹<http://opus.nlpl.eu/Tanzil.php>

¹⁰<http://opus.nlpl.eu/Europarl.php>

¹¹<https://opus.nlpl.eu/GlobalVoices.php>

¹²<http://opus.nlpl.eu/MultiUN.php>

One exception, though, is we do not include training data from Tatoeba¹³ (Tiedemann, 2012) as they do, since this domain is also in the bitext mining evaluation set. The amount of data used to train each of our models is shown in Table 1.

Hyperparameters. For all models, we fix the batch size to 128, margin δ to 0.4, and the annealing rate to 150.¹⁴ We set the size of the sentencepiece vocabulary to 50,000, using a shared vocabulary for the models trained on bitext. If a word is not in vocabulary, we simply exclude it, unless the text only consists of unknown words in which case we use a single unknown-word token. We optimize our models using Adam (Kingma and Ba, 2014) with a learning rate of 0.001 and train models for 25 epochs.

For training on the bilingual corpora, we tune each model on the 250 example 2017 English STS task (Cer et al., 2017). We vary dropout on the embeddings over $\{0, 0.1, 0.3\}$ and the mega-batch size M over $\{60, 100, 140\}$.

¹³<https://opus.nlpl.eu/Tatoeba.php>

¹⁴Annealing rate is the number of minibatches that are processed before the megabatch size is increased by 1.

For training on ParaNMT, we fix the hyperparameters in our model due to the increased data size making tuning more expensive. We use a megabatch size M of 100 and set the dropout on the embeddings to 0.0.

5.2 Evaluation

We evaluate sentence embeddings using the SemEval semantic textual similarity (STS) tasks from 2012 to 2016 (Agirre et al., 2012, 2013, 2014, 2015, 2016) as was done initially for sentence embeddings in (Wieting et al., 2016b). Given two sentences, the aim of the STS tasks is to predict their similarity on a 0-5 scale, where 0 indicates the sentences are on different topics and 5 means they are completely equivalent. As our test set, we report the average Pearson’s r over each year of the STS tasks from 2012-2016 as is convention.

Most work evaluating accuracy on STS tasks has averaged the Pearson’s r over each individual dataset for each year of the STS competition. However, Reimers and Gurevych (2019) computed Spearman’s ρ over concatenated datasets for each year of the STS competition. To be consistent with previous work, we re-ran their model and calculated results using the standard method, and thus our results are not the same as those reported Reimers and Gurevych (2019). However, we also include results using their approach for completeness. One other difference between these two ways of calculating the results is the inclusion of the SMT dataset of the 2013 task, which we also exclude when replicating the approach in Reimers and Gurevych (2019).

For cross-lingual semantic similarity and semantic similarity in non-English languages, we evaluate on the STS tasks from SemEval 2017. This evaluation contains Arabic-Arabic, Arabic-English, Spanish-Spanish, Spanish-English, and Turkish-English datasets. The datasets were created by translating one or both pairs of an English STS pair into Arabic (ar), Spanish (es), or Turkish (tr). Following convention, we report results with Pearson’s r for all systems, but also include results in Spearman’s ρ for LASER, LaBSE, and P-SP.

We also evaluate on the Tatoeba bitext mining task introduced by Artetxe and Schwenk (2019). The dataset consists of up to 1,000 English-aligned sentence pairs for over 100 languages. The aim of the task is to find the nearest neighbor for each sentence in the other language according to cosine

Language	LASER	XLM-R	mBART	CRISS	LaBSE	P-SP
ar	7.8	52.5	61.0	22.0	9.1	8.8
de	1.0	11.1	13.2	2.0	0.7	1.5
es	2.1	24.3	39.6	3.7	1.6	2.4
fr	4.3	26.3	39.6	7.3	4.0	5.4
ru	5.9	25.9	31.6	9.7	4.7	5.6
tr	2.6	34.3	48.8	7.1	1.6	1.4
Avg.	4.0	29.1	39.0	8.6	3.6	4.2

Table 4: Results on the Tatoeba bitext mining task (Artetxe and Schwenk, 2019). Results are measured in error rate $\times 100$.

similarity. Performance is measured by computing the error rate.

6 Results

English Semantic Similarity. The results for our English semantic similarity evaluation are shown in Table 2. Our P-SP model has the best performance across each year of the task, significantly outperforming all prior work. We outperform methods that use large pre-trained models including Sentence-BERT which is supervised, as it is trained on NLI data (Bowman et al., 2015).

We also include results from SimCSE (Gao et al., 2021). We compare to the unsupervised version, since our model is also unsupervised. We evaluate using the Spearman’s ρ of the concatenation of the datasets for each year, and find our average performance over the 2012-2016 datasets to be 76.9, compared to 77.4 and 77.9 for the RoBERTa-base (Liu et al., 2019) and RoBERTa-large versions of SimCSE. While our performance is slightly lower, we note that they tune their model on the dev set of the STS Benchmark (Cer et al., 2017), which contains a subset of the data from STS tasks which we use for evaluation. Therefore, they are tuning on a subset of the evaluation data, and it is unclear how tuning on this test data affects model performance.

Cross-Lingual Semantic Similarity. The results for the non-English and cross-lingual semantic similarity evaluation are shown in Table 3. From the results, our model again outperforms all prior work using sentence embeddings. The only systems that have better performance are the top (non-embedding based) systems from SemEval 2017 for Spanish-English.¹⁵

¹⁵The top systems for this task used supervision and relied on state-of-the-art translation models to first translate the non-English sentences to English.

Bitext Mining. The results on the Tatoeba bitext mining task from Artetxe and Schwenk (2019) are shown in Table 4. The results show that our embeddings are competitive, but have slightly higher error rates than LASER. The models are so close that the difference in error rate for the two models across the 6 evaluations is 0.2, corresponding to a difference of about 2 mismatched sentence pairs per dataset. LaBSE performs a bit better, but was trained on much more data than both LASER and our method. We also compare to mBART, XLM-R, and CRISS.¹⁶

This bitext mining result is in contrast to the results on cross-lingual semantic similarity, suggesting that our embeddings account for a less literal semantic similarity, making them more adept at detecting paraphrases but slightly weaker at identifying translations. It is also worth noting that LASER was trained on Tatoeba data outside the test sets, which could also account for some of the slight improvement over our model.

7 Speed Analysis

Model	GPU	CPU
P-SP	13,863	12,776
LASER	6,033	26
Sentence-Bert	288	2
InferSent	4,445	16

Table 5: Speed as measured in sentences/second on both GPU (Nvidia 1080 TI) and CPU (single core).

We analyze the speed of our models as well as selected popular sentence embedding models from prior work. To evaluate inference speed, we measure the time required to embed 120,000 sentences from the Toronto Book Corpus (Zhu et al., 2015). Preprocessing of sentences is not factored into the timing, and each method sorts the sentences by length prior to computing the embeddings to reduce padding and extra computation. We use a batch size of 64 for each model. The number of sentences embedded per second is shown in Table 5.

From the results, we see that our model is easily the fastest on GPU, sometimes by an order of magnitude. Interestingly, using a single core of CPU, we achieve similar speeds to inference on GPU, which is not the case for any other model. Moreover, we repeated the experiment, this time

¹⁶Results are copied from (Tran et al., 2020).

using 32 cores and achieved a speed of 15,316 sentences/second. This is even faster than when using a GPU and indicates that our model can effectively be used at scale when GPUs are not available. It also suggests our model would be appropriate for use on embedded devices.

8 Conclusion

In this paper, we present a system for the learning and inference of paraphrastic sentence embeddings in any language for which there is paraphrase or bilingual parallel data. Additionally, we release our trained sentence embedding models in English, as well as Arabic, German, Spanish, French, Russian, Turkish, and Chinese. These models are trained on tens of million of sentence pairs resulting in models that achieve state-of-the-art performance on unsupervised English semantic similarity and are state-of-the-art or competitive on non-English semantic similarity, cross-lingual semantic similarity, and bitext mining.

Moreover, our models are significantly faster than prior work owing to their simple architecture. They can also be run on CPU with little to no loss in speed from running them on GPU—something that no strong models from prior work are able to do. Lastly, we release our code that has been modified to make training and inference easier, with support for training on large corpora, preprocessing paraphrase and bilingual corpora and evaluation data, as well as scripts for easy inference that can generate embeddings or semantic similarity scores for sentences supplied in a text file.

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, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. A simple but tough-to-beat baseline for sentence embeddings. In *Proceedings of the International Conference on Learning Representations*.
- Mikel Artetxe and Holger Schwenk. 2019. [Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond](#). *Transactions of the Association for Computational Linguistics*, 7:597–610.
- 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.
- 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.
- 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 (ReplANLP-2019)*, pages 250–259, Florence, Italy. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- 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.
- 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.
- Cristina Espana-Bonet, Adám Csaba Varga, Alberto Barrón-Cedeño, and Josef van Genabith. 2017. An empirical analysis of nmt-derived interlingual embeddings and their use in parallel sentence identification. *IEEE Journal of Selected Topics in Signal Processing*, 11(8):1340–1350.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2022. [Language-agnostic BERT sentence embedding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.
- Zhe Gan, Yunchen Pu, Ricardo Henao, Chunyuan Li, Xiaodong He, and Lawrence Carin. 2017. [Learning generic sentence representations using convolutional neural networks](#). In *Proceedings of the 2017*

- Conference on Empirical Methods in Natural Language Processing*, pages 2390–2400, Copenhagen, Denmark. Association for Computational Linguistics.
- 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*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Francis Grégoire and Philippe Langlais. 2018. **Extracting parallel sentences with bidirectional recurrent neural networks to improve machine translation**. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1442–1453, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Mandy Guo, Qinlan Shen, Yinfei Yang, Heming Ge, Daniel Cer, Gustavo Hernandez Abrego, Keith Stevens, Noah Constant, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. **Effective parallel corpus mining using bilingual sentence embeddings**. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 165–176, Brussels, Belgium. Association for Computational Linguistics.
- 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.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. **Adversarial example generation with syntactically controlled paraphrase networks**. 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 1875–1885, New Orleans, Louisiana. Association for Computational Linguistics.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2017. **Billion-scale similarity search with GPUs**. *arXiv preprint arXiv:1702.08734*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. **Bag of tricks for efficient text classification**. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2019. **Generalization through memorization: Nearest neighbor language models**. *arXiv preprint arXiv:1911.00172*.
- Diederik Kingma and Jimmy Ba. 2014. **Adam: A method for stochastic optimization**. *arXiv preprint arXiv:1412.6980*.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. **Skip-thought vectors**. In *Advances in Neural Information Processing Systems* 28, pages 3294–3302.
- Kalpesh Krishna, John Wieting, and Mohit Iyyer. 2020. **Reformulating unsupervised style transfer as paraphrase generation**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 737–762, Online. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. **SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Quoc V. Le and Tomas Mikolov. 2014. **Distributed representations of sentences and documents**. *arXiv preprint arXiv:1405.4053*.
- Patrick Lewis, Yuxiang Wu, Linqing Liu, Pasquale Minervini, Heinrich Küttler, Aleksandra Piktus, Pontus Stenetorp, and Sebastian Riedel. 2021. **PAQ: 65 million probably-asked questions and what you can do with them**. *Transactions of the Association for Computational Linguistics*, 9:1098–1115.
- Pierre Lison and Jörg Tiedemann. 2016. **OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles**. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. **Multilingual denoising pre-training for neural machine translation**. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- 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*.
- Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. 2017. **Unsupervised learning of sentence embeddings using compositional n-gram features**. *arXiv preprint arXiv:1703.02507*.
- Nghia The Pham, Germán Kruszewski, Angeliki Lazaridou, and Marco Baroni. 2015. **Jointly optimizing word representations for lexical and sentential tasks with the C-PHRASE model**. In *Proceedings of the 53rd Annual Meeting of the Association*

- for *Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 971–981, Beijing, China. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence embeddings using Siamese BERT-networks**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Holger Schwenk. 2018. **Filtering and mining parallel data in a joint multilingual space**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 228–234, Melbourne, Australia. Association for Computational Linguistics.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2021. **Wiki-Matrix: Mining 135M parallel sentences in 1620 language pairs from Wikipedia**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1351–1361, Online. Association for Computational Linguistics.
- Holger Schwenk and Matthijs Douze. 2017. **Learning joint multilingual sentence representations with neural machine translation**. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 157–167, Vancouver, Canada. Association for Computational Linguistics.
- Richard Socher, Eric H. Huang, Jeffrey Pennington, Andrew Y. Ng, and Christopher D. Manning. 2011. **Dynamic pooling and unfolding recursive autoencoders for paraphrase detection**. In *Advances in Neural Information Processing Systems*.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J Pal. 2018. **Learning general purpose distributed sentence representations via large scale multi-task learning**. *arXiv preprint arXiv:1804.00079*.
- Jörg Tiedemann. 2012. **Parallel data, tools and interfaces in OPUS**. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- Chau Tran, Yuqing Tang, Xian Li, and Jiatao Gu. 2020. **Cross-lingual retrieval for iterative self-supervised training**. *arXiv preprint arXiv:2006.09526*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. **Attention is all you need**. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2016a. **Charagram: Embedding words and sentences via character n-grams**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1504–1515, Austin, Texas. Association for Computational Linguistics.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2016b. **Towards universal paraphrastic sentence embeddings**. In *Proceedings of the International Conference on Learning Representations*.
- John Wieting, Taylor Berg-Kirkpatrick, Kevin Gimpel, and Graham Neubig. 2019a. **Beyond BLEU: training neural machine translation with semantic similarity**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4344–4355, Florence, Italy. Association for Computational Linguistics.
- John Wieting and Kevin Gimpel. 2017. **Revisiting recurrent networks for paraphrastic sentence embeddings**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2078–2088, Vancouver, Canada. Association for Computational Linguistics.
- John Wieting and Kevin Gimpel. 2018. **ParaNMT-50M: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Melbourne, Australia. Association for Computational Linguistics.
- John Wieting, Kevin Gimpel, Graham Neubig, and Taylor Berg-Kirkpatrick. 2019b. **Simple and effective paraphrastic similarity from parallel translations**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4602–4608, Florence, Italy. Association for Computational Linguistics.
- John Wieting, Jonathan Mallinson, and Kevin Gimpel. 2017. **Learning paraphrastic sentence embeddings from back-translated bitext**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 274–285, Copenhagen, Denmark. Association for Computational Linguistics.
- John Wieting, Graham Neubig, and Taylor Berg-Kirkpatrick. 2020. **A bilingual generative transformer for semantic sentence embedding**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1581–1594, Online. Association for Computational Linguistics.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. **Aligning books and movies: Towards story-like visual explanations by watching movies and reading books**. In *The IEEE International Conference on Computer Vision (ICCV)*.