

Paraphrase Detection on Noisy Subtitles in Six Languages

Eetu Sjöblom

Mathias Creutz

Mikko Aulamo

Department of Digital Humanities, Faculty of Arts
University of Helsinki

Unioninkatu 40, FI-00014 University of Helsinki, Finland

{eetu.sjoblom, mathias.creutz, mikko.aulamo}@helsinki.fi

Abstract

We perform automatic paraphrase detection on subtitle data from the Opusparcus corpus comprising six European languages: German, English, Finnish, French, Russian, and Swedish. We train two types of supervised sentence embedding models: a word-averaging (WA) model and a gated recurrent averaging network (GRAN) model. We find out that GRAN outperforms WA and is more robust to noisy training data. Better results are obtained with more and noisier data than less and cleaner data. Additionally, we experiment on other datasets, without reaching the same level of performance, because of domain mismatch between training and test data.

1 Introduction

This paper studies automatic paraphrase detection on subtitle data for six European languages. Paraphrases are a set of phrases or full sentences in the same language that mean approximately the same thing. Automatically finding out when two phrases mean the same thing is interesting from both a theoretical and practical perspective. Theoretically, within the field of distributional, compositional semantics, there is currently a significant amount of interest in models and representations that capture the meaning of not just single words, but sequences of words. There are also practical implementations, such as providing multiple alternative correct translations when evaluating the accuracy of machine translation systems.

To our knowledge, the present work is the first published study of automatic paraphrase detection based on data from *Opusparcus*, a recently published paraphrase corpus (Creutz, 2018)¹. *Opusparcus* consists of *sentential* paraphrases, that is, pairs of full sentences that convey approximately

¹Opusparcus is available for download at: <http://urn.fi/urn:nbn:fi:lb-201804191>

the same meaning. *Opusparcus* provides data for six European languages: German, English, Finnish, French, Russian, and Swedish. The data sets have been extracted from OpenSubtitles2016 (Lison and Tiedemann, 2016), which is a collection of translated movie and TV subtitles.²

In addition to *Opusparcus*, experiments are performed on other well known paraphrase resources: (1) *PPDB*, the Paraphrase Database (Ganitkevitch et al., 2013; Ganitkevitch and Callison-Burch, 2014; Pavlick et al., 2015), (2) *MSRPC*, the Microsoft Research Paraphrase Corpus (Quirk et al., 2004; Dolan et al., 2004; Dolan and Brockett, 2005), (3) *SICK* (Marelli et al., 2014), and (4) *STS14* (Agirre et al., 2014).

We are interested in movie and TV subtitles because of their conversational nature. This makes subtitle data ideal for exploring dialogue phenomena and properties of everyday, colloquial language (Paetzold and Specia, 2016; van der Wees et al., 2016; Lison et al., 2018). We would also like to stress the importance of working on other languages beside English. Unfortunately, many language resources contain English data only, such as *MSRPC* and *SICK*. In other datasets, the quality of the English data surpasses that of the other languages to a considerable extent, as in the multilingual version of *PPDB* (Ganitkevitch and Callison-Burch, 2014).

Although our subtitle data is very interesting data, it is also noisy data, in several respects. Since the subtitles are *user-contributed data*, there are misspellings both due to human mistake and due to errors in optical character recognition (OCR). OCR errors emerge when textual subtitle files are

²OpenSubtitles2016 is extracted from www.opensubtitles.org. OpenSubtitles2016 is in itself a subset of the larger OPUS collection (“... the open parallel corpus”): opus.lingfil.uu.se, and provides a large number of sentence-aligned parallel corpora in 65 languages.

produced by “ripping” (scanning) the subtitle text from DVDs using various tools. Furthermore, movies are sometimes not tagged with the correct language, they are encoded in various character encodings, and they come in various formats. (Tiedemann, 2007, 2008, 2016)

A different type of errors emerge because of misalignments and issues with sentence segmentation. Opusparcus has been constructed by finding pairs of sentences in one language that have a common translation in at least one other language. For example, English “*Have a seat.*” is potentially a paraphrase of “*Sit down.*” because both can be translated to French “*Asseyez-vous.*” (Creutz, 2018) To figure out that “*Have a seat.*” is a translation of “*Asseyez-vous.*”, English and French subtitles for the same movie can be used. English and French text that occur at the same time in the movie are assumed to be translations of each other. However, there are many complications involved: Subtitles are not necessarily shown as entire sentences, but as snippets of text that fit on the screen. There are numerous partial overlaps when comparing the contents of subtitle screens across different languages, and the reconstruction of proper sentences may be difficult. There may also be timing differences, because of different subtitle speeds and different time offsets for starting the subtitles. (Tiedemann, 2007, 2008) Furthermore, Lison et al. (2018) argue that “[subtitles] should better be viewed as boiled down transcriptions of the same conversations across several languages. Subtitles will inevitably differ in how they ‘compress’ the conversations, notably due to structural divergences between languages, cultural differences and disparities in subtitling traditions/conventions. As a consequence, sentence alignments extracted from subtitles often have a higher degree of insertions and deletions compared to alignments derived from other sources.”

We tackle the paraphrase detection task using a sentence embedding approach. We experiment with sentence encoding models that take as input a single sentence and produce a vector representing the semantics of the sentence. While models that rely on sentence pairs as input are able to use additional information, such as attention between the sentences, the sentence embedding approach has its advantages: Embeddings can be calculated also when no sentence pair is available, and large numbers of embeddings can be precalculated, which

allows for fast comparisons in huge datasets.

Sentence representation learning has been a topic of growing interest recently. Much of this work has been done in the context of general-purpose sentence embeddings using unsupervised approaches inspired by work on word embeddings (Hill et al., 2016; Kiros et al., 2015) as well as approaches relying on supervised training objectives (Conneau et al., 2017a; Subramanian et al., 2018). While the paraphrase detection task is potentially useful for learning general purpose embeddings, we are mainly interested in paraphrastic sentence embeddings for paraphrase detection and semantic similarity tasks.

Closest to the present work is that of Wieting and Gimpel (2017), who study sentence representation learning using multiple encoding architectures and two different sources of training data. It was found that certain models benefit significantly from using full sentences (SimpWiki) instead of short phrases (PPDB) as training data. However, the SimpWiki data set is relatively small, and this leaves open the question how much the approaches could benefit from very large corpora of sentential paraphrases. It is also unclear how well the approaches generalize to languages other than English.

The current paper takes a step forward in that experiments are performed on five other languages in addition to English. We also study the effects of noise in the training data sets.

2 Data

Opusparcus (Creutz, 2018) contains so-called training, development and test sets for each of the six languages it covers. The training sets, which consist of millions of sentence pairs, have been created automatically and are orders of magnitude larger than the development and test sets, which have been annotated manually and consist of a few thousands of sentence pairs. The development and test sets have different purposes, but otherwise they have identical properties: the development sets can be used for optimization and extensive experimentation, whereas the test sets should only be used in final evaluations.

The development and test sets are “clean” (in principle), since they have been checked by human annotators. The annotators were shown pairs of sentences, and they needed to decide whether the two sentences were paraphrases (that is, meant

the same thing), on a four-grade scale: *dark green* (good), *light green* (mostly good), *yellow* (mostly bad), or *red* (bad). Two different annotators checked the same sentence pairs and if the annotators were in full agreement or if they chose different but adjacent categories, the sentence pair was included in the data set. Otherwise the sentence pair was discarded.

There was an additional choice for the annotators to explicitly discard bad data. Data was to be discarded, if there were spelling mistakes, bad grammar, bad sentence segmentation, or the language of the sentences was wrong. The highest “trash rate” of around 11 % occurred for the French data, apparently because of numerous grammatical mistakes in French spelling, which is known to be tricky. The lowest “trash rate” of below 3 % occurred for Finnish, a language with highly regular orthography. Interestingly, English was second best after Finnish, with less than 4 % discarded sentence pairs. Although English orthography is not straightforward, there are few diacritics that can go wrong (such as accents on vowels), and English benefits from the largest amounts of data and the best preprocessing tools. Table 1 displays a breakdown of the error types in the English and Finnish annotated data.

Type	English		Finnish	
Not grammatical	64	(54%)	35	(36%)
OCR error	13	(11%)	22	(23%)
Wrong language	28	(24%)	12	(13%)
Actually correct	14	(12%)	27	(28%)
Total	119	(100%)	96	(100%)

Table 1: The numbers and proportions of different error types in the data discarded by the annotators. Note that some of the sentence pairs that have been discarded are actually correct and have been mistakenly removed by the annotators.

The Opusparcus training sets need to be much larger than the development and test sets in order to be useful. However, size comes at the expense of quality, and the training sets have not been checked manually. The training sets are assumed to contain noise to the same extent as the development and test sets. On one hand, when it comes to spelling and OCR errors, this may not be too bad, as a paraphrase detection model that is robust to noise is a good thing. On the other hand, when we train a supervised paraphrase de-

tection model, we would like to know which of the sentence pairs in the training data are actual paraphrases and which ones are not. Since the training data has not been manually annotated, we cannot be sure. Instead we need to rely on the automatic ranking presented by Creutz (2018) that is supposed to place the sentence pairs that are most likely to be true paraphrases first in the training set and the sentence pairs that are least likely to be paraphrases last.

In the current paper, we investigate whether it is more beneficial to use less and cleaner training data or more and noisier training data. We also compare different models in terms of their robustness to noise.

In addition to the Opusparcus data, we use other data sources. In Section 4.3 we experiment with a model trained on PPDB, a large collection of noisy, automatically extracted and ranked paraphrase candidates. PPDB has been successfully used in paraphrase models before (Wieting et al., 2015, 2016; Wieting and Gimpel, 2017), so we are interested in comparing the performance of models trained on Opusparcus and those trained on PPDB.

We also evaluate our models on MSRPC, a well-known paraphrase corpus. While Opusparcus contains mostly short sentences of conversational nature, and PPDB contains mostly short phrases and sentence fragments, the MSRPC data comes from the news domain. MSRPC was created by automatically extracting potential paraphrase candidates, which were then checked by human annotators.

Lastly, two semantic textual similarity data sets, SICK and STS14 are used for evaluation in a transfer learning setting. SICK contains sentence pairs from image captions and video descriptions annotated for relatedness with scores in the $[0, 5]$ range. It consists of about 10,000 English sentences which are descriptive in nature. STS14 comprises five different subsets, ranging over multiple genres, also with human-annotated scores within $[0, 5]$.

3 Embedding models

We use supervised training to produce sentence embedding models, which can be used to determine how similar sentences are semantically and thus if they are likely to be paraphrases.

3.1 Models

In our models, there is a sequence of words (or subword units) to be embedded: $s = (w_1, w_2, \dots, w_n)$. The embedding of a sequence s is $g(s)$, where g is the embedding function.

The word embedding matrix is $W \in \mathbb{R}^{d \times |V|}$, where d is the dimensionality of the embeddings and $|V|$ is the size of the vocabulary. W^{w_i} is used to denote the embedding for the token w_i .

We use a simple word averaging (WA) model as a baseline. In this model the phrase is embedded by averaging the embeddings of its tokens:

$$g(s) = \frac{1}{n} \sum_{i=1}^n W^{w_i}$$

Despite its simplicity, the WA model has been shown to achieve good results in a wide range of semantic textual similarity tasks. (Wieting et al., 2016)

Our second model is a variant of the gated recurrent averaging network (GRAN) introduced by Wieting and Gimpel (2017). GRAN extends the WA model with a recurrent neural network, which is used to compute gates for each word embedding before averaging. We use a gated recurrent unit (GRU) network (Cho et al., 2014). The hidden states (h_1, \dots, h_n) are computed using the following equations:

$$\begin{aligned} r_t &= \sigma(W_r W^{w_t} + U_r h_{t-1}) \\ z_t &= \sigma(W_z W^{w_t} + U_z h_{t-1}) \\ \tilde{h}_t &= z_t \circ f(W_h W^{w_t} + U_h (r_t \circ h_{t-1}) + b_h) \\ h_t &= (1 - z_t) \circ h_{t-1} + \tilde{h}_t \end{aligned}$$

Here W_r , W_z , W_h , U_r , U_z , and U_h are the weight matrices, b_h is a bias vector, σ is the sigmoid function, and \circ denotes the element-wise product of two vectors.

At each time step t we compute a gate for the word embedding and elementwise-multiply the gate with the word embedding to acquire the new word vector a_t :

$$\begin{aligned} g_t &= \sigma(W_x W^{w_t} + W_h h_t + b) \\ a_t &= W^{w_t} \circ g_t \end{aligned}$$

Here W_x and W_h are weight matrices. The final sentence embedding is computed by averaging the word vectors:

$$g(s) = \frac{1}{n} \sum_{i=1}^n a_i$$

3.2 Training

Our training data consists of pairs of sequences (s_1, s_2) and associated labels $y \in \{0, 1\}$ indicating whether the sequences are paraphrases or not. Because the Opusparcus data contains ranked paraphrase candidates and not labeled pairs, we take the following approach to sampling the data: The desired number of paraphrase pairs (positive examples) are taken from the beginning of the data sets. That is, the highest ranking pairs, which are the most likely to be proper paraphrases according to Creutz (2018), are labeled as paraphrases, although not all of them are true paraphrases. The non-paraphrase pairs (negative examples) are created by randomly pairing sentences from the training data. It is possible that a positive example is created this way by accident, but we assume the likelihood of this to be low enough for it not to have noticeable effect on performance. We sample an equal number of positive and negative pairs in all experiments. In the rest of this paper, when mentioning training set sizes, we indicate the number of (assumed) positive pairs sampled from the data. There is always an equal amount of (assumed) negative pairs.

During training we optimize the following margin-based loss function:

$$\begin{aligned} L(\theta) &= y(\max(0, m - d(g(s_1), g(s_2))))^2 \\ &\quad + (1 - y)d(g(s_1), g(s_2)) \end{aligned}$$

Here m is the margin parameter, $d(g(s_1), g(s_2))$ is the cosine distance between the embedded sequences, and g is the embedding function. The loss function penalizes negative pairs with a cosine distance smaller than the margin (first term) and encourages positive pairs to be close to each other (second term).

We use the Adam optimizer (Kinga and Ba, 2015) with a learning rate of 0.001 and a batch size of 128 samples in all experiments. Variational dropout (Gal and Ghahramani, 2016) is used for regularization in the GRAN model. The hyperparameters were tuned in preliminary experiments for development set accuracy and, with the exception of keep probability in dropout, kept constant in all experiments.

The embedding matrix W is initialized to a uniform distribution over $[-0.01, 0.01]$. In our experiments we found that initializing with pre-trained embeddings did not improve the paraphrase detection results. The layer weights in the GRU

network are initialized using Xavier initialization (Glorot and Bengio, 2010), and we use the leaky ReLU activation function.

4 Experiments

Our initial experiment addresses the effects of unsupervised morphological segmentation on the results of the paraphrase detection task.

Next, we tackle our main question on the trade-off between the amount of noise in the training data and the data size. In particular, we try to see if an optimal amount of noise can be found, and whether the different models have different demands in this respect.

Finally, we evaluate the English-language models on out-of-domain semantic similarity and paraphrase detection tasks.

All evaluations on the Opusparcus are conducted in the following manner: Each sentence in the sentence pair is embedded using the sentence encoding model. The resulting vectors are concatenated and passed on to a multi-layer perceptron classifier with a single hidden layer of 200 units. The classifier is trained on the development set, and the final results are reported on the unseen test set in terms of classification accuracy.

4.1 Segmentation

We work on six different European languages, some of which are morphologically rich (that is, the number of possible word forms in the language is high). In the case of languages like Finnish and Russian, the vocabularies without any kind of morphological preprocessing can grow very large even with small amounts of data.

In our approach we train Morfessor Baseline (Creutz and Lagus, 2002; Virpioja et al., 2013), an unsupervised morphological segmentation algorithm, on the whole Opusparcus training data available. Segmentation approaches that result in fixed-size vocabularies, such as byte-pair encoding (BPE) (Sennrich et al., 2016), have been gaining popularity in some natural language processing tasks. We decided to use Morfessor instead, which also appeared to outperform BPE in preliminary experiments. However, we will not focus on segmentation quality, but use segmentation simply as a preprocessing step to improve downstream performance.

The results are shown in the WA-M and WA columns of Table 2. The differences in perfor-

	AP	WA-M	WA	GRAN
de	74.3	77.0	82.3	83.2
en	72.8	87.4	86.4	89.2
fi	61.0	74.7	80.3	80.1
fr	68.6	74.0	76.7	76.8
ru	65.4	61.4	70.9	69.7
sv	54.8	78.1	84.1	83.2

Table 2: Classification accuracies on the Opusparcus test sets for models trained on 1 million positive sentence pairs. AP (all paraphrases) is the majority baseline, which is the accuracy obtained if all sentence pairs in the test data are labeled as paraphrases. Consistent improvement is obtained by the WA model without segmentation (WA-M: “WA without Morfessor”) and further by the WA model with segmentation. Whether the GRAN model outperforms WA is hard to tell from these figures, but this is further analyzed in Section 4.2.

mance between the WA models with segmentation (called just WA) and without segmentation (called WA-M) clearly indicate that this is a necessary preprocessing step when working on languages with complex morphology. The effect of segmentation for the GRAN model (not shown) is similar, with the exception of English also improving by a few points instead of worsening. Based on these results we will use Morfessor as a preprocessing step in all of the remaining experiments.

4.2 Data selection

We next investigate the effects of data set size and the amount of noise in the data on model performance. We are interested in finding an appropriate amount of training data to be used in training the paraphrase detection models, as well as evaluating the robustness of different models against noise in the data.

For each language, data sets containing approximately 80%, 70%, or 60% clean paraphrase pairs are created. These percentages are the proportions of assumed positive training examples; the negative examples are created using the approach outlined in Section 3.2.

Estimates of the quality of the training sets exist for all languages in Opusparcus.³ The quality estimates were used to approximate the numbers of phrase pairs corresponding to the noise levels. Because the data sets for different languages are not

³The figures used to approximate the data set sizes can be found in the presentation slides (slides 12-13) at https://helda.helsinki.fi/bitstream/handle/10138/237338/creutz20181rec_slides.pdf

equal in size, the number of phrase pairs at a certain noise level differs from language to language. The different data set sizes for all noise levels and languages are shown in Table 3.

Table 3 shows the results for the GRAN model. The results indicate that the GRAN model is rather robust to noise in the data. For five out of six languages, the best results are achieved using either the 70% or 60% data sets. That is, even when up to 40% of the positive examples in the training data are incorrectly labeled, the model is able to maintain or improve its performance.

The results for the WA model are very different. The last row of Table 3 shows the accuracies of the WA model at different levels of noise for English. The model’s performance decreases significantly as the number of noisy pairs increases, and the results are similar for the other languages as well. We hypothesize these differences to be due to differences in model complexity. The GRAN model incorporates a sequence model and contains more parameters than the simpler WA model.

4.2.1 Further analysis of differences between models

Some qualitative differences between the WA and GRAN models are illustrated in Tables 4 and 5 as well as Figure 1. Table 4 shows which ten sentences in the English development set are closest to one target sentence “*okay, you don’t get it, man.*” according to the two models. The comparison is performed by computing the cosine similarity between the sentence embedding vectors. A similar example is shown for German in Table 5: “*Kann gut sein.*” (in English: “*That may be.*”)⁴

The result suggests that the WA (word averaging) models produce “bag of synonyms”. Sentences are considered similar if they contain the same words or similar words. This, however, makes the WA model perform weakly when a sentence should not be interpreted literally word by word. German “*Kann gut sein.*” is unlikely to literally mean “*Can be good.*”; yet sentences with that meaning are at the top of the WA ranking. By contrast, the GRAN model comes up with very different top candidates, sentences expressing modality, such as: “*Possibly*”, “*Yes, he might*”, “*You’re probably right*”, “*As naturally as possible*”, and “*I think so*”.

⁴Further examples of similar sentences can be found in the supplemental material.

Figure 1 provides some additional information on the English sentence “*okay, you don’t get it, man.*”. Distributions of the cosine similarities of a much larger number of sentences have been plotted (10 million sentences from English OpenSubtitles). In the plots, similar sentences are on the right and dissimilar sentences on the left. In the case of the GRAN model we see a huge mass of dissimilar sentences smoothing out in a tail of similar sentences. In the case of the WA model, there is clearly a second, smaller bump to the right. It turns out that the “bump” mainly contains negated sentences, that is, sentences that contain synonyms of “*don’t*”. A second look at Table 4 validates this observation: the common trait of the sentences ranked at the top by WA is that they contain “*don’t*” or “*not*”. Thus, according to WA, the main criterion for a sentence to be similar to “*okay, you don’t get it, man.*” is that the sentence needs to contain negation. Again, the GRAN model stresses other, more relevant aspects, in this case, whether the sentence refers to *not knowing* or *not understanding*.

4.3 PPDB as training data

We also train the GRAN model on PPDB data. Wieting and Gimpel (2017) found that models trained on PPDB achieve good results on a wide range of semantic textual similarity tasks, thus, good performance could be expected on the Opusparcus test sets.

For English we use the PPDB 2.0 release, for languages other than English we use the 1.0 release, as the 2.0 is not available for those languages. The phrasal paraphrase packs are used for all languages. We pick the number of paraphrase pairs in such a way that the training data contains as close to an equal number of tokens as the Opusparcus training data with 1 million positive examples. This ensures that the amount of training data is as similar as possible in both settings. The training setup is otherwise identical to that outlined above.

The results are shown in Table 6. There is a significant drop in performance when moving from in-domain training data (Opusparcus) to out-of-domain training data (PPDB). One possible explanation for this is that the majority of the phrase pairs in the PPDB dataset contain sentence fragments rather than full sentences.

	1M	80%	70%	60%
de	83.2 (90%)	86.7 (4)	85.3 (6)	85.6 (12)
en	89.2 (97%)	90.2 (5)	92.1 (20)	90.9 (34)
fi	80.1 (83%)	81.4 (2.5)	82.5 (3.5)	81.5 (9)
fr	76.8 (95%)	76.2 (5)	77.1 (13)	77.9 (22)
ru	69.7 (85%)	60.3 (2)	70.3 (5)	66.8 (15)
sv	83.2 (85%)	71.7 (1.2)	73.0 (1.8)	82.1 (5)
en (WA)	86.4 (97%)	79.5 (5)	77.9 (20)	77.2 (34)

Table 3: Results on Opusparcus for GRAN (all languages) and WA (English only). The first six rows show the accuracies of the GRAN model at different estimated levels of correctly labeled positive training pairs: 80%, 70%, and 60%. In each entry in the table, the first number is the classification accuracy and the number in brackets is the number of assumed positive training pairs in millions. For comparison, the 1M column to the left repeats the values from Table 2, in which the size of the training set was the same for each language, regardless of noise levels; the estimated proportion of truly positive pairs in these setups are shown within brackets. The last row of the Table shows the performance of the WA model for English.




Figure 1: Distributions of similarity scores between the target sentence “okay, you don’t get it, man.” and 10 million English sentences from OpenSubtitles. Cosine similarity between sentence embedding vectors are used. A sentence that is very close to the target sentence has a cosine similarity close to 1, whereas a very dissimilar sentence has a value close to -1. (Some of the similarity values are below -1 because of rounding errors in Faiss: <https://github.com/facebookresearch/faiss/issues/297>.) Section 4.2.1 discusses differences in the distributions between the GRAN and WA models.

4.4 Transfer learning

We also evaluate our English models on other data sets. Because we are primarily interested in paraphrastic sentence embeddings, we choose to evaluate our models on the MSRRC paraphrase corpus, as well as two semantic textual similarity tasks, SICK-R and STS14. The data represent a range of genres, and hence offer a view of the potential of Opusparcus for out-of-domain use and transfer learning. Because of the similarities between paraphrase detection and the semantic textual similarity tasks, we believe the two tasks to be mutually supportive.

We present results for the WA model as well as the best GRAN model from Section 4.2. The eval-

uations are conducted using the SentEval toolkit (Conneau and Kiela, 2018). To obtain comparable results, we use the recommended default configuration for the SentEval parameters. The results are shown in Table 7.

We first note that our models fall short of the state-of-the-art results by a rather large margin. We hypothesize the discrepancy between the performance on MSRRC of our models and the BiLSTM-Max model of Conneau et al. (2017b) to be due to differences in the genre of training data. The conversational language of subtitles is vastly different from the news domain of MSRRC. Although the NLI data used by Conneau et al. (2017b) is derived from an image-captioning task

	okay , you don ’t get it , man .	
	you don ’t understand .	0.98
	no , you don ’t understand .	0.98
	you can ’t know that .	0.92
G	you do not really know .	0.90
R	no , i don ’t think you understand	0.88
A	you know , nobody has to know .	0.86
N	you don ’t got it .	0.82
	no one will ever know .	0.82
	and no one will know .	0.81
	we don ’t know yet .	0.81
	you don ’t got it .	0.91
	don ’t go over .	0.91
	do not beat yourself up about that .	0.90
	please don ’t .	0.89
W	well ... not everything .	0.89
A	not all of it .	0.88
	you don ’t have to .	0.87
	no , you don ’t understand .	0.87
	one it ’s not up to you .	0.86
	okay , that ’s not necessary .	0.84

Table 4: The ten most similar sentences to “*okay, you don’t get it, man.*” in the Opusparcus English development set, based on sentence embeddings produced by the GRAN and WA models, respectively. Cosine similarities are shown along with the sentences. (The annotated “correct” paraphrase is “*you don’t understand.*”)

and thus does not represent the news domain, it is at least closer to MSRPC in terms of the vocabulary and sentence structure. Most interesting is the difference between our WA model and the Paragram-phrase model of [Wieting et al. \(2016\)](#). These are essentially the same model, but trained on two different data sets. While the performance on SICK-R is comparable, our model significantly underperforms on STS14. Overall the results indicate that our models tend to overfit the domain of the Opusparcus data and consequently do not perform as well on out-of-domain data.

5 Discussion and Conclusion

Our results show that even a large amount of noise in training data is not always detrimental to model performance. This is a promising result, as automatically collected, large but noisy data sets are often easier to come by than clean, manually collected or annotated data sets. Our results can also guide model selection when noise in training data is a consideration.

	Kann gut sein .	
	Möglicherweise .	0.93
	Ja , könnte er .	0.92
	Hast wohl Recht .	0.92
G	So natürlich wie möglich .	0.91
R	Ihr habt natürlich recht .	0.91
A	Sie haben recht , natürlich .	0.88
N	Ich denke , doch .	0.88
	Ja , ich denke schon .	0.87
	Wahrscheinlich schon .	0.87
	Ich bin mir sicher .	0.87
	Das ist doch gut .	0.83
	Na , das ist gut .	0.81
	Ist in Ordnung .	0.81
	Dir geht es gut .	0.81
W	Ihnen geht es gut .	0.81
A	Sie ist in Ordnung .	0.81
	Ich kann es fühlen .	0.80
	Es ist alles gut .	0.79
	Mir geht ’s gut .	0.79
	Sie is okay .	0.79

Table 5: The ten most similar sentences to “*Kann gut sein.*” in the Opusparcus German development set, based on sentence embeddings produced by the GRAN and WA models, respectively. The annotated “correct” paraphrase is here “*Wahrscheinlich schon.*” (“*Probably yes*”).

In future work we would like to explore how to most effectively leverage possibly noisy paraphrase data in learning general-purpose sentence embeddings for a wide range of transfer tasks. Investigating training procedures and encoding architectures that allow for robust models with the capability for generalization is a topic for future research.

	GRAN
de	78.1
en	83.4
fi	70.4
fr	74.8
ru	67.7
sv	76.4

Table 6: Results on Opusparcus test sets for models trained on PPDB.

	MSRPC	SICK-R	STS14
GRAN	69.5/80.6	.717	.40/.44
WA	67.1/79.1	.710	.54/.53
BiLSTM-Max	76.2/83.1	.884	.70/.67
Paragram-phrase	-	.716	.71/-
FastSent	72.2/80.3	-	.63/.64

Table 7: Transfer learning results on MSRPC, SICK-R and STS14. GRAN and WA denote our models. We also show results for a selection of models from the transfer learning literature. We use the evaluation measures that are customarily used in connection with these data sets. For MSRPC, the accuracy (left) and F1-score (right) are reported. For SICK-R we report Pearson’s r , and for STS14 Pearson’s r (left) and Spearman’s rho (right). For all these measures a higher value indicates a better result.

References

- 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.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder–decoder approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111. Association for Computational Linguistics.
- Alexis Conneau and Douwe Kiela. 2018. Senteval: An evaluation toolkit for universal sentence representations. *arXiv preprint arXiv:1803.05449*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017a. 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. Association for Computational Linguistics.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017b. 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. Association for Computational Linguistics.
- Mathias Creutz. 2018. Open Subtitles Paraphrase Corpus for Six Languages. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Mathias Creutz and Krista Lagus. 2002. Unsupervised discovery of morphemes. In *Proceedings of the ACL workshop on Morphological and Phonological Learning (SIGPHON)*, pages 21–30, Philadelphia, PA, USA.
- Bill Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*. Asia Federation of Natural Language Processing.
- Bill Dolan, Chris Quirk, and Chris Brockett. 2004. Unsupervised construction of large paraphrase corpora: Exploiting massively parallel news sources. In *Proceedings of the 20th International Conference on Computational Linguistics, COLING ’04*, Geneva, Switzerland. Association for Computational Linguistics.
- Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 1019–1027. Curran Associates, Inc.
- Juri Ganitkevitch and Chris Callison-Burch. 2014. The multilingual paraphrase database. In *The 9th edition of the Language Resources and Evaluation Conference*, Reykjavik, Iceland. European Language Resources Association.
- Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2013. PPDB: The paraphrase database. In *Proceedings of NAACL-HLT*, pages 758–764, Atlanta, Georgia. Association for Computational Linguistics.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 249–256.
- 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. Association for Computational Linguistics.
- Diererik P. Kinga and Jimmy Lei Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 3294–3302. Curran Associates, Inc.

- Pierre Lison and Jörg Tiedemann. 2016. OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC 2016)*, Portorož, Slovenia.
- Pierre Lison, Jörg Tiedemann, and Milen Kouylekov. 2018. OpenSubtitles2018: Statistical Rescoring of Sentence Alignments in Large, Noisy Parallel Corpora. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, Roberto Zamparelli, et al. 2014. A sick cure for the evaluation of compositional distributional semantic models. In *Proceedings of the 9th International Conference on Language Resources and Evaluation (LREC 2014)*, Reykjavik, Iceland.
- Gustavo Henrique Paetzold and Lucia Specia. 2016. Collecting and exploring everyday language for predicting psycholinguistic properties of words. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 669–1679, Osaka, Japan.
- Ellie Pavlick, Pushpendre Rastogi, Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2015. PPDB 2.0: Better paraphrase rankings, fine-grained entailment relations, word embeddings, and style classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Short Papers)*, pages 425–430, Beijing, China. Association for Computational Linguistics.
- Chris Quirk, Chris Brockett, and William B. Dolan. 2004. Monolingual machine translation for paraphrase generation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP2004)*, pages 142–149, Barcelona, Spain.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725. Association for Computational Linguistics.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J. Pal. 2018. Learning general purpose distributed sentence representations via large scale multi-task learning. In *International Conference on Learning Representations*.
- Jörg Tiedemann. 2007. Building a multilingual parallel subtitle corpus. In *Proceedings of the 17th Conference on Computational Linguistics in the Netherlands (CLIN 17)*, Leuven, Belgium.
- Jörg Tiedemann. 2008. Synchronizing translated movie subtitles. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC 2008)*, Marrakech, Morocco.
- Jörg Tiedemann. 2016. Finding alternative translations in a large corpus of movie subtitles. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC 2016)*, Portorož, Slovenia.
- Sami Virpioja, Peter Smit, Stig-Arne Grönroos, and Mikko Kurimo. 2013. Morfessor 2.0: Python implementation and extensions for Morfessor Baseline. Technical Report 25/2013, Aalto University publication series SCIENCE + TECHNOLOGY, Aalto University, Helsinki.
- Marlies van der Wees, Arianna Bisazza, and Christof Monz. 2016. Measuring the effect of conversational aspects on machine translation quality. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2571–2581, Osaka, Japan.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2015. From paraphrase database to compositional paraphrase model and back. *Transactions of the Association for Computational Linguistics*, 3:345–358.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2016. Towards universal paraphrastic sentence embeddings. In *Proceedings of the International Conference on Learning Representations*.
- 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. Association for Computational Linguistics.