

PAWS-X: A Cross-lingual Adversarial Dataset for Paraphrase Identification

Yinfei Yang*, Yuan Zhang*, Chris Tar and Jason Baldridge

Google Research, Mountain View, 94043

{yinfeiy, zhangyua, ctar, jasonbaldridge}@google.com

Abstract

Most existing work on adversarial data generation focuses on English. For example, PAWS (Paraphrase Adversaries from Word Scrambling) (Zhang et al., 2019) consists of challenging English paraphrase identification pairs from Wikipedia and Quora. We remedy this gap with PAWS-X, a new dataset of 23,659 *human* translated PAWS evaluation pairs in six typologically distinct languages: French, Spanish, German, Chinese, Japanese, and Korean. We provide baseline numbers for three models with different capacity to capture non-local context and sentence structure, and using different multilingual training and evaluation regimes. Multilingual BERT (Devlin et al., 2019) fine-tuned on PAWS English plus machine-translated data performs the best, with a range of 83.1-90.8 accuracy across the non-English languages and an average accuracy gain of 23% over the next best model. PAWS-X shows the effectiveness of deep, multilingual pre-training while also leaving considerable headroom as a new challenge to drive multilingual research that better captures structure and contextual information.

1 Introduction

Adversarial examples have effectively highlighted the deficiencies of state-of-the-art models for many natural language processing tasks, e.g. question answering (Jia and Liang, 2017; Chen et al., 2018; Ribeiro et al., 2018), textual entailment (Zhao et al., 2018; Glockner et al., 2018), and text classification (Alzantot et al., 2018; Iyyer et al., 2018). Zhang et al. (2019) introduce PAWS, which has adversarial paraphrase identification pairs with high lexical overlap, like *flights from New York to Florida* and *flights from Florida to New York*. Such pairs stress the importance of modeling sentence structure and context because they have high

word overlap ratio but different semantic meaning. In addition to revealing failures of state-of-the-art models, research on adversarial examples has generally shown that augmenting training data with good adversarial examples can boost performance for some models—providing greater clarity to the modeling landscape as well providing new headroom for further improvements.

Most previous work focuses only on English despite the fact that the problems highlighted by adversarial examples are shared by other languages. Existing multilingual datasets for paraphrase identification, e.g. Multi30k (Elliott et al., 2016) and Opusparcus (Creutz, 2018), lack challenging examples like PAWS. The lack of high-quality adversarial examples in other languages makes it difficult to benchmark model improvements. We bridge this gap by creating Cross-lingual PAWS (PAWS-X), an extension of the Wikipedia portion of the PAWS evaluation and test examples to six languages: Spanish, French, German, Chinese, Japanese, and Korean. This new corpus consists of 23,659 *human* translated example pairs with paraphrase judgments in each target language. Like previous work on multilingual corpus creation (Conneau et al., 2018), we machine translate the original PAWS English training set (49,401 pairs). Note that all translated pairs still have high word overlap and they inherit semantic similarity labels from the original PAWS examples; thus, the resulting dataset preserves the ability of probing structure and context sensitivity for models. We also machine translate the evaluation pairs of each language into English to establish the baseline performance of a translate-then-predict strategy. The PAWS-X dataset, including both the new human translated pairs and the machine translated examples, is available for download at <https://github.com/google-research-datasets/paws>.

*equal contribution

Language	Text
<i>Original Pair (id: 000005309_9438, label: not-paraphrasing)</i>	
en	However, in order to defeat Slovak, Derek must become a vampire attacker. However, in order to become Slovak, Derek must defeat a vampire assassin.
<i>Human Translated Pairs</i>	
fr	Toutefois, pour battre Slovak, Derek doit devenir un vampire attaquant. Cependant, pour devenir Slovak, Derek doit vaincre un vampire assassin.
es	Sin embargo, para derrotar a Slovak, Derek debe convertirse en un atacante vampiro. Sin embargo, para poder convertirse en Slovak, Derek debe derrotar a un asesino de vampiros.
de	Um Slovak zu besiegen, muss Derek jedoch zum Vampirjäger werden. Um jedoch Slowake zu werden, muss Derek einen Vampirjäger besiegen.
zh	但为击败斯洛伐克，德里克必须成为吸血鬼攻击者。 然而，为了成为斯洛伐克人，德里克必须击败吸血鬼刺客。
ja	ただし、スロバークを倒すためには、デレクは吸血鬼アタッカーになる必要があります。 しかし、デレクがスロバークになるには、バンパイア・アサシンを倒さなければならない。
ko	하지만 Slovak이 되기 위해 Derek은 반드시 뱀파이어 암살자를 물리쳐야만 합니다. 하지만 Slovak을 물리치기 위해 Derek은 뱀파이어 사냥꾼이 되어야만 했습니다.

Table 1: Examples of human translated pairs for each of the six languages.

Our experiments show that PAWS-X effectively measures the multilingual adaptability of models and how well they capture context and word order. The state-of-the-art multilingual BERT model (Devlin et al., 2019) obtains a 32% (absolute) accuracy improvement over a bag-of-words model. We also show that machine translation helps and works better than a zero-shot strategy. We find that performance on German, French, Spanish is overall better than Chinese, Japanese and Korean.

2 PAWS-X Corpus

The core of our corpus creation procedure is to translate the Wikipedia portion of the original PAWS corpus from English (en) to six languages: French (fr), Spanish (es), German (de), Chinese (zh), Japanese (ja), and Korean (ko). To this end, we hire human translators to translate the development and test sets, and use a neural machine translation (NMT) service¹ to translate the training set.

We choose translation instead of repeating the PAWS data generation approach (Zhang et al., 2019) to other languages. This has at least three advantages. First, human translation does not require high-quality multilingual part-of-speech taggers or named entity recognizers, which play a key role in the data generation process used in Zhang et al. (2019). Second, human translators are trained to produce the target sentence while preserving meaning, thereby ensuring high data quality. Third, the resulting data can provide a new testbed for cross-lingual transfer techniques because examples in all languages are translated from the same sources. For example, PAWS-X

¹<https://cloud.google.com/translate/>

	fr	es	de	zh	ja	ko
dev	1,992	1,962	1,932	1,984	1,980	1,965
test	1,985	1,999	1,967	1,975	1,946	1,972

Table 2: Examples translated per language.

could be used to evaluate whether a German or French sentence is a paraphrase of a Chinese or Japanese one.

Translating Evaluation Sets We obtain human translations on a random sample of 4,000 sentence pairs from the PAWS development set for each of the six languages (48,000 translations). The manual translation is performed by 10-20 in-house professionals that are native speakers of each language. A randomly sampled subset is presented and validated by a second worker. The final delivery is guaranteed to have less than 5% word level error rate. The sampled 4,000 pairs are split into new development and test sets, 2,000 pairs for each.

Due to time and cost constraints, we could not translate all 16,000 examples in both of original PAWS development and test set. Each sentence in a pair is presented independently so that translation is not affected by context. In our initial studies we noticed that sometimes it was difficult to translate an entity mention. We therefore ask translators to translate entity mentions, but different translators may have different preferences according to their background knowledge. Table 1 gives example translated pairs in each language.

Resulting Corpus Some sentences could not be translated. Table 2 shows the final counts translated to each language. Most of the untranslated

sentences were due to incompleteness or ambiguities, such as *It said that Easipower was*, and *Park Green took over No*. These sentences are likely from the adversarial generation process when creating PAWS. On average less than 2% of the pairs are not translated, and we simply exclude them.

The authors further verified translation quality for a random sample of ten pairs in each language. PAWS-X includes 23,459 human-translated pairs, including 11,815 and 11,844 pairs in development and test, respectively. Finally, original PAWS labels (paraphrase or not paraphrase) are mapped to the translations. Positive pairs account for 44.0% of development sets and 45.4% of test respectively—close to the PAWS label distribution.

Translation brings new challenges to the paraphrasing identification task. An entity can be translated differently, such as *Slovak* and *Slowake* (Table 1) and models need to capture that these refer to the same entity. In a more challenging example, *Four Rivers*, *Audubon* and *Shawnee Trails* are translated in just one of the sentences:

- en** s1 From the merger of the Four Rivers Council and the Audubon Council, the Shawnee Trails Council was born.
s2 Shawnee Trails Council was formed from the merger of the Four Rivers Council and the Audubon Council.
- zh** s1 Four Rivers 委员会与 Audubon 委员会合并后, Shawnee Trails 委员会得以问世.
s2 肖尼小径(Shawnee Trails) 委员会由合并四河(Four Rivers) 委员会和 奥杜邦(Audubon) 委员会成立.

In the zh-s2 example, the parentheses give English glosses of Chinese entity mentions.

3 Evaluated Methods

The goal of PAWS-X is to probe models’ ability to capture structure and context in a multilingual setting. We consider three models with varied complexity and expressiveness. The first baseline is a simple bag-of-words (**BOW**) encoder with cosine similarity. It uses unigram to bigram token encoding as input features and takes a cosine value above 0.5 as a paraphrase. The second model is **ESIM**, Enhanced Sequential Inference Model (Chen et al., 2017). Following Zhang et al. (2019), ESIM encodes each sentence using a BiLSTM, and passes the concatenation of encodings through a feed-forward layer for classification. The additional layers allow ESIM to capture more complex sentence interaction than cosine similarity. Third,

	BOW	ESIM	BERT
Non-local context	×	✓	✓
Word interaction	×	×	✓
Translate Train	✓	✓	✓
Translate Test	✓	✓	✓
Zero Shot	×	×	✓
Merged	×	×	✓

Table 3: Complexity of each evaluated model and the training/evaluation strategies being tested.

we evaluate **BERT**, Bidirectional Encoder Representations from Transformers (Devlin et al., 2019), which recently achieved state-of-the-art results on eleven natural language processing tasks.

We evaluate all models with two strategies (Conneau et al., 2018): (1) **Translate Train**: the English training data is machine-translated into each target language to provide data to train each model and (2) **Translate Test**: train a model using the English training data, and machine-translate all test examples to English for evaluation.

Multilingual BERT is a single model trained on 104 languages, which enables experiments with cross-lingual training regimes. (1) **Zero Shot**: the model is trained on the PAWS English training data, and then directly evaluated on all others. Machine translation is **not** involved in this strategy. (2) **Merged**: train a multilingual model on all languages, including the original English pairs and machine-translated data in all other languages.

Table 3 summarizes the models with respect to whether they represent non-local contexts or support cross-sentential word interaction, plus which strategies are evaluated for each model.

4 Experiments and Results

We use the latest public multilingual BERT base model with 12 layers² and apply the default fine-tuning strategy with batch size 32 and learning rate 1e-5. For BOW and ESIM, we use our own implementations and 300 dimensional multilingual word embeddings from fastText.³ We allow fine-tuning word embeddings during training, which gives better empirical performance.

We use two metrics: classification accuracy and area-under-curve scores of precision-recall curves (AUC-PR). For BERT, probability scores for the positive class is used to compute AUC-PR. For BOW and ESIM a cosine threshold of 0.5 is used to compute accuracy. In all experiments, the best

²<http://goo.gl/language/bert>

³<https://fasttext.cc/>

Method	Accuracy							AUC-PR						
	en	fr	es	de	zh	ja	ko	en	fr	es	de	zh	ja	ko
BOW														
Translate Train	55.8	51.7	47.9	50.2	54.5	55.1	56.7	41.1	48.9	46.8	46.4	50.0	48.7	49.3
Translate Test	–	54.9	54.7	55.2	55.3	55.9	55.2	–	46.3	45.5	45.8	50.9	46.8	48.5
ESIM														
Translate Train	67.2	66.2	66.0	63.7	60.3	59.6	54.2	69.6	67.0	64.2	59.2	58.2	56.3	50.5
Translate Test	–	66.2	66.3	66.0	62.0	62.3	60.6	–	68.4	69.5	68.2	62.3	61.8	60.3
BERT														
Translate Train	93.5	89.3	89.0	85.3	82.3	79.2	79.9	97.1	93.6	92.4	92.0	87.4	81.4	82.4
Translate Test	–	88.7	89.3	88.4	79.3	75.3	72.6	–	93.8	93.1	92.9	85.1	80.9	80.1
Zero shot	–	85.2	86.0	82.2	75.8	70.5	71.7	–	91.0	90.5	89.4	79.6	72.7	75.5
Merged	93.8	90.8	90.7	89.2	85.4	83.1	83.9	96.5	94.0	92.9	92.9	88.9	86.0	86.3

Table 4: Accuracy (%) and AUC-PR (%) of each approach. Best numbers in each column are marked in bold.

Method		Averaged							
		Accuracy	AUC-PR	#	0	1-2	3-4	5-6	7
BOW	Translate Train	52.7	48.4						
	Translate Test	55.2	47.3						
ESIM	Translate Train	61.7	59.2						
	Translate Test	63.9	65.1						
BERT	Translate Train	84.2	88.2						
	Translate Test	82.3	87.6						
	Zero Shot	78.6	83.1						
	Merged	87.2	90.2						

Table 5: Average Accuracy (%) and AUC-PR (%) over the six languages.

model checkpoint is chosen based on accuracy on development sets and report results on testing sets.

Results Table 4 shows the performance of all methods and languages. Table 5 summarizes the average results for the six non-English languages.

Model Comparisons: On both Translate Train and Translate Test, BERT consistently outperforms both BOW and ESIM by a substantial margin (>15% absolute accuracy gains) across all seven languages. BERT Translate Train achieves an average 20% accuracy gain. This result demonstrates that PAWS-X effectively measures models’ sensitivity to word order and syntactic structure.

Training/Evaluation Strategies: As Table 4 and 5 show, the Zero Shot strategy yields the lowest performance compared to other strategies on BERT. This is evidence that machine-translated data helps in the multilingual scenario. Indeed, when training on machine-translated examples in all languages (Merged), the model achieves the best performance, with 8.6% accuracy and 7.1% AUC-PR average gains over Zero Shot.

BERT and ESIM show different performance patterns on Translate Train and Translate Test. Translate Test appears to give consistently better performance than Translate Train on ESIM, but not on BERT. This may be because multilin-

	0	1-2	3-4	5-6	7
#	32	52	140	542	1234
%	1.6	2.6	7.0	27.1	61.7

Table 6: The count of examples by number of languages (of 7) that agree with the gold label in test set.

gual BERT is pre-trained on over one hundred languages; hence BERT provides better initialization for non-English languages than ESIM (which relies on fastText embeddings). The gap between training on English and on other languages is therefore smaller on BERT than on ESIM, which makes Translate Train work better on BERT.

Language Difference: Across all models and approaches, performance on Indo-European languages (German, French, Spanish) is consistently better than CJK (Chinese, Japanese, Korean). The performance difference is particularly noticeable on Zero Shot. This can be explained from two perspectives. First, the MT system we used works better on Indo-European languages than on CJK. Second, the CJK family is more typologically and syntactically different from English. For example, in table 1, *Slowake* in German is much closer to the original term *Slovak* in English, compared with its Chinese translation 斯洛伐克. This at least partly explains why performance on CJK is particularly poor in Zero Shot.

Error Analysis: To gauge the difficulty of each example for the best model (BERT-merged), Table 6 shows the count of examples based on how many languages for the same pair are assigned the correct label in test set. The majority of the examples are easy, with 61.7% correct in all languages. Of the 32 examples that failed in all languages, most are hard or highly ambiguous. Some have incorrect gold labels or were generated incorrectly in the original PAWS data.

The following is a sample of these.

- a1 On July 29, 1791, Sarah married Lea Thomas Wright Hill (1765–1842) at St. Martin’s Church in Birmingham and had 8 children.
- a2 Thomas Wright Hill married Sarah Lea (1765–1842) on 29 July 1791 at St Martin’s Church, Birmingham and had 8 children. **match**
- b1 He established himself eventually in the northwest of Italy, apparently supported by Guy, where he probably comes “title”.
- b2 He eventually established himself in northwestern Italy, apparently supported by Guy, where he probably received the title of “comes”. **not_match**

We also considered examples that are correctly predicted in just half of the languages. Some of these failed because of translation noise, e.g. inconsistent entity translations (as shown in §2).

5 Conclusion

We introduce PAWS-X, a challenging paraphrase identification dataset with 23,659 *human* translated evaluation pairs in six languages. Our experimental results showed that PAWS-X effectively measures sensitivity of models to word order and the efficacy of cross-lingual learning approaches. It also leaves considerable headroom as a new challenging benchmark to drive multilingual research on the problem of paraphrase identification.

Acknowledgments

We would like to thank our anonymous reviewers and the Google AI Language team, especially Luheng He, for the insightful comments that contributed to this paper. Many thanks also to the translate team, especially Mengmeng Niu, for the help with the annotations.

References

- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. [Generating natural language adversarial examples](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, Brussels, Belgium. Association for Computational Linguistics.
- Hongge Chen, Huan Zhang, Pin-Yu Chen, Jinfeng Yi, and Cho-Jui Hsieh. 2018. [Attacking visual language grounding with adversarial examples: A case study on neural image captioning](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2587–2597, Melbourne, Australia. Association for Computational Linguistics.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. [Enhanced LSTM for natural language inference](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668, Vancouver, Canada. Association for Computational Linguistics.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Mathias Creutz. 2018. [Open subtitles paraphrase corpus for six languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*, Miyazaki, Japan. European Languages Resources Association (ELRA).
- 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.
- Desmond Elliott, Stella Frank, Khalil Sima’an, and Lucia Specia. 2016. [Multi30K: Multilingual English-German image descriptions](#). In *Proceedings of the 5th Workshop on Vision and Language*, pages 70–74, Berlin, Germany. Association for Computational Linguistics.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. [Breaking NLI systems with sentences that require simple lexical inferences](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 650–655, Melbourne, Australia. 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.
- Robin Jia and Percy Liang. 2017. [Adversarial examples for evaluating reading comprehension systems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.

- Marco Tullio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. [Semantically equivalent adversarial rules for debugging NLP models](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 856–865, Melbourne, Australia. Association for Computational Linguistics.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. [PAWS: Paraphrase adversaries from word scrambling](#). 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 1298–1308, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zhengli Zhao, Dheeru Dua, and Sameer Singh. 2018. [Generating natural adversarial examples](#). In *International Conference on Learning Representations*.