

# PAWS: Paraphrase Adversaries from Word Scrambling

Yuan Zhang   Jason Baldridge   Luheng He

Google AI Language

{zhangyua, jridge, luheng}@google.com

## Abstract

Existing paraphrase identification datasets lack sentence pairs that have high lexical overlap without being paraphrases. Models trained on such data fail to distinguish pairs like *flights from New York to Florida* and *flights from Florida to New York*. This paper introduces PAWS (Paraphrase Adversaries from Word Scrambling), a new dataset with 108,463 well-formed paraphrase and non-paraphrase pairs with high lexical overlap. Challenging pairs are generated by controlled word swapping and back translation, followed by fluency and paraphrase judgments by human raters. State-of-the-art models trained on existing datasets have dismal performance on PAWS (<40% accuracy); however, including PAWS training data for these models improves their accuracy to 85% while maintaining performance on existing tasks. In contrast, models that do not capture non-local contextual information fail even with PAWS training examples. As such, PAWS provides an effective instrument for driving further progress on models that better exploit structure, context, and pairwise comparisons.

## 1 Introduction

Word order and syntactic structure have a large impact on sentence meaning. Even small perturbation in word order can completely change interpretation. Consider the following related sentences.

- (1) Flights from New York to Florida.
- (2) Flights to Florida from NYC.
- (3) Flights from Florida to New York.

All three have high bag-of-words (BOW) overlap. However, (2) is a paraphrase of (1), while (3) has a very different meaning from (1).

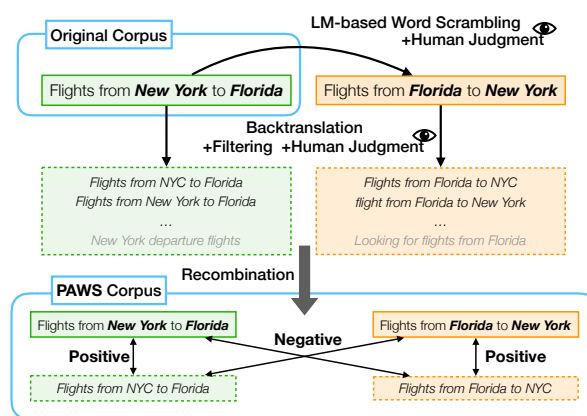


Figure 1: PAWS corpus creation workflow.

Existing datasets lack non-paraphrase pairs like (1) and (3). The Quora Question Pairs (QQP) corpus contains 400k real world pairs, but its negative examples are drawn primarily from related questions. Few have high word overlap, and of the  $\sim 1,000$  pairs with the same BOW, only 20% are not paraphrases. This provides insufficient representative examples to evaluate models' performance on this problem, and there are too few examples for models to learn the importance of word order. Table 1 shows that models trained on QQP are inclined to mark any sentence pairs with high word overlap as paraphrases despite clear clashes in meaning. Models trained or evaluated with only this data may not perform well on real world tasks where such sensitivity is important.

To address this, we introduce a workflow (outlined in Figure 1) for generating pairs of sentences that have high word overlap, but which are balanced with respect to whether they are paraphrases or not. Using this process, we create PAWS (Paraphrase Adversaries from Word Scrambling), a dataset constructed from sentences in Quora and

Sentence 1	Sentence 2	Gold	BOW	BERT	BERT+PAWS
(1) Can a bad person become good?	Can a good person become bad?	N	Y	Y	N
(2) Which is the cheapest flight from anywhere in South America to Europe?	Which is the cheapest flight from anywhere in Europe to South America?	N	Y	N	N
(3) “Taunton Castle” was on August 1 in Rio de Janeiro and on October 31 in Penang.	“Taunton Castle” was at Penang on 1 August and Rio de Janeiro on 31 October.	N	Y	Y	N
(4) Although interchangeable, the body pieces on the 2 cars are not similar.	Although similar, the body parts are not interchangeable on the 2 cars.	N	Y	Y	N
(5) Katz was born in Sweden in 1947 and moved to New York City at the age of 1.	Katz was born in 1947 in Sweden and moved to New York at the age of one.	Y	Y	Y	Y
(6) It was not the sales manager who hit the bottle that day, but the office worker with the serious drinking problem.	That day the office manager, who was drinking, hit the problem sales worker with a bottle, but it was not serious.	N	Y	Y	N

Table 1: Paraphrase/Non-paraphrase (Y/N) pairs with high bag-of-words (BOW) overlap. (1)-(5) are drawn from PAWS, and (6) is from [Mitchell and Lapata \(2008\)](#). Both a simple BOW and the state-of-the-art BERT ([Devlin et al., 2018](#)) models, if trained/fine-tuned on the Quora Question Pairs (QQP) corpus, classify all of them (with one exception) as paraphrases (Y). A BERT model fine-tuned on both QQP and PAWS examples (BERT+PAWS), however, is able to get them correct.

Wikipedia. Examples are generated from controlled language models and back translation, and given five human ratings each in both phases. A final rule recombines annotated examples and balances the labels. Our final PAWS dataset will be released publicly with 108,463 pairs at <https://g.co/dataset/paws>.

We show that existing state-of-the-art models fail miserably on PAWS when trained on existing resources, but some perform well when given PAWS training examples. BERT ([Devlin et al., 2018](#)) fine-tuned on QQP achieves over 90% accuracy on QQP, but only 33% accuracy on PAWS data in the same domain. However, the accuracy on PAWS boosts to 85% by including 12k PAWS training pairs (without reducing QQP performance). Table 1 also shows that the new model is able to correctly classify challenging pairs. Annotation scale is also important: our learning curves show strong models like BERT improve with tens of thousands of training examples.

Our experimental results also demonstrate that PAWS effectively measures sensitivity of models to word order and structure. Unlike BERT, a simple BOW model fails to learn from PAWS training examples, demonstrating its weakness at capturing non-local contextual information. Our experiments show that the gains from PAWS examples correlate with the complexity of models.

## 2 Related Work

Existing data creation techniques have focused on collecting paraphrases, e.g. from co-captions for

images ([Lin et al., 2014](#)), tweets with shared URLs ([Lan et al., 2017](#)), subtitles ([Creutz, 2018](#)), and back translation ([Iyyer et al., 2018](#)). Unlike all previous work, we emphasize the collection of challenging negative examples.

Our work closely relates to the idea of crafting adversarial examples to break NLP systems. Existing approaches mostly focused on adding label-preserving perturbations to inputs, but with the effect of distracting systems from correct answers. Example perturbation rules include adding noise to inputs ([Jia and Liang, 2017](#); [Chen et al., 2018](#)), word replacements ([Alzantot et al., 2018](#); [Ribeiro et al., 2018](#)), and syntactic transformation ([Iyyer et al., 2018](#)). A notable exception is [Glockner et al. \(2018\)](#): they generated both entailment and contradiction examples by replacing words with their synonyms or antonyms. Our work presents two main departures. We propose a novel method that generates challenging examples with balanced class labels and more word reordering variations than previous work. In addition, we release to public a large set of 108k example pairs with high-quality human labels. We believe the new dataset will benefit future research on both adversarial example generation and improvement of model robustness.

In our work, we demonstrate the importance of capturing non-local contextual information in the problem of paraphrase identification. This relates to prior work on probing sentence representations for their linguistic properties, such as how much syntactic information is encoded in representations ([Conneau et al., 2018](#); [Tenney et al.,](#)

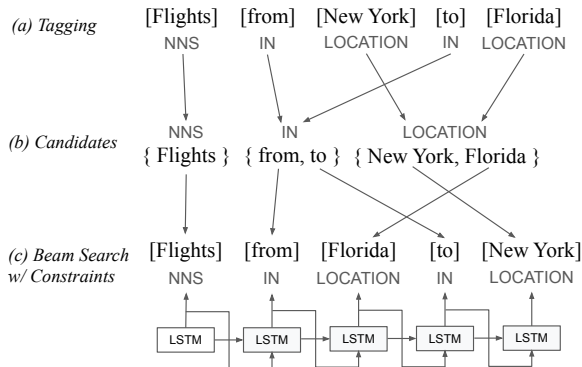


Figure 2: Illustration of the generation method in three steps. (a) Tag words and phrases with part-of-speech (POS) and named entities. (b) Build candidate sets by grouping words and phrases with the same tag. (c) Under the constraints of tag sequence template and candidate sets, find sentences with high language model scores using beam search.

2019; Ettinger et al., 2018). There also exists prior work that directly uses structural information in modeling (Filice et al., 2015; Liu et al., 2018). All these prior approaches were evaluated on existing datasets. In contrast, we perform studies on PAWS, a new dataset that emphasizes the importance of capturing structural information in representation learning. While developing new models is beyond the scope of this paper, this new dataset can facilitate research in this direction.

### 3 PAWS Example Generation

We define a **PAWS pair** to be a pair of sentences with high bag-of-words (BOW) overlap but different word order. In the Quora Question Pairs corpus, 80% of such pairs are paraphrases. Here, we describe a method to automatically generate non-trivial and well-formed PAWS pairs from real-world text in any domain (this section), and then have them annotated by human raters (Section 4).

Our automatic generation method is based on two ideas. The first swaps words to generate a sentence pair with the same BOW, controlled by a language model. The second uses back translation to generate paraphrases with high BOW overlap but different word order. These two strategies generate high-quality, diverse PAWS pairs, balanced evenly between paraphrases and non-paraphrases.

#### 3.1 Word Swapping

Our first phase generates well-formed sentences by swapping words in real world text. Most text generation models rely on large amount of training

data (Iyyer et al., 2018; Guu et al., 2018; Gupta et al., 2018; Li et al., 2018), which is unfortunately not available in our case. We thus propose a novel generation method based on language modeling and constrained beam search. The goal is to find a sentence that achieves high language model score as well as satisfying all constraints. High scores indicate that generated sentences are natural and well-formed, and constraints ensure generated pairs have the same BOW.

Figure 2 illustrates the generation procedure. First, given an input sentence, a CRF-based part-of-speech tagger tags each word. We further detect person names, locations, and organizations using a named entity recognizer, and replace POS with entity tags if probability scores are above 95%.<sup>1</sup> The sequence of tags of words and phrases form a template for the input.

Our beam search method then fills in each slot of the template from left to right, scoring each state by a language model trained on one billion words (Chelba et al., 2014). The candidate words and phrases for each slot are drawn from the input based on its tag. In Figure 2, for example, the second slot must be filled with a LOCATION from two candidate *New York* and *Florida*. Candidates are drawn without replacement so the generated sentence and the input have exactly the same bag-of-words. Note that this template-based constraint is more restrictive than the BOW requirement, but we choose it because it significantly reduces the search space. With this constraint, the method achieves high generation quality without a large beam. In practice, beam size is set to 100, which produces near-optimal results in most cases.

Let  $s'$  be the best sentence in the beam other than the input sentence  $s$ , and  $LM(\cdot)$  be their log-likelihood by the language model. We take  $(s, s')$  as a good word-swapping pair if  $LM(s') \geq LM(s) - t$ .<sup>2</sup> We manually pick the threshold  $t=3.0$  for a good balance between generation quality and coverage. Examples (1) and (2) in Table 2 are representative examples from this generation method.

<sup>1</sup>We pick this threshold to achieve about 95% precision.

<sup>2</sup>In a preliminary stage, we noticed that many pairs were simply a permutation of a list, like “A and B” changed to “B and A”. For the diversity of the dataset, 99% of these are pruned via hand-crafted, heuristic rules.

Sentence 1	Sentence 2	Generation Type
(1) Can a <b>bad</b> person become <b>good</b> ?	Can a <b>good</b> person become <b>bad</b> ?	Adjective swap
(2) <b>Jerry</b> looks over <b>Tom</b> 's shoulder and gets punched.	<b>Tom</b> looks over <b>Jerry</b> 's shoulder and gets punched.	Named entity swap
(3) The team also toured in Australia <b>in 1953</b> .	<b>In 1953</b> , the team also toured in Australia.	Temporal phrase swap
(4) Erikson <b>formed</b> the rock band Spooner with two fellow musicians.	Erikson <b>founded</b> the rock band Spooner with two fellow musicians.	Word replacement

Table 2: Examples of typical types of generation. (1) and (2) are from the word swapping method, while (3) and (4) are from the back translation method. Boldface indicates changes in each example.

### 3.2 Back Translation

Because word order impacts meaning, especially in English, the swapping method tends to produce non-paraphrases. Our preliminary results showed that the distribution of paraphrase to non-paraphrases from this method is highly imbalanced (about 1:4 ratio). However, we seek to create a balanced dataset, so we use an additional strategy based on back translation—which has the opposite label distribution and also produces greater diversity of paraphrases while still maintaining a high BOW overlap.

The back translation method takes a sentence pair and label  $(s_1, s_2, l)$  as input. For each sentence, the top- $k$  translations are obtained from an English-German neural machine translation model (NMT); then each of these is translated back to English using another German-English NMT model, providing a resulting top- $k$  results. We chose German as the pivot language because it produced more word reordering variations than other languages and the translation quality was good. Both models have the same architecture (Wu et al., 2016) and are trained on WMT14. This results in  $k^2$  back translations before deduplication. We chose  $k=5$ . To obtain more pairs with the PAWS property, we further filter back translations by their BOW similarities to the input and their word-order inversion rates, as described below.

We define BOW similarity as the cosine similarity  $\alpha$  between the word count vectors of a sentence pair. Pairs generated from the swapping strategy have score  $\alpha = 1.0$ , but here we relax the threshold to 0.9 because it brings more data diversity and higher coverage, while still generating paraphrases of the input with high quality.

To define the word-order inversion rate, we first compute word alignments between a sentence pair in a heuristic way by assuming they are one-to-one mapping and are always monotonic. For example, if the first sentence has three instances of *dog* and

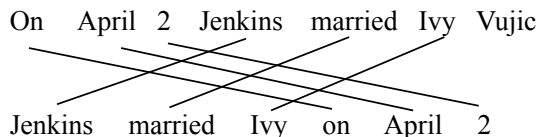


Figure 3: An example of how to compute inversion rate.

the second has two, we align the first two instances of *dog* in the same order and skip the third one. The inversion rate is then computed as the ratio of cross alignments. Figure 3 is an example pair with six alignments. There are 15 alignment pairs in total and 9 of them are crossed, e.g. alignments of *on* and *married*. The inversion rate of this example is therefore  $9/15 = 0.6$ . We sample back translation results such that at least half of the pairs have inversion rate over 0.02; this way, the final selected pairs cover interesting transformations of both word-order changes and word replacement. Examples (3) and (4) in Table 2 are representative examples from back translation.

**Label Balancing** Figure 1 illustrates the process of constructing the final label-balanced set based on human annotations. The set first includes all pairs from back translation, which are mostly paraphrases. For each labeled pair  $(s_1, s_2)$  from swapping and a labeled pair  $(s_1, s'_1)$  from back translation, the set further includes the pair  $(s_2, s'_1)$  based on the rules: (1)  $(s_2, s'_1)$  is paraphrase if both  $(s_1, s_2)$  and  $(s_1, s'_1)$  are paraphrases; (2)  $(s_2, s'_1)$  is non-paraphrase if exactly one of  $(s_1, s_2)$  and  $(s_1, s'_1)$  is non-paraphrase; (3) otherwise  $(s_2, s'_1)$  is not included because its label is unknown. We also consider pairs  $(s'_2, s_1)$  and  $(s'_2, s'_1)$  in the similar way if  $s'_2$  is a back translation of  $s_2$  with human labels.

## 4 PAWS Dataset

Using the example generation strategies described in Section 3 combined with human paraphrase an-



	Quora	Wikipedia
# Raw pairs	16,280	50,000
<i>Sentence correction</i>		
# Accepted pairs	10,699	39,903
# Fixed pairs	3,626	7,387
# Rejected pairs	1,955	2,710
<i>Paraphrase identification</i>		
Total # pairs	14,325	47,290
paraphrase	4,693	5,725
non-paraphrase	9,632	41,565
Human agreement	92.0%	94.7%
<i>After post-filtering</i>		
Total # pairs	12,665	43,647
Human agreement	95.8%	97.5%

Table 3: Detailed counts for examples created via the swapping strategy, followed by human filtering and paraphrase judgments.

notations, we create a large new dataset, **PAWS** that contains both paraphrase and non-paraphrase pairs that have both high bag-of-words overlap and word reordering. Source sentences are drawn from both the Quora Question Pairs (QQP) corpus (Iyer et al., 2017) and Wikipedia.<sup>3</sup> From these, we produce two datasets, PAWS<sub>QQP</sub> and PAWS<sub>Wiki</sub>.

We start by producing swapped examples from both QQP and Wikipedia. Both sources contain naturally occurring sentences covering many topics. On both corpora only about 3% of candidates are selected for further processing—the rest are filtered because there is no valid generation candidate that satisfies all swapping constraints or because the language model score of the best candidate is below the threshold. The remaining pairs (16,280 for QQP and 50k for Wikipedia) are passed to human review.

**Sentence correction** The examples generated using both of our strategies are generally of high quality, but they still need to be checked with respect to grammar and coherence. Annotators evaluate each generated sentence without seeing its source sentence. The sentence is accepted as is, fixed, or rejected. Table 3 shows the number of pairs of each action on each domain. Most of fixes are minor grammar corrections like *a apple*→*an apple*. Accepted and fixed sentences are then passed to the next stage for paraphrase annotation.

<sup>3</sup><https://dumps.wikimedia.org>

Total # back translation pairs	26,897
paraphrase	25,521
non-paraphrase	1,376
Human agreement	94.8%

Table 4: Paraphrase judgments on example pairs generated by back translation on Wikipedia sentences.

Overall 88% of generated examples passed the human correction phase on both domains.

**Paraphrase identification** Sentence pairs are presented to five annotators, each of which gives a binary judgment as to whether they are paraphrases or not. We choose binary judgments to make our dataset have the same label schema as the QQP corpus. Table 3 shows aggregated annotation statistics on both domains, including the number of paraphrase (positive) and non-paraphrase (negative) pairs and human agreement, which is the percentage ratio of agreement between each individual label and the majority vote of five labels on each example pair. Overall, human agreement is high on both Quora (92.0%) and Wikipedia (94.7%) and each label only takes about 24 seconds. As such, answers are usually straightforward to human raters.

To ensure the data is comprised of clearly paraphrase or non-paraphrase pairs, only examples with four or five raters agreeing are kept.<sup>4</sup> An example of low agreement is *Why is the 20th-century music so different from the 21st music?* v.s. *Why is the 21st century music so different from the 20th century music?*, where three out of five raters gave negative labels on this pair. The bottom block of Table 3 shows the final number of pairs after this filtering, and human agreement further goes up to over 95%. Finally, source and generated sentences are randomly flipped to mask their provenance.

The swapping strategy generally produces non-paraphrase examples—67% for QQP and 88% for Wikipedia. Because (a) the label imbalance is less pronounced for QQP and (b) NMT models perform poorly on Quora questions due to domain mismatch, we only apply the back translation strategy to Wikipedia pairs. Doing so creates 26,897 candidate example pairs after filtering. As before, each pair is rated by five annotators on the paraphrase identification task.<sup>5</sup> Table 4 shows that

<sup>4</sup>We exclude low agreement pairs from our experiments, but we include them in our data release for further study.

<sup>5</sup>Sentence correction was not necessary for these because

	Train	Dev	Test	Yes%
PAWS <sub>QQP</sub>	11,988	677	–	31.3%
PAWS <sub>Wiki</sub>	49,401	8,000	8,000	44.2%
PAWS <sub>Wiki-Swap</sub>	30,397	–	–	9.6%

Table 5: Counts of experimental split for each PAWS dataset. The final column gives the proportion of paraphrase (positive) pairs. There are 108,463 PAWS pairs in total.

most of the examples (94.9%) are paraphrases (as expected), with high human agreement (94.8%). Finally, we expand the pairs using the the rules described in Section 3.2.

Table 5 provides counts for each split in the final PAWS datasets. The training portion of PAWS<sub>QQP</sub> is a subset of the QQP training set; however, PAWS<sub>QQP</sub>’s development set is a subset of both QQP’s development and test sets because there are only 677 pairs. PAWS<sub>Wiki</sub> randomly draws 8,000 pairs for each of its development and test sets and takes the rest as its training set, with no overlap of source sentences across sets. Finally, any trivial pairs with identical sentences from development and test sets are removed.<sup>6</sup> The final PAWS<sub>QQP</sub> has a total of 12,665 pairs (443k tokens), where 31.3% of them have positive labels (paraphrases). PAWS<sub>Wiki</sub> has a total of 65,401 pairs (2.8m tokens), where 44.2% of them are paraphrases.

Note that we have human annotations on 43k pairs generated by the word swapping method on Wikipedia, but 30k of them have no back translation counterparts and therefore they are not included in our final PAWS<sub>Wiki</sub> dataset. Nevertheless, they are high-quality pairs with manual labels, so we include them as an auxiliary training set (PAWS<sub>Wiki-Swap</sub> in Table 5), and empirically show its impact in Section 6.

**Unlabeled PAWS<sub>Wiki</sub>** In addition to the fully labeled PAWS<sub>Wiki</sub> dataset, we also construct an unlabeled PAWS<sub>Wiki</sub> set at large scale. The idea is to simply treat all pairs from word swapping as non-paraphrases and all pairs from back translation as paraphrase, and construct the dataset in the same way as labeled PAWS<sub>Wiki</sub>. The result is a total of 656k pairs with silver labels. We show empirically

NMT generates fluent output.

<sup>6</sup>Such trivial examples exist because annotators sometimes fix a swapped sentence back to its source. We keep such examples in the training set (about 8% of the corpus) because otherwise a trained model would actually predict low similarity scores to identical pairs.

	BOW	BiLSTM & ESIM	DecAtt	DIIN & BERT
Non-local context	×	✓	×	✓
Word interaction	×	×	✓	✓

Table 6: Complexity of each evaluated model.

the impact of using this silver set in pre-training in Section 6.

## 5 Evaluated Models

PAWS is designed to probe models’ ability to go beyond recognizing overall sentence similarity or relatedness. As noted in the introduction, models—even the best available—trained on existing resources tend to classify any example with high BOW overlap as a paraphrase. Can any of these models learn finer structural sensitivity when provided with PAWS examples as part of their training?

We consider six different models that cover a wide range of complexity and expressiveness: two baseline encoders and four recent advanced models that achieved state-of-the-art or strong performance on paraphrase identification. Table 6 summarizes the models with respect to whether they represent non-local contexts or support cross-sentential word interaction.

The baseline models use cosine similarity with simple sentence encoders: a bag-of-words (**BOW**) encoder based on token unigram and bigram encodings and a bi-directional LSTM (**BiLSTM**) that produces a contextualized sentence encoding. A cosine value above .5 is taken as a paraphrase.

**ESIM.** The Enhanced Sequential Inference Model (Chen et al., 2017) achieved competitive performance on eight sentence pair modeling tasks (Lan and Xu, 2018). It encodes each sentence using a BiLSTM, concatenates the encodings for each sentence in the pair, and passes them through a multi-layer perceptron (MLP) for classification. The additional layers allow ESIM to capture more complex sentence interaction than cosine similarity in the baseline models.

**DecAtt.** The Decomposable Attention Model (Parikh et al., 2016) is one of the earliest models to introduce attention for paraphrase identification. It computes word pair interaction between two sentences and aggregates aligned vectors for final classification. This model achieved state-of-the-art results without explicitly modeling word

order. In our experiments, we show the limitations of this modeling choice on PAWS pairs.

**DIIN.** The Densely Interactive Inference Network (Gong et al., 2018) adopts DenseNet (Huang et al., 2017), a 2-dimensional convolution architecture, to extract high-order word-by-word interaction between n-gram pairs. This model achieved state-of-the-art performance without relying on pre-trained deep contextualized representations like ELMo (Peters et al., 2018). It outperformed ESIM and DecAtt models by a large margin on both paraphrase identification and natural language inference tasks.

**BERT.** The Bidirectional Encoder Representations from Transformers (Devlin et al., 2018) recently obtained new state-of-the-art results on eleven natural language processing tasks, including pushing the GLUE benchmark to 80.4% (7.6% absolute improvement). BERT involves pre-training a Transformer encoder (Vaswani et al., 2017) on a large corpus with over three billion words. This large network is then fine-tuned with just one additional output layer.

## 6 Experiments

We seek to understand how well models trained on standard datasets perform on PAWS pairs and to see which models are most able to learn from PAWS pairs. A strong model should improve significantly on PAWS when trained on PAWS pairs without diminishing performance on existing datasets like QQP. Overall, both DIIN and BERT prove remarkably able to adapt to PAWS pairs and perform well on both PAWS<sub>QQP</sub> and PAWS<sub>Wiki</sub> while the other models prove far less capable.

### 6.1 Experimental Setup

We use two metrics: classification accuracy and area-under-curve (AUC) scores of precision-recall curves. For all classification models, 0.5 is the threshold used to compute accuracy. We report results on testing sets for QQP and PAWS<sub>Wiki</sub>, and on the development set for PAWS<sub>QQP</sub> (which has no test set).

For BERT, we use the implementation provided by the authors<sup>7</sup> and apply their default fine-tuning configuration. We use the provided BERT<sub>BASE</sub> pre-trained model instead of BERT<sub>LARGE</sub> due to GPU memory limitations. For all other models, we use our own (re-)implementations that matched

reported performance on QQP. We use 300 dimensional GloVe embeddings (Pennington et al., 2014) to represent words and fix them during training.

### 6.2 Results

**Main Results on PAWS<sub>QQP</sub>** Table 7 summarizes results on the Quora domain. We first train models on the Quora Question Pairs (QQP) training set, and column “QQP→QQP” shows that all models achieve over 83% accuracy on QQP. However, when evaluating on PAWS<sub>QQP</sub>, all models, including BERT, obtain abysmal accuracy under 40% (column “QQP→PAWS<sub>QQP</sub>”).

We hypothesize the performance on PAWS<sub>QQP</sub> relies on two factors: the number of representative training examples, and the capability of models to represent complex interactions between words in each sentence and across the sentences in the pair. To verify that, we further train models on a combination of QQP and PAWS<sub>QQP</sub> training sets and the last two columns of Table 7 show the results on PAWS<sub>QQP</sub>. As expected, all models benefit from new training examples, but to different extents. Gains are much larger on state-of-the-art models like BERT, while the BOW model learns almost nothing from new examples. As a consequence, performance changes are more drastic on PAWS<sub>QQP</sub> than on QQP. For example, the absolute difference between BiLSTM and BERT is 4.2% on QQP, but it goes up to 27% on PAWS<sub>QQP</sub>, which is a 60% relative reduction in error.

It is also noteworthy that adding PAWS<sub>QQP</sub> training examples has no negative impact to QQP performance at all. For example, a BERT model fine-tuned on QQP+PAWS<sub>QQP</sub> achieves the same 90.5% classification accuracy as training on QQP alone. We therefore obtain a single model that performs well on both datasets.

**Main Results on PAWS<sub>Wiki</sub>** In our second experiment we train and evaluate models on our PAWS<sub>Wiki</sub> dataset. Table 8 presents the results. DIIN and BERT outperform others by a substantial margin (>17% accuracy gains). This observation gives more evidence that PAWS data effectively measures models’ sensitivity to word order and syntactic structure.

One interesting observation is that DecAtt performs as poorly as BOW on this dataset. This is likely due to the fact that DecAtt and BOW both consider only local context information. We there-

<sup>7</sup><https://github.com/google-research/bert>

MODELS	QQP→QQP		QQP→PAWS <sub>QQP</sub>		QQP+PAWS <sub>QQP</sub> →PAWS <sub>QQP</sub>	
	(Acc)	(AUC)	(Acc)	(AUC)	(Acc)	(AUC)
BOW	83.2	89.5	29.0	27.1	30.0 (+1.0)	27.3 (+0.2)
BiLSTM	86.3	91.6	34.8	<b>37.9</b>	57.6 (+22.9)	52.3 (+14.5)
ESIM (Chen et al., 2017)	85.3	92.8	<b>38.9</b>	26.9	66.5 (+27.7)	48.1 (+17.2)
DecAtt (Parikh et al., 2016)	87.8	93.9	33.3	26.3	67.4 (+34.1)	51.1 (+24.9)
DIIN (Gong et al., 2018)	89.2	95.2	32.8	32.4	83.8 (+51.1)	77.8 (+45.5)
BERT (Devlin et al., 2018)	<b>90.5</b>	<b>96.3</b>	33.5	35.1	<b>85.0</b> (+51.5)	<b>83.1</b> (+48.0)

Table 7: Accuracy (%) of classification and AUC scores (%) of precision-recall curves on Quora Question Pairs (QQP) testing set and our PAWS<sub>QQP</sub> development set. QQP→PAWS<sub>QQP</sub> indicates that models are trained on QQP and evaluated on PAWS<sub>QQP</sub>. Other columns are defined in a similar way. QQP+PAWS<sub>QQP</sub> is a simple concatenation of the two training sets. Boldface numbers indicate the best accuracy for each testing scenario. Numbers in parentheses indicate absolute gains from adding PAWS<sub>QQP</sub> training data.

MODELS	Supervised		Pretrain+Fine-tune	
	(Acc)	(AUC)	(Acc)	(AUC)
BOW	55.8	41.1	55.6	44.9
BiLSTM	71.1	75.6	80.8	87.6
ESIM	67.2	69.6	81.9	85.8
DecAtt	57.1	52.6	55.8	45.4
+BiLSTM	68.6	70.6	88.8	92.3
DIIN	88.6	91.1	91.8	<b>94.4</b>
BERT	<b>90.4</b>	<b>93.7</b>	<b>91.9</b>	94.3

Table 8: Accuracy (%) and AUC scores (%) of different models on PAWS<sub>wiki</sub> testing set. **Supervised** models are trained on human-labeled data only, while **Pretrain+Fine-tune** models are first trained on noisy unlabeled PAWS<sub>wiki</sub> data and then fine-tuned on human-labeled data.

fore tested an enhancement of DecAtt by replacing its word representations with encodings from a BiLSTM encoder to capture non-local context information. The enhanced model significantly outperforms the base, yielding an 11.5% (57.1% vs. 68.6%) absolute gain on accuracy.

We further evaluate the impact of using silver PAWS<sub>wiki</sub> data in pre-training, as discussed in Section 4. The last two columns of Table 8 show the results. Comparing to supervised performance, pre-training with silver data gives consistent improvements across all models except BOW and vanilla DecAtt. Perhaps surprisingly, adding silver data gives more than 10% absolute improvements on AUC scores for BiLSTM and ESIM, much higher than the gains on DIIN and BERT.

**Size of Training Set** To analyze how many PAWS examples are sufficient for training, we

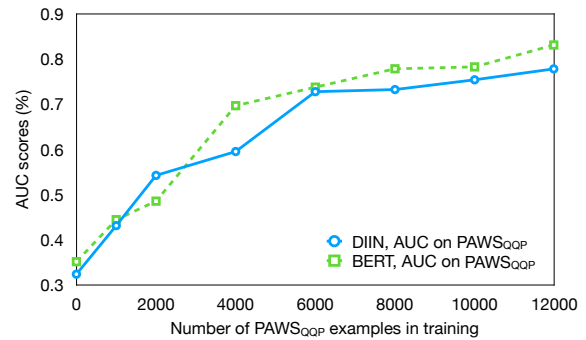


Figure 4: AUC scores ( $y$ -axis) as a function of the number of PAWS<sub>QQP</sub> examples in the training set ( $x$ -axis).

train multiple models on QQP plus different number of PAWS<sub>QQP</sub> examples. Figure 4 plots AUC score curves of DIIN and BERT as a function of the number of PAWS<sub>QQP</sub> training examples.  $x = 0$  corresponds to models trained on QQP only, and the rightmost points correspond to models trained on QQP and full PAWS<sub>QQP</sub>. Both models improve from 30% to 74% AUC scores with 6,000 PAWS<sub>QQP</sub> examples. Furthermore, neither curve reaches convergence, so they would likely still benefit from more PAWS training examples.

**Cross-domain Results** The PAWS datasets cover two domains: Quora and Wikipedia. Here we demonstrate that a model trained on one domain also generalizes to another domain, although not as well as training on in-domain data. Table 9 shows that a DIIN model trained on Quora (QQP+PAWS<sub>QQP</sub>) achieves 70.5% AUC on the Wikipedia domain. This is lower than training on in-domain data (92.9%), but higher than the model trained without any PAWS data (46.0%). We also observe similar patterns when



TRAINING DATA	QQP (Test)	PAWS <sub>QQP</sub> (Dev)	PAWS <sub>Wiki</sub> (Test)
QQP (Train)	95.2	32.4	46.0
QQP+PAWS <sub>QQP</sub>	<b>95.3</b>	77.8	70.5
QQP+PAWS <sub>Wiki</sub>	<b>95.3</b>	58.5	92.9
+PAWS <sub>Wiki-Swap</sub>	<b>95.3</b>	70.6	93.5
QQP+PAWS <sub>QQP+Wiki</sub>	95.1	87.0	93.4
+PAWS <sub>Wiki-Swap</sub>	<b>95.3</b>	<b>89.9</b>	<b>93.8</b>

Table 9: AUC scores (%) when training DIIN models on different sets of training data. Boldface numbers indicate the best accuracy for each testing set.

training on Wikipedia (QQP+PAWS<sub>Wiki</sub>) and testing on PAWS<sub>QQP</sub>. Interestingly, using out-of-domain data also boosts in-domain performance. As Table 9 shows, training on both domains (QQP+PAWS<sub>QQP+Wiki</sub>) leads to 9.2% absolute AUC gains on PAWS<sub>QQP</sub> over the model trained only on QQP+PAWS<sub>QQP</sub>.

The auxiliary training set on Wikipedia (PAWS<sub>Wiki-Swap</sub>) helps further. As Table 9 shows, adding this auxiliary training set is particularly helpful to the performance on PAWS<sub>QQP</sub>, yielding a 12.1% (70.6% vs 58.5%) gain on AUC when training on QQP+PAWS<sub>Wiki</sub>. On PAWS<sub>Wiki</sub>, this addition lifts the (no pre-training) DIIN model AUC from 91.1% (Table 8) to 93.8% (Table 9).

**BERT vs DIIN** Both models achieve top scores on PAWS, but interestingly, the two models disagree on many pairs and are not correlated in their errors. For example, of 687 of BERT’s mistakes on the PAWS<sub>Wiki</sub> test set, DIIN got 280 (41%) correct. As such, performance might improve with combinations of these two existing models.

It is also worth noting that the DIIN model used in our experiments has only 590k model parameters, whereas BERT has over 100m. Furthermore, the computational cost of BERT is notably higher than DIIN. Given this, and the fact that DIIN is competitive with BERT (especially when pre-trained on noisy pairs, see Table 8), DIIN is likely the better choice in computationally constrained scenarios—especially those with strict latency requirements.

## 7 Conclusion

Datasets are insufficient for differentiating models if they lack examples that exhibit the necessary diagnostic phenomena. This has led, for example,

to new datasets for noun-verb ambiguity (Elkahky et al., 2018) and gender bias in coreference (Webster et al., 2018; Rudinger et al., 2018; Jieyu Zhao, 2018). Our new PAWS datasets join these efforts and provide a new resource for training and evaluating paraphrase identifiers. We show that including PAWS training data for state-of-the-art models dramatically improves their performance on challenging examples and makes them more robust to real world examples. We also demonstrate that PAWS effectively measures sensitivity of models on word order and syntactic structure.

## Acknowledgement

We would like to thank our anonymous reviewers and the Google AI Language team, especially Emily Pitler, for the insightful comments that contributed to this paper. Many thanks also to the Data Compute team, especially Ashwin Kakarla and Henry Jicha, for their 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.
- Ciprian Chelba, Tomas Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, Phillipp Koehn, and Tony Robinson. 2014. One billion word benchmark for measuring progress in statistical language modeling. In *INTERSPEECH*, pages 2635–2639. ISCA.
- 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.
- 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, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1657–1668.
- Alexis Conneau, Germán Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. What you can cram into a single  $\$&!#*$  vector: Probing sentence embeddings for linguistic properties. In *Proceedings of the 56th Annual Meeting of*

- the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136. 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, May 7-12, 2018*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ali Elkahky, Kellie Webster, Daniel Andor, and Emily Pitler. 2018. A challenge set and methods for noun-verb ambiguity. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2562–2572.
- Allyson Ettinger, Ahmed Elgohary, Colin Phillips, and Philip Resnik. 2018. Assessing composition in sentence vector representations. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1790–1801.
- Simone Filice, Giovanni Da San Martino, and Alessandro Moschitti. 2015. Structural representations for learning relations between pairs of texts. 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 1003–1013.
- 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. Association for Computational Linguistics.
- Yichen Gong, Heng Luo, and Jian Zhang. 2018. Natural language inference over interaction space. In *International Conference on Learning Representations*.
- Ankush Gupta, Arvind Agarwal, Prawaan Singh, and Piyush Rai. 2018. A deep generative framework for paraphrase generation. In *AAAI*, pages 5149–5156. AAAI Press.
- Kelvin Guu, Tatsunori B. Hashimoto, Yonatan Oren, and Percy Liang. 2018. Generating sentences by editing prototypes. *TACL*, 6:437–450.
- Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. 2017. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 2261–2269.
- Shankar Iyer, Nikhil Dandekar, and Kornél Csernai. 2017. First quora dataset release: Question pairs.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. Adversarial example generation with syntactically controlled paraphrase networks. In *NAACL-HLT*, pages 1875–1885. 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.
- Mark Yatskar Vicente Ordonez Kai-Wei Chang Jieyu Zhao, Tianlu Wang. 2018. Gender bias in coreference resolution: evaluation and debiasing methods. In *NAACL (short)*.
- Wuwei Lan, Siyu Qiu, Hua He, and Wei Xu. 2017. A continuously growing dataset of sentential paraphrases. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 1224–1234.
- Wuwei Lan and Wei Xu. 2018. Neural network models for paraphrase identification, semantic textual similarity, natural language inference, and question answering. In *COLING*, pages 3890–3902. Association for Computational Linguistics.
- Zichao Li, Xin Jiang, Lifeng Shang, and Hang Li. 2018. Paraphrase generation with deep reinforcement learning. In *EMNLP*, pages 3865–3878. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. 2014. Microsoft COCO: Common Objects in Context. In *CoRR*.
- Yang Liu, Matt Gardner, and Mirella Lapata. 2018. Structured alignment networks for matching sentences. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1554–1564.
- Jeff Mitchell and Mirella Lapata. 2008. Vector-based models of semantic composition. In *ACL 2008, Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics, June 15-20, 2008, Columbus, Ohio, USA*, pages 236–244.
- Ankur P. Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2249–2255.

- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 2227–2237.
- Marco Tulio 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.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. Gender bias in coreference resolution. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 8–14, New Orleans, Louisiana. Association for Computational Linguistics.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R. Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R. Bowman, Dipanjan Das, and Ellie Pavlick. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. In *International Conference on Learning Representations*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 6000–6010.
- Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge. 2018. Mind the gap: A balanced corpus of gendered ambiguous pronouns. In *Transactions of the ACL*, page to appear.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.