

PARAPHRASUS: A Comprehensive Benchmark for Evaluating Paraphrase Detection Models

Andrianos Michail
University of Zurich
andrianos.michail@cl.uzh.ch

Simon Clematide
University of Zurich
simon.clematide@cl.uzh.ch

Juri Opitz
University of Zurich
jurialexander.opitz@cl.uzh.ch

Abstract

The task of determining whether two texts are paraphrases has long been a challenge in NLP. However, the prevailing notion of paraphrase is often quite simplistic, offering only a limited view of the vast spectrum of paraphrase phenomena. Indeed, we find that evaluating models in a paraphrase dataset can leave uncertainty about their true semantic understanding. To alleviate this, we create **PARAPHRASUS**, a benchmark designed for multi-dimensional assessment, benchmarking and selection of paraphrase detection models. We find that paraphrase detection models under our fine-grained evaluation lens exhibit trade-offs that cannot be captured through a single classification dataset. Furthermore, **PARAPHRASUS** allows prompt calibration for different use cases, tailoring LLM models to specific strictness levels. **PARAPHRASUS** includes 3 challenges spanning over 10 datasets, including 8 repurposed and 2 newly annotated; we release it along with a benchmarking library at <https://github.com/impresso/paraphrasus>

1 Introduction

Our study was set in motion by a serendipitous finding. Like many other researchers at the time, we were benchmarking large language models (LLMs). We had *particular interest in the paraphrasing detection task*¹ and evaluated Llama and others LLMs (Dubey et al., 2024; Alves et al., 2024) on the adversarial paraphrase detection test set *PAWS-X* (Yang et al., 2019). Interestingly, LLMs appear to underperform much on *PAWS-X*. Their performance lies slightly above random coin flip, and they are vastly outperformed by smaller BERT-based models fine-tuned on the training split of *PAWS-EN*.

¹We consider paraphrasing as an important and challenging task, hand in hand with Quintilianus (95, p. 117) who lived 2,000 years ago, and of course, with many papers presented at CL and ML conferences such as Bhagat and Hovy (2013); Zhou and Bhat (2021); Krishna et al. (2023).




Figure 1: Percentage of paraphrases predicted on the Semantic Text Similarity Dataset (STSBenchmark) dataset (Cer et al., 2017), binned by scores from 0 (completely dissimilar) to 5 (completely equivalent). Human annotation comes from the **STS-H** human annotation we perform.

Therefore, we investigated whether from analyzing results on PAWS-X, we could make conclusions about the paraphrase detection understanding (or lack thereof) of different models. As what was merely intended as a sanity check at first, we repurposed data from the semantic text similarity (STS) task, and investigated the distribution of predicted paraphrases against the fine-grained similarity labels, where only the highest label of 5 denotes true paraphrases. In fact, moving away from 5, the semantic similarity decreases rapidly, so there clearly are no paraphrases below a certain level (e.g., 4). We confronted different models with this simple test and were surprised by the results (see Figure 1). Suddenly, the predictions of LLMs appeared more reasonable than what we took away from the benchmarking results, whereas the trained models that performed well on PAWS-X appeared naïve and much too over-confident, assigning many paraphrase predictions to sentences that are semantically different. Indeed, “predicting paraphrases is not easy” (Vahtola et al., 2022).

Inspired by these pilot findings, we began a deeper empirical exploration into paraphrase notions and paraphrase detection models, the results of which are presented in the remainder of this paper. A main outcome of this is our **PARAPHRASUS** benchmark that allows the community to test paraphrase detection models in different aspects, building on existing, novel, and carefully repurposed data, eliciting diverse notions of paraphrases. With **PARAPHRASUS**, we can gather a fine-grained picture of paraphrase modes performance that reveals particular strengths and weaknesses. Our contributions are the following:

1. We present **PARAPHRASUS**, a multi-faceted evaluation benchmark for paraphrase detection, including datasets of human-written sentence pairs of varying semantic and lexical similarity repurposed for testing the generalization of paraphrase detection models.
2. Within our benchmark, we also contribute two novel datasets. i) We annotate a test set of 338 semantically similar sentence pairs for paraphrase classification, creating a challenging test set that paraphrase detection methods plateau at 60% accuracy. ii) We exploit guidelines for Abstract Meaning Representation annotations, extracting expert examples of true paraphrases.
3. To showcase the usefulness of **PARAPHRASUS**, we test LLMs and trained models under different setups, revealing new insights about paraphrase detection models and paraphrase data. A main insight is that evaluating a model on a single paraphrase dataset, as commonly done before, can yield a misleading picture about actual generalization performance. The effect is especially marked for trained models.
4. We perform quantitative and qualitative analyses, learning about suitable model training regimes and human perception of paraphrases.

2 Related Work

Paraphrasing. We find early excitement about paraphrases in the rhetorician **Quintilianus (95)**'s book, written 2000 years ago. **Quintilianus** suggests paraphrasing as an exercise for classrooms, praising the task as “valuable in virtue of its difficulty”, the goal is “to rival and vie with the original in the expression of the same thoughts” (p. 117). A

useful taxonomy of different types of paraphrases can be found later in **Bhagat and Hovy (2013)**'s work, outlining 25 means of creating one, e.g., by metaphor substitution or function word variation (“Pat gave a nice demo”, “Pat’s demo was nice”).

Clearly, there are many NLP applications where paraphrasing and understanding paraphrases plays a crucial role (**Mallinson et al., 2017**). Some of these are semantic search (**Reimers and Gurevych, 2019**), style transfer (**Krishna et al., 2020**), machine translation (**Madnani et al., 2007**), plagiarism (**Barrón-Cedeño et al., 2013; Sharjeel et al., 2016**), and text reuse (**Büchler et al., 2012**). The evaluation of text generation is often interested in whether a candidate and the reference are paraphrases (e.g., **Freitag et al., 2020; Thompson and Post, 2020; Nawrath et al., 2024**).

Paraphrase Datasets Recognizing the importance of paraphrasing, prior research has developed methods for constructing datasets, including mining large corpora for multilingual paraphrases (**Ganitkevitch et al., 2013; Ganitkevitch and Callison-Burch, 2014**) and using back-translation to generate 50 million paraphrase pairs (**Wieting and Gimpel, 2018**). Such large-scale datasets are especially attractive for training models for each of the diverse paraphrase-related applications mentioned above. On the other hand, smaller datasets have been proposed, with a stronger focus on quality and evaluation of paraphrase judgment models. One of the earliest of such datasets is the **MRPC** corpus (**Dolan and Brockett, 2005**), this paraphrase classification dataset was constructed by mining the web and filtering the most promising candidates using lexicon-based classifiers. Two non-specialist human judges then labelled the sentence pairs, resulting in a test set of 1,730 sentence pairs, of which 66.5% were labelled positive.

Recently, the **PAWS** dataset (**Zhang et al., 2019**) was introduced, providing an adversarial dataset for paraphrase classification through the application of word scrambling. They create challenging pairs by controlled word swapping and back-translation, followed by human quality check for fluency and paraphrase semantic. The dataset is split in 49,401 training sentence pairs and 8,000 test sentence pairs. **PAWSX** (**Yang et al., 2019**) takes 2000 of the **PAWS** Wiki test samples and creates a 7-way multilingual test set in German, Spanish, French, Chinese, Japanese, and Korean using human translators. While these minimal pairs are

interesting to perform a focused evaluation, its adversarial mode of dataset creation can introduce limitations on the coverage of linguistic formulation variation and hinder generalization abilities.

Typically, there is a quantity-quality tradeoff. With **PARAPHRASUS** we aim to get the best of the worlds of quantity and quality, releasing a comprehensive and trustworthy multifold dataset for evaluating paraphrase detection models and analyzing various sorts of paraphrases.

Creating benchmarks. Benchmarking is a very active topic in NLP (Gardent et al., 2017; Zhu et al., 2018), and a key tool for study and development of LLMs (Srivastava et al., 2023). These benchmarks are useful since they help to assess and select models within a broader picture, and thus encourage their further development (Gehrmann et al., 2021). A recent critique that has emerged is due to the way how LLMs are trained, as they are often trained on large corpora crawled from the web. Hence, benchmark data may have leaked into the model, and the performance is harder to assess (Deng et al., 2024). A way to at least partially mitigate this concern is by repurposing data from related tasks. Specifically, while Honovich et al. (2022) repurpose data for their factuality benchmark from existing factuality datasets from different domains, we take a step further and source data not only from different domains, but also from related tasks, projecting them as paraphrase classification. This further increases trust in evaluation results based on our proposed **PARAPHRASUS** benchmark.

Semantic similarity datasets. STS (Agirre et al., 2013; Cer et al., 2017) and SICK (Marelli et al., 2014) elicited human ratings of sentence similarity on a Likert scale. While STS annotates *semantic* similarity, SICK annotates *semantic* relatedness. These two aspects are highly related, but not exactly the same (Budanitsky and Hirst, 2006; Kolb, 2009). However, the highest scores on the Likert scales of SICK and STS denote the equivalence of meaning of two sentences. We exploit this property for our **PARAPHRASUS** paraphrase benchmark by carefully selecting lower similarity pairs as non-paraphrases for controlled tests.

3 Proposed Benchmark

Paraphrases come in different flavors, and distinguishing types of paraphrases can be fuzzy (Bhagat and Hovy, 2013). Our benchmark therefore

Dataset	Pairs	Paraphrase		¬Paraphrase	
		abs.	rel.	abs.	rel.
Classify!					
PAWSX	14,000	6,160	44%	7,840	56%
MRPC	1,730	1,159	67%	571	33%
STS-H	338	109	32%	229	68%
Group	16,068	7,427	48%	8,641	52%
Minimize!					
SNLI	6,632	0	0%	6,632	100%
ANLI	798	0	0%	798	100%
XNLI	16,700	0	0%	16,700	100%
STS	706	0	0%	706	100%
SICK	2,305	0	0%	2,305	100%
Group	27,141	0	0%	27,141	100%
Maximize!					
TRUE	167	167	100%	0	0%
SIMP	600	600	100%	0	0%
Group	767	767	100%	0	0%
Total	43,976	8,194	34%	35,782	66%

Table 1: Overview of the datasets in the **PARAPHRASUS** benchmark along with their label distribution.

adopts an empirical approach to evaluating paraphrase detection models, according to three desiderata: i) It should reflect a broad spectrum of domains, to lower the impact of any superficial training that may have been acquired by a model. ii) It should reflect a broad spectrum of different flavors of paraphrases, e.g., varying strictness levels of paraphrases, negative examples with contradiction or neutral semantic relation. iii) It should help us select models for different paraphrase challenges, or select a single model that performs consistently across multiple categories. Table 1 gives an overview of the size and label distribution of each of the ten evaluation sets. More statistics are available in Appendix C.

3.1 The data: Ten Parts with Three Objectives

Our benchmark consists of 10 parts, with eight of them repurposed from various types of NLP tasks, such as natural language inference (NLI) and meaning representation annotation guidelines. By distributing the data across different domains, we aim to capture a wide range of paraphrase phenomena. An additional benefit of repurposing existing data could be that it helps alleviate concerns that the LLMs being evaluated may have incorporated

this data during training. In addition to the repurposed data, three datasets are genuine paraphrase datasets, one of which is a human-annotated set of particularly challenging examples created by us.

The benchmark is divided into three challenges, along which we structure its description:

Classify! We evaluate the models on three binary paraphrase classification datasets, where each model must determine whether a given pair of sentences constitutes a paraphrase or not.

- **PAWSX:** As mentioned in Section 2, PAWSX is a multilingual dataset containing lexically similar pairs of paraphrases and non-paraphrases.
- **MRPC:** As also mentioned in Section 2, MRPC is one of the first paraphrase detection dataset that has been released (Dolan and Brockett, 2005).
- **STS-H:** We propose using challenging sentence pairs without relying on adversarial strategies. To achieve this, we select highly similar sentence pairs from the STS dataset, which are annotated on a 5-point Likert scale. Potential paraphrase candidates are situated at the upper end of this scale. We select all pairs from 4-5 and carefully create a high-quality set of 338 paraphrase annotations.

A semantics expert and a student annotated this subset independently, resulting in moderate-to-good Kappa of 0.63 (Cohen, 1960). Finally, disagreements were discussed and adjudicated. For the few cases where disagreements could not be resolved (13 out of 56 total), we assigned the non-paraphrase class. This decision was based on the reasoning that it is easier to argue that two sentences are not paraphrases than to prove the opposite.

Minimize positive predictions! In a dataset that contains no paraphrases, the primary objective of the model is to minimize the prediction of paraphrases. However, to ensure the task remains challenging, the paired texts should exhibit some degree of similarity. In this study, we repurpose the following five datasets.

- **SNLI, ANLI, XNLI**— Natural Language Inference (NLI) repurposed: NLI asks whether a hypothesis follows a given premise, labeling the relationship with *entailed*, *neutral*, or *contradiction*. We select those pairs that either

stand in a neutral relationship, or in a contradiction. Since typically the premise is longer than the hypothesis, we control this bias by adding the flipped pairs. The resulting data are denoted by **SNLI**, sampled from the first large-scale NLI dataset (Bowman et al., 2015); **ANLI** taken from an adversarial version of the task (Nie et al., 2020); and the cross-lingual **XNLI** (Conneau et al., 2018)² that allows us to study paraphrase detection models in the minimization task in different languages.

- **STS, SICK**— Similarity repurposed: Orthogonally to how we created the **STS-H** data by using only extremely similar pairs, we now select only pairs that we know that they are not paraphrases. To ensure maximum data quality and lower the possibility of annotation confusions in the upper spectrum of the similarity Likert scale, we select only values 0-3 as negative pairs, excluding the range of 3 to 5. The resulting data are gathered from two datasets (Marelli et al., 2014; Cer et al., 2017).

Maximize positive predictions! If we know that a dataset contains only pairs of paraphrases, the natural expectation towards a model would be to maximize its paraphrase detection rate. We selected two datasets to test this behavior:

- **TRUE:** We create a dataset of simple, guaranteed paraphrases. The objective for a paraphrase model using this dataset is to correctly identify as many paraphrases as possible. To ensure the accuracy of these examples, we leverage Abstract Meaning Representation annotation guidelines (Banarescu et al., 2013).³ AMR has the goal of mapping the same meaning to the same structure, and highlights this in its annotation guidelines by presenting graphs together with a group of semantically equivalent (but structurally different) sentences. For every meaning graph, we extract the n -sized set of example sentences and create $\frac{n^2-n}{2}$ true paraphrase pairs.⁴ As the guidelines succinctly treat diverse phenomena, there can be short phrases, like “20 km” and “20 kilometers”, but most of them are short sentences,

²Limited to EN, DE, FR, ES, ZH due to compute costs.

³<https://github.com/amrisi/amr-guidelines>

⁴Since AMR currently captures many, but not all semantic aspects (Sadeddine et al., 2024), we exclude any parts of the annotation guidelines where AMR conflates two aspectually divergent sentences into a single representation.

e.g., “The boy desires the girl to believe him.” and “The boy has a desire to be believed by the girl.”⁵

- **SIMP**: From the AMR guidelines, we retrieved simple and accurate paraphrases; however, the manually constructed nature of the data limits its diversity. To create a more diverse dataset of paraphrases, albeit at the cost of introducing more noise, we repurpose parallel sentence pairs of different readability levels (Advanced, Intermediate, Elementary) from Vajjala and Lučić (2018)’s OneStopEnglish.

3.2 Evaluation Metric

To achieve a final ranking of systems, we need to select a metric. While there are many choices for a metric, and each has their benefits and drawbacks (Opitz, 2024), we wish for an interpretable and robust measure. For each dataset, we calculate a classification *error* percentage. For datasets from the classify challenge, this is the ratio of wrong predictions. For datasets where the goal is to minimize the amount of positive predictions, like the neutral and contradictory pairs of the NLI data, this is the ratio of positive predictions (in that case, equivalent to false positive rate). For the datasets in the maximization challenge, such as **TRUE**, this is the ratio of negative predictions (equivalent to false negative rate). To achieve a robust final score, we calculate an unweighted average per each of the three challenges, and then another unweighted average over these averages:

$$\overline{Err} = \frac{\sum_{\text{challenge} \in \text{benchmark}} \left(\frac{\sum_{\text{dataset} \in \text{challenge}} \text{error}(\text{dataset})}{|\text{challenge}|} \right)}{|\text{benchmark}|}$$

The \overline{Err} models the probability of observing a wrong prediction, given we randomly select one of the three challenges, and randomly pick an example from a random dataset of the task. Since all scores on **PARAPHRASUS** are error measurements (or their averages), lower numbers are better.

4 Testing Paraphrase Detection Models

In this section, we introduce the methods and models used for systematically testing our **PARAPHRASUS** benchmark to assess the paraphrase detection capabilities. We focus on two research questions:

⁵The license of the AMR guidelines allows us to publicly release this subset of true paraphrases.

Zero Shot Paraphrase Detection Prompt

Are the following sentences {paraphrase notion}?

Sentence 1: {sentence1}

Sentence 2: {sentence2}

Answer with 'Yes' or 'No'

Figure 2: For **P1**, **P2**, and **P3**, the paraphrase notions we ask for are “paraphrases”, “semantically equivalent” and “expressing the same content” respectively. For the ICL expanded prompt, see Appendix A.

1. Given that **PAWSX** is a widely used dataset for measuring the Natural Language Understanding of LLMs and is relatively large, we ask: What insights can be gained from training on it? Specifically, we aim to determine how useful **PAWSX** is for training smaller, more efficient models.
2. Given the increasing use of LLMs, we aim to test a reference open-source LLM to examine its ability to recognize different types of paraphrasing in **PARAPHRASUS**. We also explore different ways of describing the concept of paraphrasing to assess the best methods for prompting LLMs to classify paraphrases.

4.1 What can we learn from PAWSX?

We replicate the approach from the original **PAWSX** study by Yang et al. (2019) and fine-tune a multilingual encoder, XLM-RoBERTa *base* (**XLM-R**) (Conneau et al., 2020) on the PAWS Wiki English training set. Using default fine-tuning hyperparameters, we train for 6 epochs. During initial evaluation, we observed considerable variance between different fine-tuning seeds on **PARAPHRASUS**. To reduce the impact of this variance on our results, we report the average performance based on three high-performing checkpoints, each corresponding to a different seed, yielding a total of nine predictions.

4.2 What can we learn from LLMs?

Recent research has indicated that LLMs have developed a strong understanding of human language, which can be accessed through natural language prompts. In this study, we prompt Llama3 Instruct 8B (Dubey et al., 2024) quantized to 4 bits with

Classification Method	Classify!			Minimize!				Maximize!			Averages			
	PAWSX	STS-H	MRPC	SNLI	ANLI	XNLI	STS	SICK	TRUE	SIMP	Clfy!	Min!	Max!	\overline{Err}
XLM-R ← PAWS-EN Train	15.2	54.1	33.4	32.4	7.2	26.7	46.6	37.0	31.4	5.3	34.2	30.0	18.3	27.5
Llama3 Instruct P1	44.7	56.2	23.6	7.3	13.0	12.3	12.9	0.9	9.0	14.7	41.5	9.3	11.8	*20.9
Llama3 Instruct P2	40.7	37.6	45.9	1.0	1.2	1.4	2.4	0.1	34.7	47.3	41.4	1.2	41.0	27.9
Llama3 Instruct P3	38.1	41.7	37.5	1.3	1.7	1.3	3.5	0.0	35.3	37.5	39.1	1.6	36.4	25.7
Llama3 Instruct P1-ICL_K4	39.0	44.7	33.2	1.9	2.0	2.8	3.5	0.3	29.9	33.3	39.0	2.1	31.6	24.2
Llama3 Instruct P2-ICL_K4	34.1	41.7	45.2	0.8	0.8	0.3	3.1	0.0	40.1	42.3	40.3	1.0	41.2	27.5
Llama3 Instruct P3-ICL_K4	33.2	39.1	46.7	0.5	0.8	0.3	2.4	0.0	50.9	45.5	39.7	0.8	48.2	29.6

Table 2: Main results on PARAPHRASUS. For description of prompts P1, P2, P3, and their ICL variants, see §4.2. All results are error percentages: Lower numbers are better. The classification method with the best overall performance is marked with an asterisk (*).

three different paraphrase notions along with their in-context learning equivalents.

Expressing different paraphrase notions.

Given that multiple “personalities” can be elicited from LLMs (Chan et al., 2024), we leverage this property to design three distinct, minimalistic prompts that vary only in how the concept of “paraphrase” expressed. The first prompt (henceforth denoted as P1) is most straightforward, just asking the model to judge whether two sentences are *paraphrases*. The second prompt P2 intends to emulate a person with background in semantics, asking the model whether the following two sentences are *semantically equivalent*. We speculate that this prompt triggers a model mode that very strictly judges the pairs, refraining from assigning a positive label if there is only a minor semantic difference. The third prompt P3 aims to kindle a person who is primed from Information Retrieval: The model is asked whether the following two sentences are about *the same content*. We speculate that this prompt provides a suitable balance between the straightforward P1 prompt and the presumed strict prompting P2.

In-context learning with PAWS-X examples.

Our second version of these prompt personas leverages the “emergent capability” of in-context learning (Brown et al., 2020), and is henceforth denoted by the ICL_K4 suffix. Compared to the original, during inference time, we sample two positive and two negative sentence pairs from the training set of PAWS-Wiki and prepend them to the prompt.

5 Main Benchmark Results

The main results on PARAPHRASUS are shown in Table 2.

What model is best overall? Interestingly, our most simple LLM setup (P1) demonstrates the best overall result (21% errors on average) on PARAPHRASUS, outperforming the second best model that uses in-context examples from PAWSX (P1-ICL_K4) by 3.2 percentage points (pp). While the smaller fine-tuned XLM-R model excels on PAWSX with an error rate of only 15.2% it struggles to generalize across the other objectives and ranks second to last, showing an overall error of 27.5%.

What model is best for specific objectives?

Findings on PARAPHRASUS show that each model exhibits unique strengths. When examining the averages for each objective (right side of Table 2), Llama with prompt P3-ICL_K4 (Min!, 0.8% average error) outperforms others in the minimization objective, while Llama3 with prompt P1 (Max!, 11.8% average error) excels in the maximization objective. In the classification objective, the trained XLM-R model achieves the best performance. However, it is important to note that this particular average is skewed due to the model’s training on PAWSX. In the other two classification datasets, XLM-R is outperformed by Llama3 prompts, with up to a 14.8 pp error difference on STS-H (using prompt P2) and a 9.8 pp difference on MRPC (using prompt P1).

Zero Shot or In Context Learning

In the zero-shot experiments, the different paraphrase notions in the prompt lead to different levels of strictness and overall result, whereas in the ICL experiments the three prompts behave stricter. Adding in-context examples to P1 seems to calibrate the model to reduce false positives, but at the cost of a 19.8 pp error difference in the maximization average, and a lower overall result.

Results are error percentages ↓	Classify!			Minimize!				Maximize!		Averages				
Model ← PAWS-EN % Negatives to EasyNegs.	PAWSX	STS-H	MRPC	SNLI	ANLI	XNLI	STS	SICK	TRUE	SIMP	Clfy!	Min!	Max!	\overline{Err}
XLM-R ← Original Negatives (44% 56%)	15.2	54.1	33.4	32.4	7.2	26.7	46.6	37.0	31.4	5.3	34.2	30.0	18.3	27.5
XLM-R ← 25% Easy Negatives (44% 56%)	15.1	54.8	32.1	32.8	5.9	20.8	42.6	32.7	32.0	4.0	34.0	27.0	18.0	* 26.3
XLM-R ← 50% Easy Negatives (44% 56%)	15.9	58.5	30.8	43.6	7.7	29.8	58.9	50.7	14.5	2.3	35.1	38.1	8.4	27.2
XLM-R ← 75% Easy Negatives (44% 56%)	18.0	63.1	33.2	75.8	38.0	67.4	85.3	78.9	5.3	1.1	38.1	69.1	3.2	36.8

Table 3: Ablation study A: XLM-R fine-tuned on PAWS EN Wiki training samples where a percentage of the second sentences of the negative pairs was exchanged with a random sentence. All results are error percentages: Lower numbers are better. The classification method with the best overall performance is marked with an asterisk (*).

Results are error percentages ↓	Classify!			Minimize!				Maximize!		Averages				
Model ← PAWS EN Train Labeled by Prompt	PAWSX	STS-H	MRPC	SNLI	ANLI	XNLI	STS	SICK	TRUE	SIMP	Clfy!	Min!	Max!	\overline{Err}
XLM-R ← Original Labels (44% 56%)	15.2	54.1	33.4	32.4	7.2	26.7	46.6	37.0	31.4	5.3	34.2	30.0	18.3	* 27.5
XLM-R ← Llama3 P1 (79% 21%)	45.0	63.2	33.4	77.9	93.1	89.9	70.4	63.7	23.2	2.9	47.2	79.0	13.0	46.4
XLM-R ← Llama3 P2 (59% 41%)	34.1	58.8	32.7	60.3	63.7	70.8	66.4	50.6	21.3	6.9	41.9	62.4	14.1	39.5
XLM-R ← Llama3 P3 (56% 44%)	30.0	56.4	32.8	47.5	58.7	62.6	52.6	34.5	29.0	8.5	39.7	51.2	18.8	36.6
XLM-R ← Llama3 P1 \wedge P2 \wedge P3	28.4	54.6	34.0	33.7	49.0	55.7	40.6	19.0	38.1	10.7	39.0	39.6	24.4	34.3

Table 4: Ablation study B: XLM-R fine-tuned on PAWS Wiki training samples but instead using predicted Llama3 labels as targets (for description of prompts P1, P2, P3 see 4.2). $X \wedge Y$: The positive label is assigned if and only if it is predicted by both prompt X and Y. All results are error percentages: Lower numbers are better. The classification method with the best overall performance is marked with an asterisk (*).

6 Discussion

6.1 Can we improve the XLM-R training?

In the main results, we observed that fine-tuning a model on the PAWSX training set yields strong performance on the PAWSX test set. However, despite this success, the trained model ranks second to last overall on PARAPHRASUS. A natural question then is: Why did the model fail to generalize?

Is it the adversarial character? We hypothesize that this may be due to the adversarial character of PAWSX. While its difficulty should ideally teach the model to handle even the most challenging cases, the artificial creation of the dataset may have introduced biases. These biases could have led the model to rely on spurious correlations – similar to the "Clever Hans" effect – where the model learns to exploit patterns that do not generalize well, potentially explaining its struggles with broader generalization (Niven and Kao, 2019).

Ablation study A: Adding easy negatives to the training data. To test this hypothesis, we introduce varying amounts of easy negative samples into the training data. To maintain the original label distribution, we generate these easy negatives by replacing one part of an adversarial negative pair with a random sentence from the training set. The results are presented in Table 3. We observe

that this strategy can indeed improve the generalization ability of the trained model, but only when the amount of easy negative samples are up to 50%. Adding 25% of easy negatives reduces the model’s overall error from 27.5 to 26.3, with improvements on SICK (-4.3 pp) and XNLI (-5.9 pp).

Is it the labels? XLM-R models trained on PAWSX exhibit poor generalization on PARAPHRASUS, which could be attributed to the target labels assigned to the adversarial sentence pairs.

Ablation Study B: Relabeling the Training Data with Llama3. To examine whether the model’s generalization ability could be improved by utilizing labels generated by large language models (LLMs), we trained the model on the PAWS-Wiki training set using silver labels predicted by Llama3 through zero-shot prompting. Contrary to our hypothesis, the results in Table 4 show that models trained on these silver labels exhibit even poorer generalization performance, invalidating the assumption that relabeling enhances generalization. We eyeballed a handful of examples where all prompts agreed on the label, but disagreed with the original one. Some of these examples suggest annotation errors, such as the following: S1: *In 1923, there were 568 wz.1902 guns in the Polish inventory* S2: *In 1923, 568 Polish guns were in the inventory of wz.1902.* In this example, PAWS

assigns a negative label, while all prompts assign a positive label. The consensus seems to suggest that the LLM has realized that ‘wz.1902’ is a type of Polish gun, and all numbers do perfectly match. Though to be perfectly fair, an extremely strict interpretation could argue, e.g., that some of the guns in *S1* might have been (e.g., accidentally) put in another inventory, hence supporting the assigned label, an insight that seems to “post-hoc” support our initial hypothesis in this experiment. But even when changing only those labels to positive where all prompts agreed on, the average error rises by about 6.8 pp. to 34.3% (see Table 4, last row).

Summary of ablation studies What is the underlying cause of the model’s lack of generalization? Based on our experiments, we cannot provide a definitive answer. The first experiment indicates that incorporating simple negative samples and reducing the adversarial nature of the training data may improve generalization. This adjustment also exposes the model to a broader spectrum of semantic similarity between sentence pairs, as opposed to the original dataset, which consisted primarily of difficult positives and negatives. On the other hand, the second experiment seems to indicate that the given PAWS labels are of mostly good quality. While apparently a few annotation mistakes can be corrected by LLM relabeling, any relabeling overall worsens a trained model’s generalization.

We conclude that good performance on PAWSX does not always mean strong generalization in paraphrase detection overall. We have not yet identified a method to train an efficient model on PAWS that effectively generalizes to out-of-distribution scenarios. We speculate that over-fitting to complex adversarial data creation schemes may be at play, counteracting any potentially valuable learning signal (for now).

6.2 Human Paraphrase Understanding Study

When creating STS-H, we obtained independent annotations from two annotators before agreeing on a final gold standard. One annotator was a researcher with a background in semantics, and the other was a student annotator. This setup allows for an interesting analysis of different agreement patterns: between the semanticist and the student, between the LLM and the trained model or one of the human annotators, and other comparisons.

We calculate pairwise Cohen’s Kappa inter-annotator agreement scores (Cohen, 1960) and



Figure 3: Cohens κ between humans and systems when annotating the STS-H dataset that consists of highly similar (STS Score 4-5) sentences.

show the results in Figure 3. We see that there is moderate to high agreement within humans and each model type. Then there is low agreement of Llama with human, and basically zero agreement of XLM with both human and Llama. Interestingly, when considering the maximum STS label of 5 as indicative of a paraphrase, as intended by the task, we observe that human annotators exhibit only moderate agreement with this classification. Occasionally, both humans assigned a paraphrase to a pair from STS that has a score lower of 5.0, at other times, both humans agreed to assign a non-paraphrase label to an STS pair labeled with 5.0.

6.3 Failure modes of LLM

When analyzing LLM performance on PARAPHRASUS using datasets containing only certified negatives or certified positives, we observed that while the error rate of the LLMs was low, it was never zero. Table 5 shows three pairs each from the repurposed NLI XNLI dataset (non-paraphrases only), as well as from the AMR-guidelines repurposed data (paraphrases only). In all these cases, all LLM prompts agree – on the wrong decision. Interestingly, sometimes, the LLM seems to be tricked by seemingly simple active/passive variants and assigned a negative prediction, like *X saddened Y / Y was saddened by X* or a slight positional variation of an adverbial particle (*look up X / look X up*). At times, the model mistakenly labeled a contradiction as a paraphrase, as seen in the first pair where one sentence calls something a significant threat, while the other denies it was a threat.

Overall, we should note that LLMs generally show low error rates on the NLI-repurposed data, demonstrating robustness against hard negatives.

False Positive examples: LLMs predict paraphrase
<i>The threat that was coming was not from sleeper cells. Sleeper cells were the only threat of any significance.</i>
<i>The cover story reviews the latest research on how babies think. The cover story talks about how infants make decisions.</i>
<i>(a) Change each d or t in the target to c. After the conversion is finished the target should have exactly four c's.</i>
False Negative examples: LLMs predict non-paraphrase
<i>The girl was saddened by the disaster. The disaster saddened the girl.</i>
<i>Behavioral problems. Problems behaving.</i>
<i>The boy looked up the answer. The boy looked the answer up.</i>

Table 5: Examples of LLM failures on NLI and AMR-repurposed data. In the first three pairs, all Llama3 prompts (P1, P2, and P3) incorrectly classify the pairs as paraphrases, even though they are not. In contrast, in the last three pairs, the prompts fail to assign a paraphrase label, although all examples are indeed paraphrases.

However, the error rate is higher on hard positives in the AMR-repurposed data. In 92% of the cases where one or more LLMs make an error, at least one prompt still makes the correct decision. Nevertheless, our analysis indicates that LLMs can be misled by relatively simple linguistic variations.

7 Conclusion

A key contribution of our work is the introduction of the **PARAPHRASUS** benchmark, which includes sentence pairs from ten domains and covers three paraphrase detection challenges. We are releasing this benchmark to the research community, with the goal of facilitating further exploration of paraphrase phenomena, evaluating paraphrase detection models and zero-shot capabilities of LLMs, and pinpointing specific areas for improvement.

Our experiments using **PARAPHRASUS** revealed several interesting insights:

1. None of the tested LLM and classifier setups demonstrated strong performance across the full spectrum of paraphrases captured by the benchmark, highlighting its objective nature and the need for system development.
2. Even modern LLMs as Llama3 failed to accurately and consistently detect paraphrases – in some cases, the passivization of the verb is all it takes to confuse all tested prompt variants.
3. Improving training strategies for smaller, more efficient models is challenging, but not

impossible. Specifically, carefully inserting a certain amount of easy negatives into a large training set of adversarial pairs helped to lower the average error.

4. Among all tested LLM prompting strategies, each exhibited strengths and weaknesses depending on the particular challenge within **PARAPHRASUS**. However, the simplest prompt – directly asking whether the sentence pairs are paraphrases – performed best with an average error of 20.9% across the benchmark.

Limitations & Future Work

Our proposed benchmark includes text pairs with different levels of semantic and lexical similarity from diverse domains and covers many interesting paraphrasing phenomena. Does it cover all possible phenomena related to paraphrases and paraphrasing? Clearly not. For example, an obvious limitation is its coverage of phenomena in languages other than English. While our dataset does include a total of 25,360 non-English samples, they are not spread across all ten parts of **PARAPHRASUS** (e.g., the AMR-sourced dataset is only English). Thus, expanding **PARAPHRASUS** to include more datasets in languages other than English, particularly low-resource languages, appears as a fruitful avenue for future work. We invite community collaboration to help extend **PARAPHRASUS**.

While many resources were used to perform our evaluations, regarding our selection of LLMs, we limited ourselves to Llama3 Instruct 8B quantized to 4 bits. Similarly, the basic architecture for our experiments with trained models was limited to XLM-RoBERTa_{base}. We believe that the benchmark can be viewed as an extended reasoning evaluation of LLMs, useful for comparing capabilities of different sizes and architectures. Specifically, among all the many ablations and different setups that we ran for both model types, there was none that showed consistently good performance.

Acknowledgements

This research is conducted under the project *Impresso – Media Monitoring of the Past II Beyond Borders: Connecting Historical Newspapers and Radio*. *Impresso* is a research project funded by the Swiss National Science Foundation (SNSF 213585) and the Luxembourg National Research Fund (17498891).

References

- Eneko Agirre, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and Weiwei Guo. 2013. [*SEM 2013 shared task: Semantic textual similarity](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, pages 32–43, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Duarte M Alves, José Pombal, Nuno M Guerreiro, Pedro H Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, et al. 2024. Tower: An open multilingual large language model for translation-related tasks. *arXiv preprint arXiv:2402.17733*.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. [Abstract Meaning Representation for sembanking](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Dis-course*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.
- Alberto Barrón-Cedeño, Marta Vila, M. Antònia Martí, and Paolo Rosso. 2013. [Plagiarism meets paraphrasing: Insights for the next generation in automatic plagiarism detection](#). *Computational Linguistics*, 39(4):917–947.
- Rahul Bhagat and Eduard Hovy. 2013. What is a paraphrase? *Computational linguistics*, 39(3):463–472.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, and Prafulla et al. Dhariwal. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Marco Böhler, Gregory Crane, Maria Moritz, and Alison Babeu. 2012. Increasing recall for text re-use in historical documents to support research in the humanities. In *Theory and Practice of Digital Libraries: Second International Conference, TPD 2012, Paphos, Cyprus, September 23-27, 2012. Proceedings 2*, pages 95–100. Springer.
- Alexander Budanitsky and Graeme Hirst. 2006. Evaluating wordnet-based measures of lexical semantic relatedness. *Computational linguistics*, 32(1):13–47.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2024. [Chateval: Towards better LLM-based evaluators through multi-agent debate](#). In *The Twelfth International Conference on Learning Representations*.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational And Psychological Measurement*, 20:37–46.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau, 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.
- Chunyuan Deng, Yilun Zhao, Xiangru Tang, Mark Gestein, and Arman Cohan. 2024. [Benchmark probing: Investigating data leakage in large language models](#). In *NeurIPS 2023 Workshop on Backdoors in Deep Learning - The Good, the Bad, and the Ugly*.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Markus Freitag, George Foster, David Grangier, and Colin Cherry. 2020. [Human-paraphrased references improve neural machine translation](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1183–1192, Online. Association for Computational Linguistics.
- Juri Ganitkevitch and Chris Callison-Burch. 2014. [The multilingual paraphrase database](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 4276–4283, Reykjavik, Iceland. European Language Resources Association (ELRA).

- Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2013. [PPDB: The paraphrase database](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 758–764, Atlanta, Georgia. Association for Computational Linguistics.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. [The WebNLG challenge: Generating text from RDF data](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 124–133, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Sebastian Gehrmann, Tosin Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Anuoluwapo Aremu, Antoine Bosselut, and Khyathi Raghavi et al. Chandu. 2021. [The GEM benchmark: Natural language generation, its evaluation and metrics](#). In *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM 2021)*, pages 96–120, Online. Association for Computational Linguistics.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022. [TRUE: Re-evaluating factual consistency evaluation](#). In *Proceedings of the Second DialDoc Workshop on Document-grounded Dialogue and Conversational Question Answering*, pages 161–175, Dublin, Ireland. Association for Computational Linguistics.
- Peter Kolb. 2009. [Experiments on the difference between semantic similarity and relatedness](#). In *Proceedings of the 17th Nordic Conference of Computational Linguistics (NODALIDA 2009)*, pages 81–88, Odense, Denmark. Northern European Association for Language Technology (NEALT).
- Kalpesh Krishna, Yixiao Song, Marzena Karpinska, John Wieting, and Mohit Iyyer. 2023. [Paraphrasing evades detectors of ai-generated text, but retrieval is an effective defense](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 27469–27500. Curran Associates, Inc.
- Kalpesh Krishna, John Wieting, and Mohit Iyyer. 2020. [Reformulating unsupervised style transfer as paraphrase generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 737–762, Online. Association for Computational Linguistics.
- Timothy Liu and De Wen Soh. 2022. [Towards better characterization of paraphrases](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8592–8601, Dublin, Ireland. Association for Computational Linguistics.
- Nitin Madnani, Necip Fazil Ayan, Philip Resnik, and Bonnie Dorr. 2007. [Using paraphrases for parameter tuning in statistical machine translation](#). In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 120–127, Prague, Czech Republic. Association for Computational Linguistics.
- Jonathan Mallinson, Rico Sennrich, and Mirella Lapata. 2017. [Paraphrasing revisited with neural machine translation](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 881–893, Valencia, Spain. Association for Computational Linguistics.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. [A SICK cure for the evaluation of compositional distributional semantic models](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 216–223, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Marcel Nawrath, Agnieszka Nowak, Tristan Ratz, Danilo Walenta, Juri Opitz, Leonardo Ribeiro, João Sedoc, Daniel Deutsch, Simon Mille, Yixin Liu, Sebastian Gehrmann, Lining Zhang, Saad Mahamood, Miruna Clinciu, Khyathi Chandu, and Yufang Hou. 2024. [On the role of summary content units in text summarization evaluation](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 272–281, Mexico City, Mexico. Association for Computational Linguistics.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Timothy Niven and Hung-Yu Kao. 2019. [Probing neural network comprehension of natural language arguments](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics.
- Juri Opitz. 2024. [A Closer Look at Classification Evaluation Metrics and a Critical Reflection of Common Evaluation Practice](#). *Transactions of the Association for Computational Linguistics*, 12:820–836.
- Marcus Fabius Quintilianus. 95. *Institutio Oratoria (Book X)*, loeb classical library edition, 1920 edition. Loeb. Public domain.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages

- 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Zacchary Sadeddine, Juri Opitz, and Fabian Suchanek. 2024. [A survey of meaning representations – from theory to practical utility](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2877–2892, Mexico City, Mexico. Association for Computational Linguistics.
- Muhammad Sharjeel, Paul Rayson, and Rao Muhammad Adeel Nawab. 2016. [UPPC - Urdu paraphrase plagiarism corpus](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 1832–1836, Portorož, Slovenia. European Language Resources Association (ELRA).
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, and Adam R. Brown et al. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.
- Brian Thompson and Matt Post. 2020. [Automatic machine translation evaluation in many languages via zero-shot paraphrasing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 90–121, Online. Association for Computational Linguistics.
- Teemu Vahtola, Mathias Creutz, and Jörg Tiedemann. 2022. [It is not easy to detect paraphrases: Analysing semantic similarity with antonyms and negation using the new SemAntoNeg benchmark](#). In *Proceedings of the Fifth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 249–262, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Sowmya Vajjala and Ivana Lučić. 2018. [OneStopEnglish corpus: A new corpus for automatic readability assessment and text simplification](#). In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 297–304, New Orleans, Louisiana. Association for Computational Linguistics.
- John Wieting and Kevin Gimpel. 2018. [ParaNMT-50M: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Melbourne, Australia. Association for Computational Linguistics.
- Yinfei Yang, Yuan Zhang, Chris Tar, and Jason Baldridge. 2019. [PAWS-X: A cross-lingual adversarial dataset for paraphrase identification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3687–3692, Hong Kong, China. 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.
- Jianing Zhou and Suma Bhat. 2021. [Paraphrase generation: A survey of the state of the art](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5075–5086, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. [Texygen: A benchmarking platform for text generation models](#). In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1097–1100.

A ICL Prompt

In-Context-Learning (k=4) Paraphrase Detection Prompt	
Here are example sentence pairs that are {paraphrase notion} (Yes) and not {paraphrase notion} (No):	
Sentence 1: {RandPos1.sentence1}	
Sentence 2: {RandPos1.sentence2}	
Yes	
Sentence 1: {RandNeg1.sentence1}	
Sentence 2: {RandNeg1.sentence2}	
No	
...	
Are the following sentences {paraphrase notion}?	
Sentence 1: {sentence1}	
Sentence 2: {sentence2}	
Answer with 'Yes' or 'No'	

Figure 4: In-Context-Learning Prompt Template. For P1, P2, and P3, the questions asked are “paraphrases,” “semantically equivalent,” and “expressing the same content,” respectively.

B Per Language Performance

Model	en (E%)	de (E%)	fr (E%)	es (E%)	ja (E%)	ko (E%)	zh (E%)
XLM-R← PAWS-EN Train	6.15	12.40	11.53	11.27	23.18	23.09	19.03
Llama3 Instruct P1	34.75	43.90	43.30	43.60	49.70	50.20	47.25
Llama3 Instruct P2	33.60	37.55	39.45	38.35	45.95	46.55	43.65
Llama3 Instruct P3	29.65	35.00	35.85	34.90	45.00	45.15	40.85
Llama3 Instruct P1-ICL_K4	35.50	37.65	38.40	39.10	40.25	43.50	38.25
Llama3 Instruct P2-ICL_K4	29.90	31.20	33.05	33.85	37.45	38.70	34.70
Llama3 Instruct P3-ICL_K4	27.95	31.15	32.20	32.85	36.55	38.50	33.30

Table 6: Error rates of Table 1 for the PAWSX test set per language.

Model	en (E%)	de (E%)	fr (E%)	es (E%)	zh (E%)
XLM-R← PAWS-EN Train	26.99	24.70	29.88	29.94	21.96
Llama3 Instruct P1	15.18	12.71	13.16	11.87	8.74
Llama3 Instruct P2	2.23	1.44	1.18	1.15	0.79
Llama3 Instruct P3	2.40	1.24	1.14	1.20	0.63
Llama3 Instruct P1-ICL_K4	4.82	3.10	2.60	1.98	1.44
Llama3 Instruct P2-ICL_K4	1.00	0.19	0.15	0.12	0.04
Llama3 Instruct P3-ICL_K4	0.78	0.13	0.15	0.13	0.06

Table 7: Error rates of Table 1 for the XNLI subset dataset for the five languages we use.

C Additional Datasets Statistics




Figure 5: Average Word Position Deviation (WPD) and Lexical Diversity (LD) (Liu and Soh, 2022) of the symmetric datasets of PARAPHRASUS.




Figure 6: Average sentence lengths. Left plot shows symmetric datasets and right plot shows the asymmetric NLI datasets of PARAPHRASUS