

NMTSCORE: A Multilingual Analysis of Translation-based Text Similarity Measures

Jannis Vamvas¹ and Rico Sennrich^{1,2}

¹Department of Computational Linguistics, University of Zurich

²School of Informatics, University of Edinburgh

{vamvas, sennrich}@cl.uzh.ch

Abstract

Being able to rank the similarity of short text segments is an interesting bonus feature of neural machine translation. Translation-based similarity measures include direct and pivot translation probability, as well as *translation cross-likelihood*, which has not been studied so far. We analyze these measures in the common framework of multilingual NMT, releasing the NMTSCORE library. Compared to baselines such as sentence embeddings, translation-based measures prove competitive in paraphrase identification and are more robust against adversarial or multilingual input, especially if proper normalization is applied. When used for reference-based evaluation of data-to-text generation in 2 tasks and 17 languages, translation-based measures show a relatively high correlation to human judgments.

1 Introduction

Measures of paraphrastic similarity aim to quantify the degree to which text segments mean the same thing. Such measures can be used to identify paraphrases, and also to automatically evaluate text generation by estimating the similarity of model outputs to human-written references.

Neural machine translation (NMT) enables several similarity measures as a by-product of learning to estimate the probability of translations (Mallinson et al., 2017; Junczys-Dowmunt, 2018; Thompson and Post, 2020). These measures are promising given that they naturally leverage parallel corpora and might pay more attention to details such as word order or named entities than sentence embeddings do. For example, Zhang et al. (2019) have demonstrated that it is difficult to spot the mismatch between “*Flights from New York to Florida*” and “*Flights from Florida to New York*” purely based on pooled representations, and they have released a challenge set of such paraphrase adversaries called PAWS.

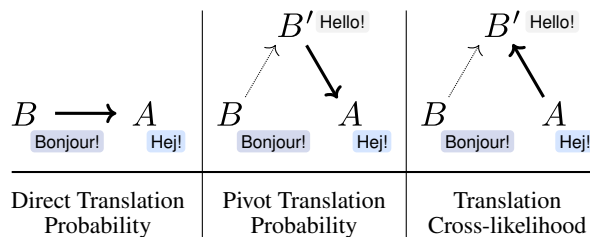


Figure 1: Three text similarity measures are analyzed in this paper. Each measure uses a specific translation probability (thick arrow) to estimate the similarity of sentence A to B . Some measures require a translation B' into an intermediate language (thin arrow).

In this paper, we consider three distinct subtypes of translation-based similarity, as visualized in Figure 1. First of all, a straightforward approach is to estimate the *direct translation probability* of sentence A when translating from sentence B (Junczys-Dowmunt, 2018; Thompson and Post, 2020). Secondly, *pivot translation probability* is an estimate of how probable it is to arrive at sentence A when pivoting through an intermediary language (Mallinson et al., 2017). Finally, we propose to estimate *translation cross-likelihood*: the likelihood that a translation of B into some language could also be a translation of A .

We release the NMTSCORE library, using many-to-many multilingual NMT to implement these measures.¹ All three measures are competitive in paraphrase identification across 9 languages, compared to other general-purpose similarity measures, and especially on adversarial examples. We find, however, that normalizing the measures with reconstruction probability is important. They also work well cross-lingually, where pivot translation probability performs best while cross-likelihood has the advantage of not requiring an explicit specification of the input’s languages.

With respect to reference-based evaluation, we

¹<https://github.com/ZurichNLP/nmtscore>

show that NMTSCORE is a competitive evaluation metric for data-to-text, according to human judgments from the WebNLG Challenge (Ferreira et al., 2020) and from a multilingual AMR-to-text evaluation (Fan and Gardent, 2020). Taken together, multilingual NMT offers a relatively precise and often complementary perspective on paraphrastic similarity, with little correlation to other metrics.

In summary, we make the following main contributions:

- We redefine translation-based similarity measures in a common multilingual framework, proposing *reconstruction normalization* and a novel *translation cross-likelihood* measure.
- We compile a multilingual paraphrase identification benchmark, showing that NMTSCORE outperforms other general-purpose measures.
- We demonstrate that NMTSCORE provides effective metrics for reference-based evaluation of data-to-text generation.

2 Translation-based Similarity Measures

The similarity of two sentences A and B can be measured in several ways by using a multilingual translation model θ . Such a model accepts multiple source languages and can translate into multiple target languages. Usually only the target language needs to be specified, e.g. with a target language token, and thus we use θ_ℓ to denote a model that is conditioned on a target language ℓ .

2.1 Direct Translation Probability

Cross-lingual If A and B are in two different languages a and b , the model can directly estimate the translation probability of A given B :

$$P_{\text{direct}}(A|B) = p_{\theta_a}(A|B)$$

This probability is sometimes called the *translational equivalence* of A and B . In practice, there are different ways how such a probability can be calculated from the token-level probabilities predicted by the model. For this and the following measures, we follow previous work (Junczys-Dowmunt, 2018; Thompson and Post, 2020) and normalize by sequence length:

$$p_{\theta_a}(A|B) := \left[\prod_{i=0}^{|A|} p_{\theta_a}(A^i|B, A^{<i}) \right]^{\frac{1}{|A|}}$$

Monolingual Since θ is a multilingual model, A and B may also be in the same language (Thompson and Post, 2020). This is because multilingual NMT enables zero-shot translation (Johnson et al., 2017), which includes any monolingual direction $\ell \rightarrow \ell$. Thompson and Post (2020) argue that monolingual translation can be seen as (non-diverse) paraphrasing, and they have demonstrated that paraphrasing probability is a useful metric for reference-based MT evaluation (called *Prism*).

2.2 Pivot Translation Probability

Paraphrastic similarity can also be estimated via translation to a pivot language (Bannard and Callison-Burch, 2005; Mallinson et al., 2017). Pivot translation requires two translation directions, θ_{pivot} and θ_a . First, a translation $B' \sim p_{\theta_{\text{pivot}}}(\cdot|B)$ is generated and then used to calculate the probability of translating B' into A :

$$P_{\text{pivot}}(A|B) = p_{\theta_a}(A|B')$$

Such an approach is typically used for monolingual sentences (*round-trip translation*), but we argue that A and B can also be in two different languages. Furthermore, the pivot language may be identical to the language of A or B , considering the zero-shot paraphrasing capability of multilingual NMT.

2.3 Translation Cross-likelihood

As an alternative measure we propose *translation cross-likelihood*, which requires only one translation direction θ_{tgt} , where *tgt* is any target language supported by the NMT model. We generate a translation $B' \sim p_{\theta_{\text{tgt}}}(\cdot|B)$ and then estimate the likelihood that B' could have been generated from A :

$$\text{Cross-likelihood}(A|B) = p_{\theta_{\text{tgt}}}(B'|A)$$

In other words, this similarity reflects the surprisal of a translation model that is conditioned on sentence A but exposed to a translation of sentence B .

Like with pivot translation, A and B may be a monolingual or a cross-lingual pair, and the target language may or may not be identical to the language of A and B .

2.4 From Probability to Similarity Measure

Normalization Similarity measures typically assign maximum similarity to indiscernible inputs. We propose to ensure this by applying the follow-

ing normalizations to translation-based measures:

$$\text{NMTSCORE-direct}(A|B) = \frac{p_{\theta_a}(A|B)}{p_{\theta_a}(A|A)}$$

$$\text{NMTSCORE-pivot}(A|B) = \frac{p_{\theta_a}(A|B')}{p_{\theta_a}(A|A')}$$

$$\text{NMTSCORE-cross}(A|B) = \frac{p_{\theta_{tgt}}(B'|A)}{p_{\theta_{tgt}}(B'|B)}$$

The two formulas for NMTSCORE-direct and NMTSCORE-pivot can be seen as a form of *reconstruction normalization*, given that $p_{\theta_a}(A|A)$ is the probability that the sentence remains identical when zero-shot paraphrasing is performed. Likewise, $p_{\theta_a}(A|A')$ is the pivot reconstruction probability of A given its pivot translation $A' \sim p_{\theta_{pivot}}(\cdot|A)$. The latter measure has also been used by Mallinson et al. (2017) for normalization.

Symmetrization Translation probabilities are directed measures, however the order of A and B is often arbitrary, such as in paraphrase identification. We thus follow previous work (Junczys-Dowmunt, 2018; Thompson and Post, 2020) and average the directed measures for both directions:

$$\text{sim}(A, B) = \frac{1}{2} \text{sim}(A|B) + \frac{1}{2} \text{sim}(B|A)$$

3 Baseline Measures

3.1 Surface Similarity Measures

Some well-known text similarity measures, especially for reference-based evaluation, rely on surface similarity. We choose CHRF (Popović, 2015) and sentence-level BLEU (Papineni et al., 2002) as surface-similarity baselines. CHRF is a character-based metric that calculates precision and recall of character n -grams. BLEU calculates the precision of word n -grams with a brevity penalty.

3.2 Embedding-based Similarity Measures

Another family of similarity measures uses the cosine similarity of text embeddings. Such embeddings are typically learned on the token level, and thus need to be aggregated in some way. In this paper, we consider two embedding baselines:

Similarity of aggregate token embeddings A typical approach is to average the token embeddings before calculating cosine similarity. However, hidden states of self-supervised Transformer language models may not be directly useful when

averaged; Reimers and Gurevych (2019) fine-tune them using a sentence pair classification objective, calling their approach Sentence-BERT.

Aggregation of token similarities An alternative is to aggregate the similarities between the individual tokens of the two sentences (Mihalcea et al., 2006). Typically, a *precision* is calculated as the average maximum cosine similarity of all tokens in A to all tokens in B , and a *recall* is calculated with A and B switched. It has been shown that when aggregated in this way, hidden states of self-supervised language models are useful for text similarity even without any fine-tuning (Mathur et al., 2019; Zhang et al., 2020). This measure is called BERTSCORE by Zhang et al. (2020).

4 Paraphrase Identification

In this section we compare the similarity measures using paraphrase identification test sets in multiple languages. The test sets contain pairs of sentences that have been annotated with whether the sentences are paraphrases or not, yielding a binary classification problem. For datasets with a validation split, we determine thresholds that optimally separate the validation set for each measure; we then apply the thresholds to the test set to compute the accuracy of the measures. If there is no validation set, we report the Area Under the Curve (AUC) on the test set.

4.1 Experimental Setup

Translation model We use the same multilingual NMT model for all three translation-based measures. Specifically, we use a 745M-parameter Transformer model (*Prism*) that was trained by Thompson and Post (2020) using Fairseq (Ott et al., 2019).² The model supports 39 languages and is not English-centric. We found no indication that its training data overlap with the datasets used in the experiments.

Intermediary language We use English as the pivot language for pivot translation and as the target language for estimating cross-likelihood. English has the largest share of training data in the models we use.

Surface Similarity baselines We compute sentence-level CHRF and sentence-level BLEU using the SacreBLEU library (Post, 2018). We

²<https://github.com/thompsonb/prism>

Language Metric	Individual datasets				PAWS-X dataset					Macro- average	
	EN Acc.	RU AUC	FI AUC	SV AUC	DE Acc.	ES Acc.	FR Acc.	JA Acc.	ZH Acc.	Avg.	
<i>Surface similarity baselines</i>											
CHRF	69.9	78.4	58.9	66.8	58.3	57.2	58.5	56.2	59.3	57.9	66.4
SENTBLEU	64.2	70.1	63.6	62.2	61.3	59.9	61.2	55.4	59.0	59.4	63.9
<i>Embedding baselines</i>											
Sentence-BERT	71.2	83.5	68.7	73.0	58.2	58.2	59.5	55.5	58.1	57.9	70.9
BERTSCORE-F1	72.6	79.7	66.6	65.9	59.8	60.2	60.5	56.7	59.5	59.4	68.8
<i>Translation-based measures</i>											
NMTSCORE-direct	72.6	84.1	72.4	70.6	73.9	73.5	75.7	66.4	68.9	71.7	74.3
NMTSCORE-pivot	72.1	84.9	70.3	70.9	77.4	76.2	76.9	68.4	70.8	74.0	74.4
NMTSCORE-cross	71.7	86.6	71.2	72.4	76.6	75.1	75.6	65.8	70.5	72.7	74.9

Table 1: Comparison of general-purpose text similarity measures on paraphrase identification in 9 languages. While the embedding baselines build on XLM-RoBERTa (Conneau et al., 2020), the translation-based measures use probability estimates of a multilingual NMT system (Prism, Thompson and Post, 2020). We report AUC if there is no validation set for a language. Results within the top significance cluster are printed in bold.

use the recommended tokenization for BLEU, tokenizing Japanese text using MeCab and splitting Chinese characters individually. When applying CHRF and BLEU to paraphrase identification, we calculate the similarity in both directions and take the average.

Embedding baselines We use pre-trained embeddings from XLM-RoBERTa, which is a multilingual masked language model pre-trained on CommonCrawl (Conneau et al., 2020). We use the large version (550M parameters) to compute BERTSCORE, specifically the 17th layer as recommended by the BERTSCORE reference implementation³. For Sentence-BERT, we use a version of size ‘base’ (270M parameters) that Reimers and Gurevych (2020) have finetuned by distilling the sentence embeddings of an English RoBERTa model. The latter has in turn been fine-tuned on 50M English paraphrase pairs, which do not overlap with our test sets. The distillation was performed using parallel sentences for 50 languages.

4.2 Datasets

We use the following datasets for our experiments (statistics are reported in Appendix D):

English MRPC (Dolan and Brockett, 2005), a corpus of sentence pairs automatically extracted from news, and annotated with binary labels.

³https://github.com/Tiiiger/bert_score

We exclude samples where a re-annotation effort by Kovatchev et al. (2018) has found inconsistent labeling.

Russian ParaPhraser (Pivovarova et al., 2018), a corpus of news headlines annotated on a three-class ordinal scale. We follow the original setup and create binary labels by merging *precise* and *near* paraphrases into a single class.

Finnish and Swedish The Finnish Paraphrase Corpus (Kanerva et al., 2021), a dataset of manually selected subtitle lines and news headlines, annotated on a four-class ordinal scale. A small test set in Swedish is likewise available. We create binary labels by categorizing all pairs with a label of 4 as positive, and all below as negative.

PAWS-X We also include PAWS-X (Yang et al., 2019), a challenge set of sentence pairs with high word overlap in a total of 7 languages. The dataset is based on English sentences extracted from Wikipedia that have been paired with automatically created derivative sentences, and annotated with binary labels (Zhang et al., 2019). Test sets in other languages have been created by manually translating the English sentences. We report results for German, Spanish, French, Japanese, Chinese and an average over these languages. We do not report results for English, since a part of the positive examples in the original English dataset have been created with automatic round-trip translation. The

Sentence A	Sentence B	Sentence-BERT	BERTSCORE	NMTSCORE
<i>Flights from New York to Florida.</i>	↔ <i>Flights from Florida to New York.</i>	1.00	0.94	0.67
	↔ <i>Flights to Florida from New York.</i>	0.99 ↓	0.89 ↓	0.71 ↑
	↔ <i>Flights from NYC to Florida.</i>	0.96 ↓	0.92 ↓	0.72 ↑
	↔ <i>Flights to Florida from NYC.</i>	0.95 ↓	0.81 ↓	0.53 ↓

Figure 2: An adversarial example originally used by Zhang et al. (2019), judged by three similarity metrics: Sentence-BERT cosine similarity, rescaled BERTSCORE-F1 and NMTSCORE-cross. The arrows indicate whether a paraphrase pair is assigned a higher or lower similarity than the pair in the first row, which is a non-paraphrase.

other languages are better suited for our analysis, because the manual translation process has broken the direct link between the original sentence and its round-trip translation. We also do not report results for Korean, since the *Prism* translation model we use for our main experiments has not been trained in that language.

4.3 Statistical Analysis

For all comparisons between the translation-based and baseline measures we broadly follow the methodology of the WMT Metrics Task (Freitag et al., 2021). We perform paired bootstrap resampling (Koehn, 2004) with 1000 repetitions to assess the statistical significance of a difference between two measures at $\alpha = 0.05$. For each dataset, we determine the *top significance cluster*, which is printed in bold. This cluster contains the top measures that are not significantly outperformed by any other measure. We also perform significance tests for average metrics over multiple datasets by combining the i th bootstrap sample of every dataset into the i th bootstrap sample of the full benchmark. Our analysis was implemented using the SacreROUGE framework (Deutsch and Roth, 2020).

4.4 Monolingual Results

The results of monolingual paraphrase identification are shown in Table 1. In the final column we report the macro-average over all datasets, i.e., the average of EN accuracy, RU AUC, FI AUC, SV AUC and average PAWS-X accuracy.

Overall, the translation-based measures perform better than the baselines that use surface similarity or embeddings. They excel on the adversarial PAWS-X dataset, with an improvement of 10–15 points over the baselines. On the other datasets the accuracy of the translation-based measures is comparable to the embedding baselines.

It should be noted that the embedding baselines and the translation-based measures are not perfectly comparable, since multilingual models with different hyperparameters and pretraining languages are used. The three translation-based measures are perfectly comparable, since they rely on the same NMT model. Translation cross-likelihood has a slightly higher accuracy than direct or pivot translation probability in monolingual paraphrase detection. Appendix A7 shows that this finding can be reproduced with an alternative multilingual NMT system in two different sizes (*M2M-100*; Fan et al., 2021).

4.5 Qualitative Analysis

A qualitative comparison suggests that translation-based measures are superior in distinguishing numbers, named entities and enumerations, but that embeddings can better capture the similarity of sentences with similar meaning but very different phrasing (Appendices F and G).

Figure 2 illustrates this observation on the example of the adversarial pair mentioned in the introduction. Both Sentence-BERT and BERTSCORE assign a higher score to the non-paraphrase with a high word overlap, and a lower score to paraphrases with a difference in word order or entity naming. NMTSCORE accurately outputs a higher score for the two paraphrases than for the non-paraphrase but still fails on a fourth sentence pair that combines both phenomena.

4.6 Cross-lingual Results

As discussed in Section 2, all three translation-based measures can be applied to both monolingual and cross-lingual sentence pairs. The same holds for the baseline measures, even though this use case has been less prominent in previous work. We rearrange the PAWS-X dataset to create a cross-lingual

	EN + DE	EN + ES	EN + FR	EN + JA	EN + ZH	DE + ES	DE + FR	DE + JA	DE + ZH	ES + FR	ES + JA	ES + ZH	FR + JA	FR + ZH	JA + ZH	Avg.
<i>Surface similarity baselines</i>																
CHRF	54.9	54.9	54.3	54.6	54.8	54.9	54.5	54.5	54.6	54.7	54.6	54.8	54.5	54.7	54.7	54.7
SENTBLEU	56.4	56.2	56.3	54.6	54.7	56.0	56.2	54.5	54.5	56.1	54.6	54.7	54.5	54.6	54.7	55.2
<i>Embedding baselines</i>																
Sentence-BERT	53.9	54.7	54.8	54.6	54.7	55.5	56.0	54.5	54.5	56.2	54.6	54.6	54.5	54.6	54.7	54.8
BERTSCORE-F1	57.4	57.5	57.0	55.1	54.8	57.2	57.3	54.6	54.5	57.6	54.7	54.9	54.6	54.6	55.0	55.8
<i>Translation-based measures</i>																
NMTSCORE-direct	76.4	76.4	76.1	68.6	68.8	73.3	74.5	66.0	66.9	74.3	66.7	66.8	66.8	67.4	64.4	70.2
NMTSCORE-pivot	77.4	76.9	77.3	68.9	70.7	75.0	76.0	67.0	69.5	75.5	67.6	69.5	67.5	69.9	66.5	71.7
NMTSCORE-cross	76.0	75.9	75.9	65.2	66.0	74.5	75.2	64.8	65.8	74.2	64.6	66.2	64.4	65.7	65.3	69.3

Table 2: Comparison of text similarity measures on cross-lingual paraphrase identification using the PAWS-X dataset. Results within the top significance cluster are printed in bold.

	CHRF	SENTBLEU	SBERT	BERTSc.	NMTSC.-dir.	NMTSC.-pivot
SENTBLEU	0.56±.01					
Sentence-BERT	0.42±.01	0.35±.01				
BERTSCORE	0.52±.01	0.45±.01	0.46±.01			
NMTSCORE-direct	0.49±.01	0.45±.01	0.52±.01	0.54±.01		
NMTSCORE-pivot	0.44±.01	0.40±.01	0.50±.01	0.47±.01	0.77±.01	
NMTSCORE-cross	0.47±.01	0.41±.01	0.52±.01	0.52±.01	0.74±.01	0.75±.01

Table 3: Sample-level Kendall correlation between the measures analyzed in this paper, averaged across the 5 datasets in our paraphrase identification benchmark. We report confidence intervals with bootstrap resampling.

version of the benchmark that covers 15 language pairs. This is possible because the non-English versions of PAWS-X are translations from the English version. For example, we can create a cross-lingual sentence pair by pairing the English version of sentence A with the German version of sentence B.

Note that even though this process makes use of translation, simply reversing the translation process is not sufficient for a similarity measure to solve the resulting task, unlike with retrieval tasks that are constructed from parallel data. As such, the cross-lingual PAWS-X tasks are at least as hard as the original monolingual PAWS-X tasks.

Results are reported in Table 2. Compared to monolingual PAWS-X, all measures perform worse by up to 4 points on average. Overall, pivot translation probability is the most accurate similarity measure for cross-lingual paraphrase identification. It should be noted, however, that both direct and pivot translation probability require the input languages to be specified, while translation cross-likelihood (like Sentence-BERT and BERTSCORE) allows them to remain unspecified.

4.7 Correlation to Alternative Measures

Calculating the pairwise correlation between the measures allows us to learn about similarities between the measures. Table 3 visualizes the average Kendall correlations on the monolingual paraphrase identification datasets. The translation-based measures form a cluster with a high mutual correlation, but still seem to behave differently to some degree, especially cross-likelihood.

4.8 Effect of Normalization

Table 4 presents an ablation study for the normalizations that we applied to the similarity measures (Section 2.4). Overall, reconstruction normalization leads to a clear improvement. On the English dataset, normalization does not have a positive effect on pivot translation probability and cross-likelihood. However, since we also use English as an intermediary language for these measures, we do not think that this special case should affect our conclusions regarding the ablation study.

Language Metric	Individual datasets				PAWS-X dataset						Macro- average
	EN Acc.	RU AUC	FI AUC	SV AUC	DE Acc.	ES Acc.	FR Acc.	JA Acc.	ZH Acc.	Avg.	
NMTSCORE-direct	72.6	<u>84.1</u>	<u>72.4</u>	<u>70.6</u>	<u>73.9</u>	73.5	<u>75.7</u>	66.4	<u>68.9</u>	<u>71.7</u>	<u>74.3</u>
– no normalization [†]	72.3	83.3	65.8	67.8	71.8	72.8	73.1	62.0	67.1	69.3	71.7
NMTSCORE-pivot	72.1	<u>84.9</u>	<u>70.3</u>	<u>70.9</u>	<u>77.4</u>	<u>76.2</u>	<u>76.9</u>	<u>68.4</u>	<u>70.8</u>	<u>74.0</u>	<u>74.4</u>
– no normalization	73.0	81.3	64.4	66.5	70.2	71.4	70.9	61.5	63.5	67.5	70.5
NMTSCORE-cross	71.7	<u>86.6</u>	<u>71.2</u>	<u>72.4</u>	<u>76.6</u>	<u>75.1</u>	75.6	<u>65.8</u>	<u>70.5</u>	<u>72.7</u>	<u>74.9</u>
– no normalization	71.9	86.1	70.6	71.5	75.2	74.3	75.5	65.0	69.3	71.8	74.4

Table 4: Ablation study for the reconstruction normalizations proposed in Section 2.4. The measure marked with (†) corresponds to the *Prism* measure (Thompson and Post, 2020). Underlined results are results that are significantly better than the other variant; in most cases reconstruction normalization leads to a significant improvement.

5 Evaluation of Data-to-Text Generation

We now turn to a different application of sentence similarity, namely the reference-based evaluation of data-to-text generation. While the good performance of the translation-based measures on paraphrase identification is encouraging, this setting poses slightly different requirements on text similarity metrics. Specifically, it is not necessary that the similarities of paraphrases and non-paraphrases are completely separable, but only that multiple hypotheses are correctly ranked with respect to a shared reference. Moreover, it is relevant in which direction the measure is calculated. Below we separately evaluate both directions: $\text{sim}(\text{hyp}|\text{ref})$ and $\text{sim}(\text{ref}|\text{hyp})$, as well as their average.

5.1 RDF-to-text

The WebNLG 2020 challenge (Ferreira et al., 2020) includes a task that requires generating natural language sentences from RDF triple sets. Human annotators have rated system output in English and Russian using five criteria: *data coverage*, *relevance*, *correctness*, *text structure*, and *fluency*. In this paper we average the first three criteria to calculate an overall judgment of adequacy for each submitted sample (averaging first across annotators, then across individual criteria).

Since the dataset contains multiple references per RDF triple set (statistics are reported in Table A5), we need to aggregate the scores computed by the automated metrics. We follow previous work and select the maximum score across the references.

Figure 3 shows the Kendall correlation between the similarity measures and the human judgments.

We report correlation on the level of the individual samples (also called *global* correlation). Since such meta-evaluations of metrics are known to have high statistical uncertainty, we follow Deutsch et al. (2021) and estimate confidence intervals using a Boot-Both technique. As we did before, we perform pairwise hypothesis tests and visualize the top significance cluster for each language.

Overall, direct translation probability has the highest correlation to human judgments of adequacy, indicating that translation-based measures are a competitive evaluation metric for RDF-to-text generation.

5.2 AMR-to-text

While the WebNLG dataset contains the output of various systems, it only encompasses two languages. We thus complement our analysis with data collected by Fan and Gardent (2020) to evaluate a single multilingual AMR-to-text system in 15 additional European languages. This dataset contains 50 sentences per language, with ratings by up to 10 native speakers along the criteria *morphology*, *word order*, *semantic accuracy*, and *good paraphrases*. The languages are listed in Table A6.

Here, we focus on semantic accuracy. We calculate the Kendall correlation between each metric and the average human rating individually per language, and then report the average across all languages. Figure 4 visualizes the confidence intervals, showing that the translation-based measures are more reliable in judging semantic accuracy than the baseline measures.

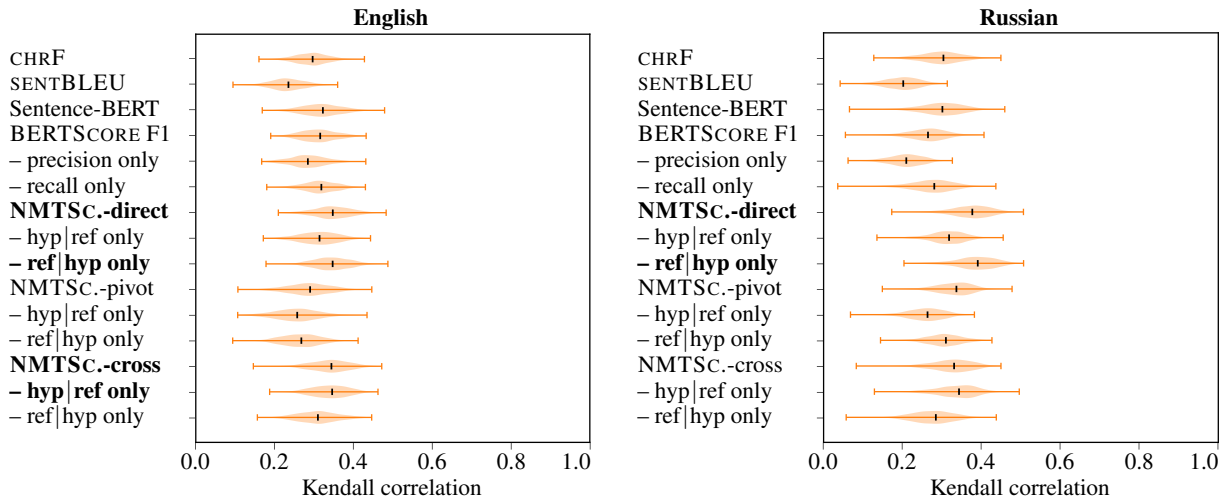


Figure 3: Sample-level Kendall correlation of text similarity measures to human judgments of WebNLG 2020 RDF-to-text submissions. A black tick denotes the correlation coefficient, and error bars denote 95% confidence intervals. Measures in the top significance cluster are printed in bold.

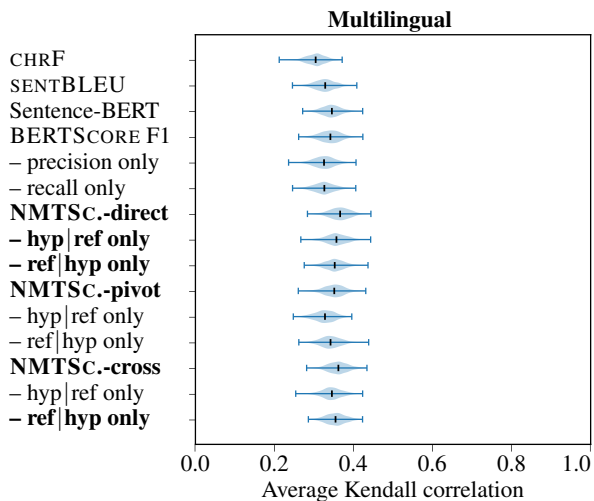


Figure 4: Sample-level Kendall correlation of text similarity measures to human judgments of AMR-to-text output (Fan and Gardent, 2020), averaged across 15 languages. A black tick denotes the average correlation, and error bars denote 95% confidence intervals. Measures in the top significance cluster are printed in bold.

6 Related Work

Various strategies have been suggested to leverage translation for paraphrastic similarity. Most related to translation cross-likelihood is perhaps the work of Barzilay and McKeown (2001), who extracted *multiple translations* from a bilingual parallel corpus assuming that sentences that are aligned to the same counterpart in the other language have similar meaning. In this paper, we revisit this idea in the context of NMT. A different approach was pursued by Bannard and Callison-Burch (2005),

who applied phrase-based statistical MT to estimate round-trip translation probability. This corresponds to the NMT approach of Mallinson et al. (2017), who also explore the use of multiple translation variants (*multi-pivoting*) or multiple pivot languages (*multilingual pivoting*). In a variation of this approach, Wieting et al. (2017) and Wieting and Gimpel (2018) used round-trip translation to generate training data for a paraphrastic sentence embedding model.

Multilingual MT allows to avoid pivoting by using zero-shot paraphrasing. This has been exploited for model analysis (Tiedemann and Scherrer, 2019) and reference-based evaluation (Thompson and Post, 2020). Agrawal et al. (2021) investigate alternative techniques to estimate direct translation probability for quality estimation. In the context of parallel corpus filtering (Junczys-Dowmunt, 2018), Chen et al. (2020) propose trie-constrained decoding to improve the efficiency of pairwise comparisons. Future work could apply their method to the other translation-based measures.

Similarity of NMT representations An alternative line of research has used representations of NMT encoders to compare sentences, ever since it has been demonstrated that such representations can be informative (Cho et al., 2014; Sutskever et al., 2014). While bilingual NMT has not been found to be particularly useful for unsupervised similarity (Hill et al., 2016; Cifka and Bojar, 2018), multilingual NMT representations have proven more successful (Schwenk and Douze 2017; Johnson et al. 2017; among others). However, rep-

resentation learning approaches that use parallel training data without an explicit translation objective are highly competitive (Wieting et al., 2019; Conneau et al., 2020; Hu et al., 2021), raising the question whether translation is indeed necessary for embedding-based similarity measures.

7 Conclusion

Our analysis highlights theoretical and empirical properties of translation-based text similarity measures in a multilingual setting. Direct translation probability is the most straightforward measure (an empirical comparison of inference times is found in Appendix A). However, it treats inputs as target sequences, and we show that accuracy on paraphrase identification can be clearly improved by normalizing with reconstruction probability.

Pivot translation probability is advantageous especially when performing cross-lingual comparisons. Finally, translation cross-likelihood has the advantage that it achieves symmetry with a single translation direction, and that the input languages need not be specified. The latter property also has interesting consequences for reference-based evaluation: The metric is expected to ignore whether the generated text matches the language of the reference. This can be seen as a rigorous disentanglement of adequacy from fluency.

In comparison to baseline measures, translation-based measures are generally slower but show high accuracy on multilingual paraphrase identification, comparatively good reliability on reference-based evaluation of data-to-text generation, and little correlation to alternative measures. Our findings thus show the usefulness of NMT translation probabilities for similarity tasks that require high attention to detail.

Limitations

The experiments in this paper are performed on mid- and high-resource languages. MT on low-resource languages might not yet be good enough for translation-based similarity measures to be useful. Still, our analysis extends to more languages than previous work, including languages that have little relatedness to English. Another limitation of translation-based text similarity measures is that the maximum sequence length supported by NMT models is often relatively short. In Appendix D we report the average character count of the text sequences used in this paper.

Acknowledgments

This work was funded by the Swiss National Science Foundation (project MUTAMUR; no. 176727). We would like to thank Farhad Nooralahzadeh for helpful feedback, Angela Fan for making human judgment data available, and Iker García-Ferrero for reporting a bug in an earlier version of the library.

References

- Sweta Agrawal, George Foster, Markus Freitag, and Colin Cherry. 2021. [Assessing reference-free peer evaluation for machine translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1158–1171, Online. Association for Computational Linguistics.
- Colin Bannard and Chris Callison-Burch. 2005. [Paraphrasing with bilingual parallel corpora](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pages 597–604, Ann Arbor, Michigan. Association for Computational Linguistics.
- Regina Barzilay and Kathleen R. McKeown. 2001. [Extracting paraphrases from a parallel corpus](#). In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics*, pages 50–57, Toulouse, France. Association for Computational Linguistics.
- Pinzhen Chen, Nikolay Bogoychev, Kenneth Heafield, and Faheem Kirefu. 2020. [Parallel sentence mining by constrained decoding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1672–1678, Online. Association for Computational Linguistics.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Ondřej Cífka and Ondřej Bojar. 2018. [Are BLEU and meaning representation in opposition?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1362–1371, Melbourne, Australia. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In

- Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Daniel Deutsch, Rotem Dror, and Dan Roth. 2021. [A Statistical Analysis of Summarization Evaluation Metrics Using Resampling Methods](#). *Transactions of the Association for Computational Linguistics*, 9:1132–1146.
- Daniel Deutsch and Dan Roth. 2020. [SacreROUGE: An open-source library for using and developing summarization evaluation metrics](#). In *Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS)*, pages 120–125, Online. Association for Computational Linguistics.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *Journal of Machine Learning Research*, 22(107):1–48.
- Angela Fan and Claire Gardent. 2020. [Multilingual AMR-to-text generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2889–2901, Online. Association for Computational Linguistics.
- Thiago Castro Ferreira, Claire Gardent, Nikolai Ilinykh, Chris Van Der Lee, Simon Mille, Diego Moussallem, and Anastasia Shimorina. 2020. [The 2020 Bilingual, Bi-Directional WebNLG+ Shared Task Overview and Evaluation Results \(WebNLG+ 2020\)](#). In *Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+)*, Dublin/Virtual, Ireland.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, George Foster, Alon Lavie, and Ondřej Bojar. 2021. [Results of the WMT21 metrics shared task: Evaluating metrics with expert-based human evaluations on TED and news domain](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 733–774, Online. Association for Computational Linguistics.
- Felix Hill, Kyunghyun Cho, and Anna Korhonen. 2016. [Learning distributed representations of sentences from unlabelled data](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1367–1377, San Diego, California. Association for Computational Linguistics.
- Junjie Hu, Melvin Johnson, Orhan Firat, Aditya Siddhant, and Graham Neubig. 2021. [Explicit alignment objectives for multilingual bidirectional encoders](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3633–3643, Online. Association for Computational Linguistics.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s Multilingual Neural Machine Translation System: Enabling Zero-Shot Translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Marcin Junczys-Dowmunt. 2018. [Dual conditional cross-entropy filtering of noisy parallel corpora](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 888–895, Belgium, Brussels. Association for Computational Linguistics.
- Jenna Kanerva, Filip Ginter, Li-Hsin Chang, Iiro Rastias, Valteri Skantsi, Jemina Kilpeläinen, Hanna-Mari Kupari, Jenna Saarni, Maija Sevón, and Otto Tarkka. 2021. [Finnish paraphrase corpus](#). In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 288–298, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Venelin Kovatchev, M. Antònia Martí, and Maria Salamó. 2018. [ETPC - a paraphrase identification corpus annotated with extended paraphrase typology and negation](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- 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.
- Nitika Mathur, Timothy Baldwin, and Trevor Cohn. 2019. [Putting evaluation in context: Contextual embeddings improve machine translation evaluation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2799–2808, Florence, Italy. Association for Computational Linguistics.

- Rada Mihalcea, Courtney Corley, and Carlo Strapparava. 2006. [Corpus-based and knowledge-based measures of text semantic similarity](#). In *AAAI*, pages 775–780.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Lidia Pivovarov, Ekaterina Pronoza, Elena Yagunova, and Anton Pronoza. 2018. [Paraphraser: Russian paraphrase corpus and shared task](#). In *Artificial Intelligence and Natural Language*, pages 211–225, Cham. Springer International Publishing.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4512–4525, Online. Association for Computational Linguistics.
- Holger Schwenk and Matthijs Douze. 2017. [Learning joint multilingual sentence representations with neural machine translation](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 157–167, Vancouver, Canada. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- 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.
- Jörg Tiedemann and Yves Scherrer. 2019. [Measuring semantic abstraction of multilingual NMT with paraphrase recognition and generation tasks](#). In *Proceedings of the 3rd Workshop on Evaluating Vector Space Representations for NLP*, pages 35–42, Minneapolis, USA. Association for Computational Linguistics.
- John Wieting and Kevin Gimpel. 2018. [ParaNMT-50M: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Melbourne, Australia. Association for Computational Linguistics.
- John Wieting, Kevin Gimpel, Graham Neubig, and Taylor Berg-Kirkpatrick. 2019. [Simple and effective paraphrastic similarity from parallel translations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4602–4608, Florence, Italy. Association for Computational Linguistics.
- John Wieting, Jonathan Mallinson, and Kevin Gimpel. 2017. [Learning paraphrastic sentence embeddings from back-translated bitext](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 274–285, Copenhagen, Denmark. Association for Computational Linguistics.
- 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.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [BERTScore: Evaluating text generation with BERT](#). In *International Conference on Learning Representations*.
- 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.

A Inference Time

Similarity measure	<i>ms</i> per pair
CHRF	0.6
SENTBLEU	0.4
Sentence-BERT	30.7
BERTSCORE-F1	4.5
NMTSCORE-direct	22.4
– without normalization	12.3
NMTSCORE-pivot	147.8
– without normalization	75.2
NMTSCORE-cross	75.0
– without normalization	75.0

Table A1: Inference time of the measures analyzed in this paper, averaged across the sentence pairs in the MRPC validation set. We measure the average time needed to compute a measure on a sentence pair in the MRPC validation set on a RTX 2080 Ti GPU. We use a batch size of 32 and compute the measures in both directions whenever this is required to make the measure symmetrical.

B Description of Models

Name	N	d_{model}	d_{ffn}	h	Param.	Vocab.	Lang.	License	URL
paraphrase-xlm-r-multilingual-v1	12	768	3072	12	278M	250k	50	Apache 2.0	↗
XLM-Roberta-large (up to layer 17)	17	1024	4096	16	472M	250k	100	MIT	↗
Prism	16	1280	12288	20	745M	64k	39	MIT	↗
m2m100_418M	24	1024	4096	16	484M	128k	100	MIT	↗
m2m100_1.2B	48	1024	8192	16	1239M	128k	100	MIT	↗

Table A2: Hyperparameters of the Transformer models used in this paper, as well as number of parameters, vocabulary size and number of supported languages.

C Metric Version Signatures

CHRF: nrefs:1|case:mixed|eff:yes|nc:6|nw:0|space:no|version:2.0.0

SENTBLEU: nrefs:1|case:mixed|eff:yes|tok:13a|smooth:exp|version:2.0.0

– JA: nrefs:1|case:mixed|eff:yes|tok:ja-mecab-0.996-IPA|smooth:exp|version:2.0.0

– ZH: nrefs:1|case:mixed|eff:yes|tok:zh|smooth:exp|version:2.0.0

BERTSCORE: xlm-roberta-large_L17_no-idf_version=0.3.11(hug_trans=4.17.0)

NMTSCORE:

NMTScore-direct|model:prism|normalized|both-directions|v0.2.0|hf4.17.0

NMTScore-pivot|pivot-lang:en|model:prism|normalized|both-directions|v0.2.0|hf4.17.0

NMTScore-cross|tgt-lang:en|model:prism|normalized|both-directions|v0.2.0|hf4.17.0

D Dataset Statistics

	Split	Positive pairs	Negative pairs	Avg. chars	License	Domains	URL
EN	Validation	239	128	109	unspecified	news	↗
	Test	1002	566	107			
RU	Test	1152	772	60	MIT License	news	↗
FI	Test	15368	5574	73	CC-BY-SA 4.0	subtitles, news	↗
SV	Test	783	298	45	CC-BY-SA 4.0	subtitles	↗
DE	Validation	831	1101	119	public domain	wikipedia	↗
	Test	895	1073	121			
ES	Validation	847	1115	117			
	Test	907	1092	118			
FR	Validation	860	1132	120			
	Test	903	1083	121			
JA	Validation	854	1126	58			
	Test	883	1063	60			
ZH	Validation	853	1131	43			
	Test	894	1081	44			

Table A3: Dataset statistics for our multilingual paraphrase identification benchmark.

	Validation		Test	
	Positive pairs	Negative pairs	Positive pairs	Negative pairs
EN+DE	1662	2202	1790	2146
EN+ES	1694	2230	1814	2184
EN+FR	1720	2264	1806	2166
EN+JA	1708	2252	1766	2126
EN+ZH	1706	2262	1788	2162
DE+ES	1640	2168	1790	2146
DE+FR	1658	2194	1788	2132
DE+JA	1646	2184	1748	2094
DE+ZH	1648	2190	1772	2126
ES+FR	1688	2220	1806	2166
ES+JA	1678	2208	1766	2124
ES+ZH	1674	2218	1788	2160
FR+JA	1702	2242	1764	2114
FR+ZH	1702	2252	1784	2142
JA+ZH	1688	2240	1744	2104

Table A4: Dataset statistics for the cross-lingual PAWS-X benchmark.

Language	Documents	Systems	Samples	Avg. references	Avg. reference characters
EN	178	16	2848	2.9	132
RU	110	7	770	2.5	123

Table A5: Statistics for the WebNLG 2020 RDF-to-text dataset of human judgments.

Language	DA	EL	ES	FI	IT	NL	PT	SV	BG	CS	ET	HU	LV	PL	RO
Samples	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50
Avg. ref. chars	56	61	60	56	62	61	59	52	130	129	42	138	129	151	130

Table A6: Statistics for the multilingual AMR-to-text dataset of human judgments.

E Other NMT Models

Language Metric	Individual datasets				PAWS-X dataset						Macro- average
	EN Acc.	RU AUC	FI AUC	SV AUC	DE Acc.	ES Acc.	FR Acc.	JA Acc.	ZH Acc.	Avg.	
m2m100_418M											
NMTSCORE-direct	72.0	83.2	71.1	71.1	71.1	69.0	72.3	61.9	65.4	67.9	73.1
NMTSCORE-pivot	72.4	84.2	68.2	70.3	73.2	72.0	72.2	64.0	67.9	69.9	73.0
NMTSCORE-cross	72.1	85.1	69.7	71.5	71.6	72.6	72.2	63.1	66.7	69.2	73.5
m2m100_1.2B											
NMTSCORE-direct	72.9	84.0	71.4	71.2	73.0	70.2	72.4	62.4	66.4	68.9	73.7
NMTSCORE-pivot	74.1	84.5	69.1	69.6	75.1	73.0	73.3	65.8	70.2	71.5	73.8
NMTSCORE-cross	72.8	85.0	70.0	71.0	74.1	73.0	73.3	66.2	69.5	71.2	74.0

Table A7: Comparison of translation-based text similarity measures when using two other multilingual NMT models (*M2M-100*; Fan et al., 2021). Overall, the accuracy of all three measures is slightly lower compared to the *Prism* NMT model but still competitive compared to the embedding baselines.

	EN + DE	EN + ES	EN + FR	EN + JA	EN + ZH	DE + ES	DE + FR	DE + JA	DE + ZH	ES + FR	ES + JA	ES + ZH	FR + JA	FR + ZH	JA + ZH	Avg.
m2m100_418M																
NMTSCORE-direct	72.5	70.9	72.2	63.3	65.0	67.9	69.0	61.0	63.1	68.4	60.6	62.0	61.4	63.2	60.6	65.4
NMTSCORE-pivot	73.8	73.1	73.6	63.9	65.4	71.9	70.5	63.0	63.8	70.1	62.8	64.3	62.7	63.5	62.0	67.0
NMTSCORE-cross	72.8	72.3	72.2	61.9	62.3	69.7	69.0	60.2	63.2	69.8	61.5	61.5	60.5	61.7	61.9	65.4
m2m100_1.2B																
NMTSCORE-direct	75.0	72.4	73.0	64.8	67.0	71.4	71.7	62.7	65.1	69.8	61.5	63.4	62.7	65.1	62.6	67.2
NMTSCORE-pivot	75.9	74.0	74.4	66.4	67.5	72.3	72.4	64.4	66.5	70.9	63.4	65.6	64.1	64.9	63.4	68.4
NMTSCORE-cross	74.8	74.1	73.8	63.0	63.6	70.9	70.5	61.4	63.5	71.3	61.6	62.5	61.8	63.8	63.3	66.7

Table A8: Comparison of translation-based text similarity measures on the cross-lingual PAWS-X dataset, using two other multilingual NMT models (*M2M-100*; Fan et al., 2021). Again, the average accuracy is lower compared to the *Prism* NMT model that we used for the main experiments, but superior to the baselines.

F MRPC Examples

Sentence Pair	Gold	SBERT	NMTSCORE
<i>The Dow Jones Industrial Average fell 0.7 per cent to 9,547.43 while the S&P 500 was 0.8 per cent weaker at 1,025.79.</i> <i>The Dow Jones industrial average fell 44 points, or 0.46 per cent, to 9,568.</i>	0	0.93	0.21
<i>So far, they have searched Pennsylvania, Ohio, Michigan, Illinois and Indiana, authorities in those state said.</i> <i>So far, authorities also have searched areas in Pennsylvania, Ohio, Indiana, and Michigan.</i>	0	0.81	0.08
<i>MEN who drink tea, particularly green tea, can greatly reduce their risk of prostate cancer, a landmark WA study has found.</i> <i>DRINKING green tea can dramatically reduce the risk of men contracting prostate cancer, a study by Australian researchers has discovered.</i>	1	0.91	0.14
<i>Bashir felt he was being tried by opinion not on the facts, Mahendradatta told Reuters.</i> <i>Bashir also felt he was being tried by opinion rather than facts of law, he added.</i>	1	0.87	0.18

Table A9: MRPC examples with a high disagreement between Sentence-BERT cosine similarity and NMTSCORE-cross.

Sentence Pair	Gold	BERTSCORE	NMTSCORE
<i>Batters faced: Sheets 28, Vizcaino 2, DeJean 4, Clement 26, Alfonseca 4, Guthrie 2, Farnsworth 4.</i> <i>Batters faced: Franklin 25, Kieschnick 7, Foster 2, Leskanic 3, DeJean 4, Prior 28, Alfonseca 2, Guthrie 2, Cruz 7, Remlinger 6.</i>	0	0.50	0.07
<i>But the technology-laced Nasdaq Composite Index was up 5.91 points, or 0.35 percent, at 1,674.35.</i> <i>The broader Standard & Poor's 500 Index .SPX was off 1.07 points, or 0.11 percent, at 1,010.59.</i>	0	0.45	0.07
<i>They also found shortness was associated with a family history of hearing loss.</i> <i>Shortness was found twice as often in those with hearing loss.</i>	1	0.46	0.09
<i>Kollar-Kotelly has scheduled another antitrust settlement compliance hearing for January.</i> <i>The judge scheduled another oversight hearing for late January.</i>	1	0.41	0.08

Table A10: MRPC examples with a high disagreement between rescaled BERTSCORE-F1 and NMTSCORE-cross.

G Cross-lingual PAWS-X Examples

Sentence Pair	Gold	SBERT	NMTSCORE
EN: <i>Write once , run anywhere</i> FR: <i>Écrivez n'importe où, une fois exécuté</i>	0	0.76	0.21
EN: <i>Worcester is a town and county city of Worcestershire in England .</i> DE: <i>Worcestershire ist eine Stadt und Kreisstadt von Worcester, England.</i>	0	0.93	0.37
EN: <i>The Jiul de Vest River is a tributary of the Jidanul River in Romania .</i> DE: <i>Der Jidanul ist ein Nebenfluss des Jiul de Vest, Rumänien.</i>	0	0.78	0.30
EN: <i>The Cugir River is a tributary of the Ghişag River in Romania .</i> DE: <i>Der Fluss Cugir ist ein Nebenfluss des Ghiaag in Rumänien.</i>	1	0.89	0.48
EN: <i>Film stars Lily Rabe , Timothée Chalamet , Lili Reinhart , Anthony Quintal , Oscar Nunez and Rob Huebel .</i> FR: <i>Le film met en vedette Oscar Nunez, Rob Huebel, Timothée Chalamet, Lily Rabe, Anthony Quintal et Lili Reinhart.</i>	1	0.80	0.49

Table A11: Cross-lingual PAWS-X examples with a high disagreement between Sentence-BERT cosine similarity and NMTSCORE-pivot.