

UMBCLU at SemEval-2024 Task 1: Semantic Textual Relatedness with and without machine translation

Shubhashis Roy Dipta[†] and Sai Vallurupalli[†]

{sroydip1,kolli}@umbc.edu

Department of Computer Science and Electrical Engineering
University of Maryland, Baltimore County
Baltimore, MD 21250 USA

Abstract

The aim of SemEval-2024 Task 1, “Semantic Textual Relatedness for African and Asian Languages” is to develop models for identifying semantic textual relatedness (STR) between two sentences using multiple languages (14 African and Asian languages) and settings (supervised, unsupervised, and cross-lingual). Large language models (LLMs) have shown impressive performance on several natural language understanding tasks such as multilingual machine translation (MMT), semantic similarity (STS), and encoding sentence embeddings. Using a combination of LLMs that perform well on these tasks, we developed two STR models, *TranSem* and *FineSem*, for the supervised and cross-lingual settings. We explore the effectiveness of several training methods and the usefulness of machine translation. We find that direct fine-tuning on the task is comparable to using sentence embeddings and translating to English leads to better performance for some languages. In the supervised setting, our model performance is better than the official baseline for 3 languages with the remaining 4 performing on par. In the cross-lingual setting, our model performance is better than the baseline for 3 languages (leading to 1st place for Africaans and 2nd place for Indonesian), is on par for 2 languages and performs poorly on the remaining 7 languages.

1 Introduction

The objective of the SemEval 2024 Task 1 is to build and evaluate models capable of identifying relatedness between a sentence pair. Sentence pairs from 14 African and Asian languages belonging to 5 language groups are annotated for relatedness and released for model development. The task is divided into 3 tracks targeting different types of model training: supervised (Track A), unsupervised (Track B), and cross-lingual (Track C). Each

track targets a different subset of languages. Extensive details about the languages, language groups, and the data collection process are provided in the task description paper (Ousidhoum et al., 2024a). A detailed description of the shared task, tracks, and datasets are provided in the shared task description paper (Ousidhoum et al., 2024b).

Semantic relatedness helps with understanding language meaning (Jarmasz and Szpakowicz, 2012; Miller, 1995; Antony et al., 2022; Osgood, 1949) and is useful in many areas of natural language processing such as word-sense disambiguation (Banerjee and Pedersen, 2003), machine translation (Bracken et al., 2017) and sentence representation (Reimers et al., 2019; Wang et al., 2022) which have numerous applications. Until recently, semantic relatedness has been mostly restricted to finding word relatedness (Feng et al., 2017; Budanitsky and Hirst, 2006), leading to a lack of sentence-relatedness datasets. At the sentence level, relatedness has been limited to similarity, providing a restricted view of STR (Abdalla et al., 2021). The current shared task aims to broaden the scope of sentence relatedness and extend it to several languages with the goal of encouraging model and resource development in these languages (Ousidhoum et al., 2024b).

Recent advancements in multi-lingual translation and the availability of models for obtaining high-quality sentence embeddings allowed us to explore the effectiveness of machine translated data. Using various sentence embedding models to encode data translated into English, we trained a model, *TranSem*, to find the relatedness score between sentence pairs. Although the task requires the sentence pair to be from the same language, our model can handle sentences from two different languages. Our second model, *FineSem* directly fine-tunes a T5 model (T5 is already fine-tuned on the STS benchmark) on the STR task using both untranslated and translated data to explore the

[†]These authors contributed equally to this work

usefulness of translation. We use both these models to evaluate languages in Track A. For Track C languages, we use a T5 model fine-tuned only on the english STR data. For evaluating the English dataset in the cross-lingual track, we use a T5 model fine-tuned on the Spanish dataset.

Our contributions to the STR task are as follows: We 1) develop unified models for STR to work with all languages. 2) participate in supervised and cross-lingual tracks. 3) explore the usefulness of machine translation. 4) explore data augmentation using machine translation. Our code is publicly available at <https://github.com/dipta007/SemEval24-Task8>

2 System Overview

After exploring models and datasets available in the languages we understand¹, we realized the dearth of resources available in these languages. To leverage resources available in English, we translated the 13 non-English languages into English. Assuming the translated data accurately reflects the semantic meaning of the source language, the derived relatedness value from our model for a translated sentence pair should reflect the STR between the sentence pair in the source language. This section describes our machine translation process and models, *TranSem* and *FineSem*. Besides using different training strategies, these models can use both translated and untranslated data.

2.1 Translation to English & Data Augmentation

We use Meta’s “No Language Left Behind (NLLB) open-source models” that provide direct high-quality translations for 200 languages with many low-resource languages (Costa-jussà et al., 2022). We use four of the translation models² ranging from 600 million to 3.3 billion parameters. Using each of the 4 models, we translated the training data for all languages in track A, except Amharic and Algerian Arabic, and obtained 4 translated datasets for each language. None of the 4 models we used supported Amharic, Algerian Arabic or Punjabi. We decided against translating track C languages with 3 of 12 unsupported languages. Using 4 different model translations gave us a 4-fold augmentation

¹Of the 14 languages, the authors are proficient in English, Hindi & Telugu languages.

²[facebook/nllb-200-3.3B](https://github.com/facebook/nllb-200-3.3B), [facebook/nllb-200-1.3B](https://github.com/facebook/nllb-200-1.3B), [facebook/nllb-200-distilled-1.3B](https://github.com/facebook/nllb-200-distilled-1.3B), [facebook/nllb-200-distilled-600M](https://github.com/facebook/nllb-200-distilled-600M)

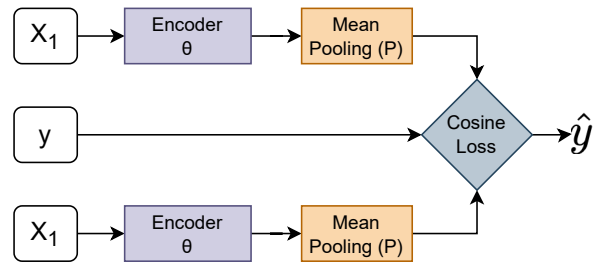


Figure 1: Overview of TranSem model architecture (Inspired by Reimers et al. (2019)). The encoder (θ) is shared, and the diamond box represents the loss function. The encoded sentence pairs (x_1, x_2) and the label (y) are the input to the cosine similarity loss.

of the training data. We translated the test data using only the largest model (*facebook/all-200-3.3 B*) to obtain the best translation features.

2.2 TranSem Model

Inspired by Reimers et al. (2019), we used the Siamese model architecture (shown in Fig. 1). For a given pair of sentences (x_1, x_2) and their semantic relatedness score (y), we encode each sentence with a sentence encoder (θ). The embeddings for the sentences go through a pooling operation (P) to produce sentence embeddings (s_1, s_2). The cosine similarity of the encoded embeddings is trained to match the semantic relatedness score using the mean-squared error loss:

$$\mathcal{L} = MSE(\cos\text{-sim}(P\theta(x_1), P\theta(x_2)), y) \quad (1)$$

We experimented with several sentence encoding models for encoding our translated and augmented training dataset. We chose DistilRoberta³ to submit results for the competition leaderboard based on our primary validation (details on §4.1.1, §4.1.2, and §4.1.3). This is a distilled version of RoBERTa (Liu et al., 2019) fine-tuned on sentence-level datasets and suitable for clustering and semantic searches, which we further fine-tuned on our translated and augmented dataset. The sentence-t5-xl embedding model was chosen to compare the effectiveness of sentence embeddings as opposed to the direct fine-tuning used in the *FineSem* model. After experimenting with different pooling mechanisms of mean, max, and CLS tokens, we found that mean pooling works well for our setting. This aligns with earlier findings, which show that mean pooling produced encodings lead to better performance on downstream tasks.

³[sentence-transformers/all-distilroberta-v1](https://github.com/sentence-transformers/all-distilroberta-v1)

2.3 FineSem Model

T5 (Raffel et al., 2020) is a transformer model that uses transfer learning; the model trained on “Colossal Clean Crawled Corpus” is fine-tuned on a mixture of 8 downstream unsupervised and supervised tasks converting them into a unified text-to-text task setting. The T5 model is available in several sizes, of which we use the base, large, and XL models ranging from (660 million to 3 billion parameters). One of the supervised tasks used in the T5 model training is the semantic textual similarity benchmark (STS-B) dataset trained as a regression classification problem. We use the STS task setting to train on the track A STR training datasets using 3 different options: separate T5 models trained on individual languages, a single model trained on all 14 languages (without translation), and a single model trained on the translated and augmented dataset (for 12 languages). These settings allow us to contrast the effectiveness of direct fine-tuning with the sentence embedding-based *transem* model and the usefulness of machine translation.

From the T5 models fine-tuned on the individual languages, we use the English and Spanish fine-tuned models (we refer to these models as the English and Spanish models) for evaluating the cross-lingual Track C languages. We use the English model to evaluate development and test data from all languages except English and the Spanish model to evaluate the English data.

3 Experimental Setup

3.1 HyperParameters

We train our models using AdamW (Loshchilov and Hutter, 2017) optimizer with a learning rate of $1e-5$ and weight decay of 0.01. We use an effective batch size of 32 (batch size 16 with 2 steps of gradient accumulation) (for TransSem) and a batch size of 16 (for FineSem).

We train our *transem* model infinitely with an early stopping patience of 10 on the validation Spearman Correlation score. We train the *fine-sem* model for 10 epochs (2 epochs for the model trained on the translated and augmented model) and checkpoint the models at the end of every epoch. We evaluate the dev sets for each language against these 10 checkpoints. We evaluate the corresponding test data using the checkpoint which provides the best performance on the dev data for a language.

3.2 Infrastructure

All experiments were conducted on an NVIDIA Quadro RTX 8000 with 48GB of VRAM and A100 80GB. We utilize the PyTorch Lightning library⁴ for conducting the experiments and Weight & Biases⁵ for logging purposes (for the *TransSem* Model) and HuggingFace Transformers (for the *FineSem* Model).

4 Results

We first report our results and analysis on Track A languages (§4.1), and then on Track C (§4.2). We use the official baseline (Ousidhoum et al. (2024b)) that used LaBSE (Feng et al., 2020) fine-tuned on the provided training dataset and refer to this model as *baseline*.

4.1 Track A Languages

This section discusses our findings using various model settings with the *TranSem* model.

4.1.1 Effect of Batch Size

Comparing performance with various batch sizes (results are shown in Table 1), we show that our batch size selections are fairly good (32 for task *TranSem* and 8 for *FineSem*).

Batch Size	A3 eng	A4 hau	A5 kin	A6 mar	A7 ary	A8 esp	A9 tel	avg	
TranSem+	2	.8102	.6857	.6886	.8515	.7376	.6519	.8342	.7514
TranSem+	4	.5415	.1355	.0643	.5394	.3038	.5998	.4231	.3725
TranSem+	8	.8132	.6519	.6642	.8348	.7340	.6355	.8187	.7360
TranSem+	16	.8093	.6377	.6997	.8429	.7500	.6462	.8257	.7445
TranSem+	64	.8092	.6589	.6456	.8271	.7247	.6234	.8131	.7289
TranSem+	128	.8129	.6787	.6659	.8417	.7411	.6396	.8324	.7446
TranSem+	256	.8152	.6716	.6589	.8357	.7197	.6353	.8189	.7365

Table 1: The effect of batch size on *TranSem* for different batch sizes {2, 4, 8, 16, 64, 128, 256}

4.1.2 Effect of Encoder Pooling

In Table 2, we compare performance using 3 different pooling operations. We used mean pooling for the official results we submitted, as it showed good performance in most languages.

Pool	A3 eng	A4 hau	A5 kin	A6 mar	A7 ary	A8 esp	A9 tel	avg	
TranSem+	CLS	.8133	.6737	.6655	.8339	.7376	.6363	.8309	.7416
TranSem	Mean	.8125	.6403	.6807	.8406	.7448	.7211	.8255	.7522
TranSem+	Max	.7960	.6157	.5809	.8227	.6643	.6075	.7997	.6981

Table 2: The effect of pooling on *TranSem* using different pooling mechanisms (CLS Token, Mean, Max)

⁴<https://lightning.ai/>

⁵<https://wandb.ai/>

4.1.3 Effect of Sentence Embedding Models

In Table 3, we provide contrastive results with several sentence embedding models used in *TranSem*. For official results, we submitted results from the distilroberta-v1 sentence embedding model (results for some of the languages are from the *FineSem* model fine-tuned on individual STR training datasets).

4.1.4 Usefulness of Machine Translation and Direct Fine-tuning

We compare the performance of the *FineSem* models fine-tuned using the 3 data options (results are shown in Table 3). *FineSem*-Individual shows the performance of T5-XL models fine-tuned on the individual datasets. Unified and Translated models show the performance of the two T5-XL models fine-tuned on the untranslated data and translated+augmented data. The model trained on untranslated data performs poorly on the Marathi dataset, but performs on par with the other models indicating that we may not need to translate all languages to English. We find that direct fine-tuning with the translated and augmented data is comparable with the *TranSem* model using various sentence embeddings.

4.2 Track C Languages

In Table 4, we compare the performance of various T5 models on the track C languages. We submitted official results using our T5-XL based *FineSem* model (*FineSem*-LB) where the results are obtained using the checkpoint after the third epoch. With the same model we also report with the approach where we use the checkpoint which results in the best performance on the development data for a given language. These results are shown as *FineSem*-XL. We compare these results with the T5-base and T5-large based *FineSem* models. We bold the best scores for easy readability but underline scores that are better than the baseline. Among our models the overall performance of the XL model is better and this model improves upon the baseline for Afrikaans, Indonesian and Spanish.

5 Related Work

In this section, we present previous research conducted in the fields of Machine Translation (see §5.1), Sentence Embedding (see §5.2) and Semantic Similarity (see §5.3).

5.1 Machine Translation

Machine translation has evolved in the last 75 years from rule-based systems to statistical-based systems to the current neural machine translation (NMT) systems. In the 10 years since the first sequence-to-sequence NMT model (Bahdanau et al., 2014), machine translation reached a point where translations from models for high-resource languages rival human translators (Läubli et al., 2018; Popel et al., 2020). This was possible due to the amount of bilingual data pairs available for training in these languages (Haddow et al., 2022). Translation systems for medium and low resource languages that lacked the scale of these resources either developed cross-lingual models (Nguyen and Chiang, 2017; Zoph et al., 2016) or developed datasets (Bañón et al., 2020; Schwenk et al., 2019). Current state-of-the-art translation models use a many-to-many approach to handle a large number of medium to low-resource languages (Costa-jussà et al., 2022).

5.2 Sentence Embedding

Generating a sentence-level embedding is useful for semantic searches and clustering. Since the first transformer model (Vaswani et al., 2017), several encoder-only models such as BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019), XLNet (Yang et al., 2018) were used to learn effective sentence embeddings that also performed well on downstream NLP tasks (Cer et al., 2018; Roy Dipta et al., 2023). Reimers et al. (2019) developed a Siamese-like network architecture with two BERT sentence embedding models that improved semantic search systems. Using T5 (Raffel et al., 2020) models in a similar architecture, sentence embeddings produced by T5 were shown to be superior to the encoder-only model embeddings with performance gains in downstream tasks (Ni et al., 2021). A more recent work (Reimers and Gurevych, 2020) showed that a teacher-student model can be efficiently used to develop a sentence embedding system for many low-resource languages.

5.3 Semantic Similarity

Semantic textual similarity captures a type of semantic relatedness requiring similarity on all aspects between a sentence pair. SemEval tasks on semantic textual similarity from 2012 to 2017 resulted in the STS benchmark (Cer et al., 2017). Recently, Deshpande et al. (2023) proposed condi-

Model	A3 eng	A4 hau	A5 kin	A6 mar	A7 ary	A8 esp	A9 tel	avg
baseline	.8300	.6900	.7200	.8800	.7700	.7000	.8200	.7729
distilroberta-v1 (TranSem)	.8125	.6403	.6807	.8406	.7448	.7211	.8255	.7522
mpnet-base-v2	.8104	.6692	.6971	.8568	.7297	.6518	.8250	.7486
roberta-large-v1	.8260	.6750	.7056	.8461	.7480	.6298	.8394	.7528
sentence-t5-xl	.8236	.6440	.6720	.8324	.7124	.6277	.8250	.7339
multi-qa-mpnet-base-dot-v1	.8111	.6852	.7041	.8482	.7163	.6586	.8245	.7497
all-MiniLM-L12-v2	.8237	.6474	.7026	.8522	.7605	.6141	.8247	.7465
FineSem-Individual	.8385	.6335	.7175	.2211	.7647	.6900	.6085	.6391
FineSem-Unified	.8438	.6369	.6837	.3878	.6265	.7040	.6993	.6546
FineSem-Translated	.8105	.6383	.7133	.8608	.7403	.6663	.8152	.7493

Table 3: Model Performance (Spearman Correlation Coefficient) on Subtask A test set. *TranSem* shows results submitted before the official deadline, *baseline* shows official baseline results, and the rest are contrastive results for our various models. The best scores within each section are **bolded**, and best scores across all sections are **underlined**.

Model	C1 Afr	C2 Arq	C3 Amh	C4 Eng	C5 Hau	C6 Hin	C7 Ind	C8 Kin	C9 Arb	C10 Ary	C11 Pan	C12 Esp	avg
Baseline	.7900	.4600	.8400	.8000	.6200	.7600	.4700	.5700	.6100	.4000	-.050	.6200	.5742
FineSem-LB	.8223	.1263	.0430	.7875	.4569	.1552	.5153	.4836	.0354	-.0375	-.0775	.6089	.3266
FineSem-XL	<u>.8164</u>	.1023	.0373	.7889	.4561	.1594	.5279	.4128	.0000	.0219	-.0817	<u>.6259</u>	.3246
FineSem-L	.8007	-.0515	.0112	.7752	.4831	.1764	.4419	.5094	.0154	.0331	-.0591	.6605	.3164
FineSem-B	.7802	.1799	.2543	.7448	.4784	.2404	.4517	.3861	.0527	.0268	-.0520	.6289	.3477

Table 4: Model performance (Spearman Correlation Coefficient) on Track C test sets. All language test sets (except English) use the FineSem models trained on the English training set. The English test set uses the FineSem models trained on the Spanish training set. The best scores among our models are **bolded**. Scores better than baseline are **Underlined**.

tional semantic textual similarity to explore semantic relatedness.

6 Conclusion & Future Work

We developed two different models and showed how the models performed in supervised and cross-domain training tasks in 14 languages. We explored using machine translation, sentence encoders, and SST-B style training with T5 models. Our models improved over the official baseline for some of the languages. For computational purposes, we have excluded using more recent models like mistral-7b (Jiang et al., 2023), which have outperformed most of the open-source and close-source models in various benchmarks (Zheng et al., 2024). For future work, we intend to explore prompting for STR and prompt-based LLMs⁶ for translation.

7 Disclaimer

We did not use AI assistants to write any part of our paper or code. All writing is original and produced

⁶<https://chat.openai.com/>

by the authors.

8 Limitations

We acknowledge our work has the following limitations. We use several pre-trained LLMs in our experiments. It is well known that these models can echo biases and misinformation either implicitly or explicitly. We did not control for any of these when training them on the STR datasets. In addition, the STR datasets may also echo several biases related to social groups, cultural groups, race, gender, behavioral, and perceptual differences of annotators. We did not explore or control for any of these biases in our work. As a result, our work carries the limitations of both the models and the datasets we used.

References

Mohamed Abdalla, Krishnapriya Vishnubhotla, and Saif M. Mohammad. 2021. [What makes sentences semantically related: A textual relatedness dataset and empirical study](#). *CoRR*, abs/2110.04845.

- James W Antony, America Romero, Anthony H Vierra, Rebecca S Luenser, Robert D Hawkins, and Kelly A Bennion. 2022. [Semantic relatedness retroactively boosts memory and promotes memory interdependence across episodes](#). *eLife*, 11:e72519.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Satanjeev Banerjee and Ted Pedersen. 2003. Extended gloss overlaps as a measure of semantic relatedness. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence, IJCAI'03*, page 805–810, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Marta Bañón, Pinzhen Chen, Barry Haddow, Kenneth Heafield, Hieu Hoang, Miquel Esplà-Gomis, Mikel Forcada, Amir Kamran, Faheem Kirefu, Philipp Koehn, et al. 2020. Paracrawl: Web-scale acquisition of parallel corpora. Association for Computational Linguistics (ACL).
- Jennifer Bracken, Tamar Degani, and Natasha Tokowicz Chelsea Eddington. 2017. [Translation semantic variability: How semantic relatedness affects learning of translation-ambiguous words](#). *Bilingualism: Language and Cognition*, 20(4):783–794.
- 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.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Marta R Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*.
- Ameet Deshpande, Carlos E. Jimenez, Howard Chen, Vishvak Murahari, Victoria Graf, Tanmay Rajpurohit, Ashwin Kalyan, Danqi Chen, and Karthik Narasimhan. 2023. [C-sts: Conditional semantic textual similarity](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. Language-agnostic bert sentence embedding. *arXiv preprint arXiv:2007.01852*.
- Yue Feng, Ebrahim Bagheri, Faezeh Ensan, and Jelena Jovanovic. 2017. [The state of the art in semantic relatedness: a framework for comparison](#). *The Knowledge Engineering Review*, 32:e10.
- Barry Haddow, Rachel Bawden, Antonio Valerio Miceli Barone, Jindřich Helcl, and Alexandra Birch. 2022. Survey of low-resource machine translation. *Computational Linguistics*, 48(3):673–732.
- Mario Jarmasz and Stanialaw Szpakowicz. 2012. [Roget’s thesaurus and semantic similarity](#). In *Recent Advances in Natural Language Processing*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Samuel Lüubli, Rico Sennrich, and Martin Volk. 2018. [Has machine translation achieved human parity? a case for document-level evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4791–4796, Brussels, Belgium. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- George A. Miller. 1995. [Wordnet: a lexical database for english](#). *Commun. ACM*, 38(11):39–41.
- Toan Q Nguyen and David Chiang. 2017. Transfer learning across low-resource, related languages for neural machine translation. *arXiv preprint arXiv:1708.09803*.
- Jianmo Ni, Gustavo Hernández Ábrego, Noah Constant, Ji Ma, Keith B. Hall, Daniel Cer, and Yinfei Yang. 2021. [Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models](#). *CoRR*, abs/2108.08877.
- C. E. Osgood. 1949. [The similarity paradox in human learning: a resolution](#). 56:132–143.
- Nedjma Ousidhoum, Shamsuddeen Hassan Muhammad, Mohamed Abdalla, Idris Abdulmumin, Ibrahim Said Ahmad, Sanchit Ahuja, Alham Fikri Aji, Vladimir Araujo, Abinew Ali Ayele, Pavan Baswani, Meriem Beloucif, Chris Biemann, Sofia Bourhim, Christine De Kock, Genet Shanko Dekebo, Oumaima

- Hourrane, Gopichand Kanumolu, Lokesh Madasu, Samuel Rutunda, Manish Shrivastava, Thamar Solorio, Nirmal Surange, Hailegnaw Getaneh Tilaye, Krishnapriya Vishnubhotla, Genta Winata, Seid Muhie Yimam, and Saif M. Mohammad. 2024a. [Semrel2024: A collection of semantic textual relatedness datasets for 14 languages.](#)
- Nedjma Ousidhoum, Shamsuddeen Hassan Muhammad, Mohamed Abdalla, Idris Abdulmumin, Ibrahim Said Ahmad, Sanchit Ahuja, Alham Fikri Aji, Vladimir Araujo, Meriem Beloucif, Christine De Kock, Oumaima Hourrane, Manish Shrivastava, Thamar Solorio, Nirmal Surange, Krishnapriya Vishnubhotla, Seid Muhie Yimam, and Saif M. Mohammad. 2024b. [SemEval-2024 task 1: Semantic textual relatedness for african and asian languages.](#) In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*. Association for Computational Linguistics.
- Martin Popel, Marketa Tomkova, Jakub Tomek, Łukasz Kaiser, Jakob Uszkoreit, Ondřej Bojar, and Zdeněk Žabokrtský. 2020. [Transforming machine translation: a deep learning system reaches news translation quality comparable to human professionals.](#) 11.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer.](#) *Journal of Machine Learning Research*, 21(140):1–67.
- 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*. Association for Computational Linguistics.
- Nils Reimers, Nils Reimers, Iryna Gurevych, and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks.](#) *arXiv: Computation and Language*.
- Shubhashis Roy Dipta, Mehdi Rezaee, and Francis Ferraro. 2023. [Semantically-informed hierarchical event modeling.](#) In *Proceedings of the 12th Joint Conference on Lexical and Computational Semantics (*SEM 2023)*, pages 353–369, Toronto, Canada. Association for Computational Linguistics.
- Holger Schwenk, Guillaume Wenzek, Sergey Edunov, Edouard Grave, and Armand Joulin. 2019. [Ccmatrix: Mining billions of high-quality parallel sentences on the web.](#) *arXiv preprint arXiv:1911.04944*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need.](#) *Advances in neural information processing systems*, 30.
- Bin Wang, C.-C. Jay Kuo, and Haizhou Li. 2022. [Just rank: Rethinking evaluation with word and sentence similarities.](#) In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6060–6077, Dublin, Ireland. Association for Computational Linguistics.
- Yinfei Yang, Steve Yuan, Daniel Cer, Sheng-yi Kong, Noah Constant, Petr Pilar, Heming Ge, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. [Learning semantic textual similarity from conversations.](#) In *Proceedings of the Third Workshop on Representation Learning for NLP*, pages 164–174, Melbourne, Australia. Association for Computational Linguistics.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. [Judging llm-as-a-judge with mt-bench and chatbot arena.](#) *Advances in Neural Information Processing Systems*, 36.
- Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. [Transfer learning for low-resource neural machine translation.](#) *arXiv preprint arXiv:1604.02201*.