

# CAILMD-23 at SemEval-2024 Task 1: Multilingual Evaluation of Semantic Textual Relatedness

Srushti Sonavane<sup>1</sup>, Sharvi Endait<sup>1</sup>, Ridhima Sinare<sup>1</sup>, Pritika Rohera<sup>1</sup>, Advait Naik<sup>1</sup>,  
and Dipali Kadam<sup>1</sup>

Pune Institute of Computer Technology, Pune<sup>1</sup>

## Abstract

The explosive growth of online content demands robust Natural Language Processing (NLP) techniques that can capture nuanced meanings and cultural context across diverse languages. Semantic Textual Relatedness (STR) goes beyond superficial word overlap, considering linguistic elements and non-linguistic factors like topic, sentiment, and perspective. Despite its pivotal role, prior NLP research has predominantly focused on English, limiting its applicability across languages. Addressing this gap, our paper dives into capturing deeper connections between sentences beyond simple word overlap. Going beyond English-centric NLP research, we explore STR in Marathi, Hindi, Spanish, and English, unlocking the potential for information retrieval, machine translation, and more. Leveraging the SemEval-2024 shared task, we explore various language models across three learning paradigms: supervised, unsupervised, and cross-lingual. Our comprehensive methodology gains promising results, demonstrating the effectiveness of our approach. This work aims to not only showcase our achievements but also inspire further research in multilingual STR, particularly for low-resourced languages. (Ousidhoum et al., 2024b)

Keywords: Natural Language Processing, Semantic Textual Relatedness, Sentence Transformers, supervised learning, unsupervised learning, cross-lingual.

## 1 Introduction

The ever-increasing diversity of online content demands robust Natural Language Processing (NLP) techniques that can grasp the nuances of meaning across diverse languages. Semantic Textual Relatedness (STR) plays a crucial role in achieving this goal by delving beyond superficial lexical similarity and capturing the deeper connections between sentences. Unlike semantic similarity, which focuses solely on the taxonomic overlap of words,

STR encompasses both linguistic elements and non-linguistic factors like the topic, point of view, and period. This richer understanding unlocks significant potential in various NLP tasks, regardless of the user's native language. Imagine searching for information online in your native language and receiving results that truly understand your intent, and not just match keywords. STR holds the key to unlocking this dream, bridging the language gap, and fostering true multilingual communication.

Despite the recognized importance of Semantic Textual Relatedness (STR) for multilingual communication, most prior NLP research has focused on semantic similarity within English due to limitations in labeled data for diverse languages (Abdalla et al., 2023). This narrow focus restricts the potential of STR applications like information retrieval across languages with different cultural contexts or machine translation that accurately captures nuances beyond direct word equivalents. Existing relatedness methods primarily target English (Hasan and Halliday, 1976), with limited exploration in languages like German, Chinese, and Japanese (Zesch et al., 2007) (Li et al., 2005) (De Saeger et al., 2010). This highlights a critical gap in Natural Language Processing (NLP): accurately measuring semantic relatedness across diverse languages.

The identified gap in multilingual STR research, with its limitations in diverse language applications, demands innovative solutions. This paper dives into the exciting realm of bridging this gap through multilingual Semantic Textual Relatedness (STR). Specifically, we explore methods to capture the semantic connections between texts in languages like English, Marathi, Hindi, and Spanish.

Our research focuses on the SemEval-2024 shared task, which provides three tracks to evaluate STR techniques:

1. **Supervised Learning:** This track focuses on building systems trained on the provided labeled datasets.

2. Unsupervised Learning: Here, the challenge lies in developing systems that learn semantic relationships without relying on any labeled data.
3. Cross-lingual Learning: This track pushes the boundaries by requiring systems to leverage knowledge from labeled data in a source language (Track A) to address a target language with limited resources.

For each track, we present a comprehensive methodology, employing diverse language models and rigorously analyzing their performance. This allows us to identify the most effective approaches for each challenge. Notably, our submissions achieved promising scores on several tracks, demonstrating the strength and potential of our proposed methods.

Looking beyond our achievements, this work aims to inspire further exploration of multilingual STR, particularly for under-resourced languages. We believe that larger datasets and broader language coverage hold immense potential to benefit the NLP community, unlocking the true potential of language understanding and empowering communication across diverse cultures.

## 2 Related Work

Semantic textual relations (STR) play an important role in natural language processing (NLP), which aims to identify the degree of semantic similarity between text groups. It forms the backbone of various NLP tasks such as information retrieval, question answering, and paraphrase detection, necessitating the assessment of similarity between sentences, phrases, or documents.

Historically, detailed STR research from the 1900s through the 2000s relied heavily on statistical methods heavily dependent on lexical databases like WordNet. However, these methods suffered from a lack of real-world knowledge integration (Gabrilovich et al., 2007). Classified translation emerged with developments such as GloVe (Pennington et al., 2014), Word2Vec, and FastText, which enabled text to be converted into word input. The current methods require converting corpora into words or sentence-embedded forms and computing connectivity scores. Notably, large language models (LLMs) such as Sentence-BERT often use Cosine Similarity in embedded sentences to measure relatedness (Gunawan et al., 2018) (Reimers

and Gurevych, 2019).

Previous methodologies have delved into both knowledge-based (ontology, classification) and corpus-based (unsupervised learning) approaches. For example, (Siblini and Kosseim, 2017) examined three approaches: semantic linkage, classification similarity, and hybrid approaches. Notably, the multilingual approach of (Hasan and Haliday, 1976), improved by 47%, confirming the potential of emphasis on the use of multilingual strategies. Furthermore, studies on less resourceful African languages highlight the need for different data types and methodologies (Delil and Kuyumcu, 2023).

A significant challenge in STR lies in the scarcity of huge-scale, promising datasets for education and assessment. Initiatives like SemEval play a pivotal role in addressing this gap through dedicated shared tasks focused on STR (Abdalla et al., 2023). These collaborative efforts foster the improvement and evaluation of STR models throughout diverse linguistic landscapes and domain names.

The current advent of the STR-2022 dataset by (Abdalla et al., 2023) marks a significant leap forward in STR studies. This annotated dataset, comprising sentence pairs with relatedness scores, serves as a precious aid for schooling and evaluating STR fashions. Covering various domains and languages, it displays the multilingual nature of STR studies (Abdalla et al., 2023).

Moreover, STR-2022 addresses biases and perceptions in relatedness judgments. Through meticulous curation, it aims to mitigate biases and ensure annotation quality, thereby fostering fair evaluations and robust model development. Additionally, the dataset highlights the relative nature of relatedness ratings, emphasizing the significance of context and assignment-precise thresholds in decoding similarity measures (Abdalla et al., 2023).

## 3 System Description

In this section, we aim to outline our system’s components for assessing semantic textual relatedness across different datasets: a) labeled datasets using supervised learning, b) unlabeled datasets employing unsupervised learning, and c) cross-lingual datasets. We’ll detail the utilized data, the models employed in each track, and the results obtained from training these models on respective datasets.

### 3.1 Data Collection

We utilized the SemRel2024 Dataset (Ousidhoum et al., 2024a) for training and evaluating our final results. This comprehensive dataset consists of semantic textual relatedness data across 14 diverse languages, including Afrikaans, Algerian Arabic, Amharic, English, Hausa, Hindi, Indonesian, Kinyarwanda, Marathi, Moroccan Arabic, Modern Standard Arabic, Punjabi, Spanish, and Telugu. From this array of languages, we focused on English, Hindi, Marathi, and Spanish datasets for our analysis.

Each entry in the dataset comprises a sentence pair along with its corresponding semantic similarity score. This score ranges from 0 to 1, where 0 signifies no similarity between the sentences, while 1 indicates complete similarity.

For the supervised track (Track A), we concentrated on English and Marathi datasets. In the unsupervised track (Track B), our attention was on English and Hindi datasets. Lastly, for Track C, we employed English and Hindi datasets, utilizing Spanish and English as their language training bases, respectively.

Language	Train	Dev	Test
English	5500	250	2500
Hindi	-	288	968
Marathi	1155	293	298
Spanish	1592	140	600

Table 1: Distribution of dataset for Training, Development, and Testing

### 3.2 Experiments

#### 3.2.1 Track A:

The SemRel2024 English and Hindi datasets were initially trained on baseline models such as Support Vector Regression and XGBoost. However, we additionally adapted the sentence-transformer-based models, such as all-mpnet-base-v2<sup>1</sup> and marathi-sentence-bert-nli<sup>2</sup> by L3Cube for English and Marathi respectively. This was done to compensate for the smaller size of the corpora available, as these sentence transformer models are trained on a larger data size initially, and this would be efficient to understand not only the n-gram sequences but also the context of the sentences that are being compared.

<sup>1</sup><https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

<sup>2</sup><https://huggingface.co/l3cube-pune/marathi-sentence-bert-nli>

Preprocessing and Feature vectorization were done using Term-frequency and inverse-document-frequency (TF-IDF) to generate vectors and preprocess the models SVR and XGBoost. Term-frequency, Inverse-Document-Frequency (TF-IDF) is a numerical statistic that determines how important a word is in a given document or a piece of textual content. This is done by multiplying two metrics: How many times a word appears in a document Inverse document frequency of the word across a set of documents. This score for word in the document  $d$  from document  $D$  is calculated as follows:

$$tfidf(t, d, D) = tf(t, d).idf(t, D) \quad (1)$$

$$tf(t, d) = \log(1 + freq(t, d)) \quad (2)$$

$$idf(t, D) = \log\left(\frac{N}{\text{count}(d \in D : t \in d)}\right) \quad (3)$$

(Chen et al., 2020).

Support vector regression (SVR) can be used for calculating semantic similarity scores between sentences. It's a supervised learning algorithm that can model the relationship between input features, (here in this case the sentences) and the output labels (semantic similarity scores in this case).

XGBoost can be used for semantic similarity score calculation between sentences as it is a powerful gradient-boosting algorithm, and it can be applied to various supervised learning tasks. This is capable of handling complex non-linear relationships between features and labels, and it's robust against overfitting.

All-mpnet-base-v2 (All- Massively Parallel Multilingual Transformer) is a sentence encoder model, given an input text, it gives a vector that collects the semantic information. The sentence vector is used for tasks such as clustering or sentence similarity tasks. This is a sentence-transformers model and it maps sentences & paragraphs to a 768-dimensional dense vector space. This is trained on the SemRel2024 dataset and additionally, it is capable of capturing long-range dependencies, and it leads to higher performance on text classification, NER, and question answering.

Marathi-Sentence-Bert-Nli (Joshi et al., 2022) is a Marathi sentence transformer model that has been trained on synthetic STS and NLI datasets. These are fine-tuned on MahaBERT, a BERT-based model that is fine-tuned on a large Marathi corpora.

### 3.2.2 Track B:

For track B, the unsupervised track, the SemRel2024 dev set, and the test set were used for testing the model selected. The languages were English and Hindi, for which Track-B dev and test sets were used. The models used for this were BERT-based uncased and Hindi-Bert v2 (Joshi, 2022) accordingly.

Hindi-BERT-v2 was roughly trained on 1.8 B tokens. Compared to general-purpose language models, this monolingual model is optimized to understand and process Hindi text effectively. Due to the larger corpus it has been trained upon this has been an accurate model to obtain results from.

BERT-based-uncased (Bidirectional Encoder Representations from Transformers) is trained on uncased text. BERT is based on the transformer architecture that relies on self-attention mechanisms to capture relationships between words in a sequence, enabling effective modeling of long-range dependencies in text data.

Algorithm:

1. The BERT model (English/Hindi) is initialized.
2. Sentence embeddings for sentence 1 are calculated.
3. Sentence embeddings for sentence 2 are calculated.
4. Calculate the cosine similarity scores of the embeddings.

### 3.2.3 Track C:

Cross-linguistic track: The English and Spanish SemRel2024 training datasets were used for training data in languages Hindi and English respectively. The dataset first underwent translation using the deep translation API. By translating the dataset from English to Hindi and Spanish to English, the training dataset for that language was available and was used for testing the development set and the test set of the SemRel 2024 dataset. The models used for training the dataset were the “all-mpnet-base-v2” sentence transformer and “hindi-sentence-bert-nli”<sup>3</sup> by L3cube. These 2 sentence transformers are discussed in Track A, above, however, this went through an additional translation pipeline before that.

<sup>3</sup><https://huggingface.co/l3cube-pune/hindi-sentence-similarity-sbert>

1. Translate sentences from Language 1 to Language 2 using an appropriate translation service or tool.
2. Initialize the model for the task, such as sentence similarity or classification.
3. Encode Sentence1 into a numerical representation using the initialized model. This involves converting the text input into a format suitable for processing by the model, typically through tokenization and embedding.
4. Similarly, encode Sentence2 into a numerical representation using the same BERT model.
5. Train the initialized model on the provided dataset. This step involves feeding the encoded sentence pairs into the model and adjusting the model’s parameters to minimize a predefined loss function, typically using techniques like backpropagation and gradient descent.
6. Evaluate the performance of the trained model on a separate evaluation dataset or through cross-validation. This step aims to assess the model’s ability to generalize to unseen data and its overall effectiveness in the task of interest, such as sentence similarity or classification.
7. If necessary, repeat steps 3 to 6 for all pairs of sentences in the dataset. This process ensures that the model learns from a diverse range of examples and improves its performance across different input scenarios.

## 4 Experimental Setup

The dimensions of the dataset splits are summarized in Table 1, indicating the number of samples allocated for training, development, and testing across different languages. The experimental setup encompassed preprocessing procedures, leveraging Hugging Face Transformers for model access, NumPy for array operations, Pandas for data manipulation, Sentence Transformers for sentence embeddings, and NLTK for various NLP tasks. Evaluation measures such as the F1 score, accuracy, and recall were employed to comprehensively assess the performance of the models across correlation, classification accuracy, and retrieval quality aspects.

## 5 Results

Tables 2, 3, and 4 present the performance results for the three setups (supervised, unsupervised, and cross-lingual) in our experiments. Each table summarizes the F1 score, accuracy, and recall achieved by various models for each language.

Sr.No.	Language	Model Name	F1	Accuracy	Recall
1	English	BERT-base-nli	0.87	0.876	0.84
2	English	SVR	0.59	0.55	0.65
3	English	XG-Boost	0.79	0.82	0.76
4	Marathi	Marathi-NLI	0.83	0.81	0.90
5	Marathi	SVR	0.63	0.61	0.60
6	Marathi	XGBoost	0.66	0.67	0.70

Table 2: Performance Metrics for Track A

This table shows the performance of models in the supervised learning setup, where labeled data was available for training. BERT-based models ("BERT-nli" and "Marathi-nli") consistently outperform other models (SVR, XGBoost) in both English and Marathi, achieving significantly higher correlation scores (0.823 and 0.871, respectively). Interestingly, the Marathi-specific nli model even surpasses the multilingual BERT performance in Marathi, suggesting the benefit of language-specific models.

Sr. No.	Language	Model Name	F1	Accuracy	Recall
1	English	sentence-t5	0.66	0.49	0.49
2	English	BERT based uncased	0.85	0.86	0.80
3	Hindi	Indic-BERT	0.66	0.50	0.50
4	Hindi	hindi-bert-v2	0.66	0.74	0.50

Table 3: Performance Metrics for Track B

This table presents the results for the unsupervised learning setup, where models were trained without relying on labeled data. In English, the BERT-based model ("BERT-base-uncased"<sup>4</sup>) outperforms the pre-trained Sentence-T5<sup>5</sup> model, possibly due to its larger size and fine-tuning on relevant NLP tasks. In Hindi, while both Indic-BERT(Kakwani et al., 2020) and hindi-bert-v2 (Joshi et al., 2022) have similar F1 scores (around 66%), the latter achieves a significantly higher correlation coefficient (0.796). This indicates that hindi-bert-v2<sup>6</sup> captures semantic relatedness more effectively despite similar overall accuracy.

<sup>4</sup><https://huggingface.co/google-bert/bert-base-uncased>

<sup>5</sup><https://huggingface.co/sentence-transformers/sentence-t5-base>

<sup>6</sup><https://huggingface.co/l3cube-pune/hindi-bert-v2>

Sr. No.	Language	Model Name	F1	Accuracy	Recall
1	Spanish to English	all-mpnet-base-v2	0.82	0.82	0.81
2	English to Hindi	hindi-sentence-bert-nli	0.71	0.77	0.92

Table 4: Performance Metrics for Track C

This table shows the performance of models in the cross-lingual learning setup, where the goal was to assess semantic relatedness across different languages. Both models used ("all-mpnet-base-v2" for Spanish-to-English and "hindi-sentence-bert-nli" for English-to-Hindi) achieve worthy correlation scores (0.786 and 0.809, respectively) demonstrating the potential of cross-lingual approaches.

Table 5 represents the results of the development phase.

Track	Language	Sp. Corr Coeff
A	English	0.812
	Marathi	0.855
B	Hindi	0.819
	English	0.825
C	Hindi	0.825
	Englsih	0.790

Table 5: Development Phase Results

## 6 Conclusion

In this paper, we presented a comparative analysis of systems for the Semantic Textual Relatedness (STR) task at SemEval-2024 Task 1. Our approaches, primarily based on language-specific transformer models, achieved top scores on several tracks, including 1st place in Unsupervised Learning for Hindi. Notably, we did not utilize any external datasets, highlighting the effectiveness of our approach despite potential variations in pre-trained model training data.

Prior research in STR has largely focused on English due to limited labeled data for diverse languages. This restricts the true prospect of STR applications like multilingual information retrieval and machine translation. We addressed this gap by exploring solutions for STR in English, Marathi, Hindi, and Spanish. We aim to inspire further research on multilingual STR, particularly for low-resourced languages. We believe larger datasets and broader language coverage hold immense potential for multilingual NLP, unlocking a deeper understanding and empowering cross-cultural communication.

## References

- Mohamed Abdalla, Krishnapriya Vishnubhotla, and Saif M. Mohammad. 2023. [What makes sentences semantically related: A textual relatedness dataset and empirical study](#).
- Weilong Chen, Xin Yuan, Sai Zhang, Jiehui Wu, Yanru Zhang, and Yan Wang. 2020. [Ferryman at semeval-2020 task 3: bert with tfidf-weighting for predicting the effect of context in word similarity](#). In *Proceedings of the fourteenth workshop on semantic evaluation*, pages 281–285.
- Stijn De Saeger, Kow Kuroda, Masaki Murata, Kentaro Torisawa, et al. 2010. [A bayesian method for robust estimation of distributional similarities](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 247–256.
- Selman Delil and Birol Kuyumcu. 2023. [Sefamerve at semeval-2023 task 12: Semantic evaluation of rarely studied languages](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 512–516.
- Evgeniy Gabrilovich, Shaul Markovitch, et al. 2007. [Computing semantic relatedness using wikipedia-based explicit semantic analysis](#). In *IJCAI*, volume 7, pages 1606–1611.
- Dani Gunawan, CA Sembiring, and Mohammad Andri Budiman. 2018. [The implementation of cosine similarity to calculate text relevance between two documents](#). In *Journal of physics: conference series*, volume 978, page 012120. IOP Publishing.
- Ruqaiya Hasan and Michael AK Halliday. 1976. [Cohe-sion in english](#). *London, 1976; Martin JR*.
- Ananya Joshi, Aditi Kajale, Janhavi Gadre, Samruddhi Deode, and Raviraj Joshi. 2022. [L3cube-mahasbert and hindsbert: Sentence bert models and benchmarking bert sentence representations for hindi and marathi](#).
- Raviraj Joshi. 2022. [L3cube-hindbert and devbert: Pre-trained bert transformer models for devanagari based hindi and marathi languages](#). *arXiv preprint arXiv:2211.11418*.
- Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, Gokul N.C., Avik Bhattacharyya, Mitesh M. Khapra, and Pratyush Kumar. 2020. [IndicNLPsuite: Monolingual Corpora, Evaluation Benchmarks and Pre-trained Multilingual Language Models for Indian Languages](#). In *Findings of EMNLP*.
- Wanyin Li, Qin Lu, and Ruifeng Xu. 2005. [Similarity based chinese synonym collocation extraction](#). In *International Journal of Computational Linguistics & Chinese Language Processing, Volume 10, Number 1, March 2005*, pages 123–144.
- 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.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). *arXiv preprint arXiv:1908.10084*.
- Reda Sibli and Leila Kosseim. 2017. [Clac: Semantic relatedness of words and phrases](#).
- Torsten Zesch, Iryna Gurevych, and Max Mühlhäuser. 2007. [Comparing wikipedia and german wordnet by evaluating semantic relatedness on multiple datasets](#). In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pages 205–208.