

BanglaParaphrase: A High-Quality Bangla Paraphrase Dataset

Ajwad Akil*, Najrin Sultana*, Abhik Bhattacharjee, Rifat Shahriyar

Bangladesh University of Engineering and Technology (BUET)

ajwadakillabib@gmail.com, nazrinshukti@gmail.com,

abhik@ra.cse.buet.ac.bd, rifat@cse.buet.ac.bd

Abstract

In this work, we present BanglaParaphrase, a high-quality synthetic Bangla Paraphrase dataset curated by a novel filtering pipeline. We aim to take a step towards alleviating the low resource status of the Bangla language in the NLP domain through the introduction of BanglaParaphrase, which ensures quality by preserving both semantics and diversity, making it particularly useful to enhance other Bangla datasets. We show a detailed comparative analysis between our dataset and models trained on it with other existing works to establish the viability of our synthetic paraphrase data generation pipeline. We are making the dataset and models publicly available at <https://github.com/csebuetnlp/banglaparaphrase> to further the state of Bangla NLP.

1 Introduction

Bangla, despite being the seventh most spoken language by the total number of speakers¹ and fifth most spoken language by native speakers² is still considered a low resource language in terms of language processing. Joshi et al. (2020) have classified Bangla in the language group that has substantial lackings of efforts for labeled data collection and preparation. This lacking is rampant in terms of high-quality datasets for various natural language tasks, including paraphrase generation.

Paraphrases can be roughly defined as pairs of texts that have similar meanings but may differ structurally. So the task of generating paraphrases given a sentence is to generate sentences with different wordings or/and structures to the original sentences while preserving the meaning. Paraphrasing can be a vital tool to assist language understanding tasks such as question answering (Pazzani and Engelman, 1983; Dong et al., 2017), style transfer (Krishna et al., 2020), semantic parsing (Cao

et al., 2020), and data augmentation tasks (Gao et al., 2020).

Paraphrase generation has been a challenging problem in the natural language processing domain as it has several contrasting elements, such as semantics and structures, that must be ensured to obtain a good paraphrase of a sentence. Syntactically Bangla has a different structure than high-resource languages like English and French. The principal word order of the Bangla language is subject-object-verb (SOV). Still, it also allows free word ordering during sentence formation. The pronoun usage in the Bangla language has various forms, such as "very familiar", "familiar", and "polite forms"³. It is imperative to maintain the coherence of these forms throughout a sentence as well as across the paraphrases in a Bangla paraphrase dataset. Following that thread, we create a Bangla Paraphrase dataset ensuring good quality in terms of semantics and diversity. Since generating datasets by manual intervention is time-consuming, we curate our BanglaParaphrase dataset through a pivoting (Zhao et al., 2008) approach, with additional filtering stages to ensure diversity and semantics. We further study the effects of dataset augmentation on a synthetic dataset using masked language modeling. Finally, we demonstrate the quality of our dataset by training baseline models and through comparative analysis with other Bangla paraphrase datasets and models. In summary:

- We present BanglaParaphrase, a synthetic Bangla Paraphrase dataset ensuring both diversity and semantics.
- We introduce a novel filtering mechanism for dataset preparation and evaluation.

2 Related Work

Paraphrase generation datasets and models are heavily dominated by high-resource languages

*These authors contributed equally to this work.

¹<https://w.wiki/Pss>

²<https://w.wiki/Psq>

³https://en.wikipedia.org/wiki/Bengali_grammar

such as English. But for low-resource languages such as Bangla, this domain is less explored. To our knowledge, only (Kumar et al., 2022) described the use of IndicBART (Dabre et al., 2021) to generate paraphrases using the sequence-to-sequence approach for the Bangla language. One of the most challenging barriers to paraphrasing research for low-resource languages is the shortage of good-quality datasets. Among recent work on low-resource paraphrase datasets, (Kanerva et al., 2021) introduced a comprehensive dataset for the Finnish language. The OpusParcus dataset (Creutz, 2018) consists of paraphrases for six European languages. For Indic languages such as Tamil, Hindi, Punjabi, and Malayalam, Anand Kumar et al. (2016) introduced a paraphrase detection dataset in a shared task. Scherrer (2020) introduced a paraphrase dataset for 73 languages, where there are only about 1400 sentences in total for the Bangla language, mainly consisting of simple sentences.

3 Paraphrase Dataset Generation and Curation

3.1 Synthetic Dataset Generation

We started by scraping high-quality representative sentences for the Bangla web domain from the RoarBangla website⁴ and translated them from Bangla to English using the state-of-the-art translation model developed in (Hasan et al., 2020) with 5 references. For the generated English sentences, 5 new Bangla translations were generated using beam search. Among these multiple generations, only those (original sentence, back-translated sentence) pairs were chosen as candidate datapoints where the LaBSE (Feng et al., 2022) similarity score for both (original Bangla and back-translated Bangla), as well as (original Bangla and translated English) were greater than 0.7⁵. After this process, there were more than 1.364M sentences with multiple references for each source.

3.2 Novel Filtering Pipeline

As mentioned in (Chen and Dolan, 2011), paraphrases must ensure the fluency, semantic similarity, and diversity. To that end, we make use of different metrics evaluating each of these aspects as **filters**, in a pipelined fashion.

⁴<https://roar.media/bangla>

⁵We chose 0.7 as the LaBSE semantic similarity threshold following (Bhattacharjee et al., 2022a)

To ensure diversity, we chose **PINC** (*Paraphrase In N-gram Changes*) among various diversity measuring metrics such as (Chen and Dolan, 2011; Sun and Zhou, 2012) as it considers the lexical dissimilarity between the source and the candidates. We name this first filter as **PINC Score Filter**. To use this metric for filtering, we determined the optimum threshold value empirically by following a plot⁶ of the data yield against the PINC score, indicating the amount of data having at least a certain amount of PINC score. We chose the threshold value that maximizes the PINC score with over 63.16% yield.

Since contextualized token embeddings have been shown to be effective for paraphrase detection (Devlin et al., 2019), we use BERTScore (Zhang et al., 2019) to ensure semantic similarity between the source and candidates. After our PINC filter, we experimented with BERTScore, which uses the multilingual BERT model (Devlin et al., 2019) by default. We also experimented with BanglaBERT (Bhattacharjee et al., 2022a) embeddings and decided to use this as our semantic filter since BanglaBERT is a monolingual model performing exceptionally well on Bangla NLU tasks. We select the threshold similar to the PINC filter by following the corresponding plot, and in all of our experiments, we used F1 measure as the filtering metric. We name this second filter as **BERTScore Filter**. Through a human evaluation⁷ of 300 randomly chosen samples, we deduced that pairs having BERTScore (with BanglaBERT embeddings) ≥ 0.92 were semantically sound and decided to use this as a starting point to figure out our desired threshold. We further validated our choice of parameters through model-generated paraphrases, with the models trained on filtered datasets using different parameters (detailed in Section 4.1).

Initially training on the resultant dataset from the previous two filters, we noticed that some of the predicted paraphrases were growing unnecessarily long by repeating parts during inference. As repeated N-grams within the corpus most likely have been the culprit behind this, attempts to ameliorate the issue were made by introducing our third filter, namely **N-gram Repetition Filter**, where we tested the target side of our dataset to see if there were any N-gram repeats with a value of N from 1 to 4. We obtained less than 200 sentences on the

⁶More details are presented in the Appendix

⁷More details are presented in the ethical considerations section

Filter Name	Significance	Filtering Parameters
PINC	Ensure diversity in generated paraphrase	0.65, 0.76, 0.80
BERTScore	Preserve semantic coherence with the source	lower 0.91 - 0.93, upper 0.98
N-gram repetition	Reduce n-gram repetition during inference	2 - 4 grams
Punctuation	Prevent generating non-terminating sentences during inference	N/A

Table 1: Filtering Scheme

target side with a 2-gram repetition and decided to use $N = 2$ for this filter. Additionally, we removed sentences without terminating punctuation from the corpus to ensure a noise-free dataset before proceeding with the training. We term this last filter as **Punctuation Filter**. The filters, along with their significance and parameters, have been summarised in Table 1.

3.3 Evaluation Metrics

Following the work of (Niu et al., 2021), we used multiple metrics to evaluate several criteria in our generated paraphrase. For **quality**, we used sacreBLEU (Post, 2018) and ROUGE-L (Lin, 2004). We used the multilingual ROUGE scoring implementation introduced by (Hasan et al., 2021) which supports Bangla stemming and tokenization. For **syntactic diversity**, we used the PINC score as we did for filtering. For measuring **semantic correctness**, we used BERTScore F1-measure with BanglaBERT embeddings. Additionally, we used a modified version of a hybrid score named BERT-iBLEU score (Niu et al., 2021) where we also used BanglaBERT embeddings for the BERTScore part. This hybrid score measures semantic similarity while penalizing syntactical similarity to ensure the diversity of the paraphrases. More details about evaluation scores can be found in the Appendix.

3.4 Diverse Dataset Generation by Masked Language Modeling

We wondered whether the dataset could be further augmented through replacing tokens from a particular part of speech with other synonymous tokens.

To that end, we fine-tuned BanglaBERT (Bhattacharjee et al., 2022a) for POS tagging with a token classification head on the (Sankaran et al., 2008) dataset containing 30 POS tags.

The idea of augmenting the dataset with masking follows the work of (Mohiuddin et al., 2021). We first tagged the parts of speech of the source side of our synthetic dataset and then chose 7 Bangla parts of speech to maximize the diversification in syntactic content. We masked the corresponding tokens

and filled them through MLM sequentially. We used both XLM-RoBERTa (Conneau et al., 2020) and BanglaBERT to perform MLM out of the box. Of these two, BanglaBERT performed mask-filling with less noise, and thus we selected the results of this model. To ensure consistency with our initial dataset, we also filtered these with our pipeline outlined in Section 3.2 by choosing the PINC score threshold of 0.7⁸ and (0.92 - 0.98) (lower and upper limit) for the BERTScore threshold, obtaining about 70K sentences. We used this dataset for training models with our initially filtered one in a separate experiment.⁹

4 Experiments and Results

4.1 Experimental Setup

We first filtered the synthetic dataset with our 4-stage filtering mechanisms and then fine-tuned mT5-small model (Xue et al., 2021), keeping the default learning rate as 0.001 for 10 epochs. In each of the experiments, we changed the dataset by keeping the model fixed as our objective was to find the threshold for the first two filters for which the metrics on both the validation and the test set of the individual dataset gave us promising results. We conducted several experiments by varying PINC scores from (0.65, 0.76, 0.80) and BERTScore from (0.91, 0.92, 0.93) and 0.98 (lower and upper limit) by following respective plots.

The evaluation metrics for each experiment were tracked, and we examined how the thresholds affected the metrics for the test set of the dataset we were experimenting with. We finally chose the effective threshold to be **0.76** for the PINC score and **0.92 - 0.98** (lower and upper limit) for BERTScore such that it provides a good balance between good automated evaluation scores and data amount, and obtained **466630** parallel paraphrase pairs. We fine-tuned mT5-small, and BanglaT5 (Bhattacharjee et al., 2022c) with the BanglaParaphrase training

⁸We lowered the threshold since this augmentation does not diversify in terms of the structure of the sentences

⁹Further details of the whole experiment can be found in the Appendix.

Test Set	Model	sacreBLEU	ROUGE-L	PINC	BERTScore	BERT-iBLEU
BanglaParaphrase	mT5-small	<u>20.9</u>	53.57	80.5	<u>94.20</u>	92.67
	mT5-small-aug	19.90	<u>53.63</u>	<u>80.72</u>	94.00	<u>92.54</u>
	BanglaT5	32.8	63.58	74.40	94.80	92.18
	BanglaT5-aug	32.5	63.43	74.41	94.80	92.18
	IndicBART	5.60	35.61	80.26	91.50	91.16
	IndicBARTSS	4.90	33.66	82.10	91.10	90.95
IndicParaphrase	mT5-small	7.3	18.66	<u>82.30</u>	<u>94.30</u>	<u>89.06</u>
	mT5-small-aug	7.0	18.27	82.80	94.10	89.00
	BanglaT5	<u>11.00</u>	19.99	74.50	94.80	87.738
	BanglaT5-aug	11.00	20.10	74.43	94.80	87.540
	IndicBART	12.00	21.58	76.83	93.30	90.65
	IndicBARTSS	10.7	<u>20.59</u>	77.60	93.10	90.54

Table 2: Test results of different models on BanglaParaphrase and IndicParaphrase Test Set where bold items indicate best results and underlined items indicate the runner up

set as well as with a MLM augmented dataset as mentioned in Section 3.4. For training, validation, and testing purposes, we randomly split the whole dataset into 80:10:10 ratios. We sampled the MLM dataset twice for the second dataset and added it to our initial training and validation set. After augmentation, the dataset consisted of **603672** parallel pairs with **551324** pairs for training and **29016** for validation. We used the same testing set consisting of **23332** parallel pairs for all the models.¹⁰ And finally we used the IndicBART and IndicBARTSS (Dabre et al., 2021) fine-tuned on the IndicParaphrase dataset (Kumar et al., 2022) to generate predictions and compute the evaluation scores for comparative analysis.

Hyperparameter Tuning We fine-tuned mT5-small for 10-15 epochs, tuning the learning rate from $3e-4$ to $1e-3$. BanglaT5 was fine-tuned for 10 epochs with a learning rate of $5e-4$ and a warmup ratio of 0.1. We chose the final models based on the validation performance of the sacreBLEU score. During inference for the mT5-small model, we used top-K (Fan et al., 2018) sampling with a value of 50 in combination with top-P sampling with a value of 0.95 along with beam search for generating multiple inferences, which we filter by PINC score of 0.74 followed by max BERTScore. For BanglaT5, the inference was simply made with a beam search with a beam length of 5.

4.2 Results and Comparison

In Table 2, we show how our trained models namely mT5-small, mT5-small-aug¹¹, BanglaT5 and BanglaT5-aug models as well as IndicBART and IndicBARTSS perform on our released test set and Indic test Set (only Bangla) from IndicParaphrase dataset. A few examples of how mT5-small performs on the BanglaParaphrase test set and a detailed comparison of the IndicParaphrase dataset with our dataset in terms of diversity and semantics can be found in the Appendix.

For the BanglaParaphrase test set, we observe that all the evaluation scores are almost similar for both mT5-small and BanglaT5 trained on the original dataset as well as the MLM augmented dataset. We find that the BanglaT5 model performs best on sacreBLEU, ROUGE-L, and BERTScore for our test set. We also observe that both the IndicBART models achieve lower scores in all the metrics except PINC, which is not sufficient enough to ensure the quality of generated paraphrases. The scores on sacreBLEU and ROUGE-L are particularly low compared to what our trained models achieved. As for the PINC score, IndicBARTSS achieved the highest value, with mT5 models slightly trailing behind. Since all other scores are lower, this high PINC score has low significance. As for the hybrid score, we find that mT5-small trained on the BanglaParaphrase training set achieves the best result on our test set, with BanglaT5 models trailing slightly lower and IndicBART models having a much lower value.

For the IndicParaphrase test set, we observe

¹⁰MLM augmented dataset is for experimental purpose only

¹¹aug means the models were trained with MLM augmented BanglaParaphrase training set

that mT5 models perform poorly in sacreBLEU and ROUGE-L scores, whereas BanglaT5 models perform very competitively with IndicBART models inspite of being only fine-tuned on our dataset, which has virtually no overlap with IndicParaphrase training set. We also observe that both mT5 and BanglaT5 trained on the BanglaParaphrase training set and augmented training set have similar performance on all the metrics for this test set. We find both the BanglaT5 models achieve the highest BERTScore, beating IndicBART and IndicBARTSS, and both mT5 models trail closely to BanglaT5. So BanglaT5 can generalize well on other datasets. As for the PINC score, we see that mT5-small-aug achieves the highest score among all the models. And finally, for the hybrid score, we find both IndicBART models achieving the best score. We believe the reason for IndicBART to have higher scores is that it has a high PINC score, i.e., less similarity with the source, which results in a higher BERT-iBLEU score.

Overall, the models trained on the BanglaParaphrase data set, specifically BanglaT5, perform competitively with the IndicBART models, even besting in terms of semantics concerning the source, while generating diverse paraphrases and thus validating that our dataset not only ensures good diversity but semantics as well.

5 Conclusion & Future Works

In this work, starting from a pure synthetic paraphrase dataset, we introduced an automated filtering pipeline to curate a high-quality Bangla Paraphrase dataset, ensuring both diversity and semantics. We trained the mT5-small and BanglaT5 models with our dataset to generate quality paraphrases of Bangla sentences. Our choice of the initial monolingual corpus has been made to include highly representative sentences for the Bangla language, which is large enough for an isolated paraphrase generation task. The corpus can easily be extended for desired pretraining tasks using a larger monolingual corpus. Furthermore, we plan on improving the MLM scheme by automating parts of speech selection and using LaBSE with BanglaBERT embeddings to compare semantics at the sentence level, which would ensure better filters and better evaluation of generated paraphrases. Though our work is language-agnostic, the extent to which our approach applies to other low-resource languages given language-specific components (datasets and

models) is subject to further experimentation. In future work, we want to investigate the viability of our synthetic data generation pipeline in the context of paraphrase datasets in different languages included in popular benchmarks such as (Gehrmann et al., 2022). Additionally, we want to investigate how our paraphrase dataset and models can be used to improve the performance of other low-resource tasks in Bangla, such as Readability detection (Chakraborty et al., 2021) and Cross-lingual summarization (Bhattacharjee et al., 2022b)

Acknowledgements

We would like to thank the Research and Innovation Centre for Science and Engineering (RISE), BUET, for funding the project.

Ethical Considerations

Dataset and Model Release The *Copy Right Act, 2000*¹² of Bangladesh allows public release and reproduction of copy-right materials for non-commercial research purposes. As valuable research work for Bangla Language, we will release the BanglaParaphrase dataset under a non-commercial license. Additionally, we will release the relevant codes and the trained models for which we know the distribution will not cause copyright infringement.

Manual Efforts The manual observations regarding the choice of primary BERTScore threshold which is reflective of high semantic quality by going through 300 randomly chosen samples were done by the native authors.

References

- M Anand Kumar, Shivkaran Singh, B Kavirajan, and KP Soman. 2016. *Shared task on detecting paraphrases in Indian languages (DPIL): An overview*. In *Forum for Information Retrieval Evaluation*, pages 128–140. Springer.
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Ahmad, Kazi Samin Mubasshir, Md Saiful Islam, Anindya Iqbal, M. Sohel Rahman, and Rifat Shahriyar. 2022a. *BanglaBERT: Language model pretraining and benchmarks for low-resource language understanding evaluation in Bangla*. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1318–1327, Seattle, United States. Association for Computational Linguistics.

¹²<http://bdlaws.minlaw.gov.bd/act-details-846.html>

- Abhik Bhattacharjee, Tahmid Hasan, Wasi Uddin Ahmad, Yuan-Fang Li, Yong-Bin Kang, and Rifat Shahriyar. 2022b. [CrossSum: Beyond English-centric cross-lingual abstractive text summarization for 1500+ language pairs](#). *arXiv:2112.08804*.
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Uddin Ahmad, and Rifat Shahriyar. 2022c. [BanglaNLG: Benchmarks and resources for evaluating low-resource natural language generation in Bangla](#). *arXiv:2205.11081*.
- Ruisheng Cao, Su Zhu, Chenyu Yang, Chen Liu, Rao Ma, Yanbin Zhao, Lu Chen, and Kai Yu. 2020. [Unsupervised dual paraphrasing for two-stage semantic parsing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6806–6817, Online. Association for Computational Linguistics.
- Susmoy Chakraborty, Mir Tafseer Nayeem, and Wasi Uddin Ahmad. 2021. [Simple or complex? learning to predict readability of Bengali texts](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12621–12629.
- David Chen and William B Dolan. 2011. [Collecting highly parallel data for paraphrase evaluation](#). In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, pages 190–200.
- 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.
- Mathias Creutz. 2018. [Open subtitles paraphrase corpus for six languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Raj Dabre, Himani Shrotriya, Anoop Kunchukuttan, Ratish Puduppully, Mitesh M Khapra, and Pratyush Kumar. 2021. [IndicBART: A pre-trained model for natural language generation of Indic languages](#). *arXiv:2109.02903*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Li Dong, Jonathan Mallinson, Siva Reddy, and Mirella Lapata. 2017. [Learning to paraphrase for question answering](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 875–886, Copenhagen, Denmark. Association for Computational Linguistics.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2022. [Language-agnostic BERT sentence embedding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.
- Silin Gao, Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020. [Paraphrase augmented task-oriented dialog generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 639–649, Online. Association for Computational Linguistics.
- Sebastian Gehrmann, Abhik Bhattacharjee, Abinaya Mahendiran, Alex Wang, Alexandros Papangelis, Aman Madaan, Angelina McMillan-Major, Anna V. Shvets, Ashish Upadhyay, Bingsheng Yao, Bryan Wilie, Chandra Bhagavatula, Chaobin You, Craig Thomson, Cristina Garbacea, Dakuo Wang, Daniel Deutsch, Deyi Xiong, Di Jin, Dimitra Gkatzia, Dragomir Radev, Elizabeth Clark, Esin Durmus, Faisal Ladhak, Filip Ginter, Genta Indra Winata, Hendrik Strobelt, Hiroaki Hayashi, Jekaterina Novikova, Jenna Kanerva, Jenny Chim, Jiawei Zhou, Jordan Clive, Joshua Maynez, João Sedoc, Juraj Juraska, Kaustubh D. Dhole, Khyathi Raghavi Chandu, Leonardo F. R. Ribeiro, Lewis Tunstall, Li Zhang, Mahima Pushkarna, Mathias Creutz, Michael White, Mihir Kale, Moussa Kamal Eddine, Nico Daheim, Nishant Subramani, Ondrej Dusek, Paul Pu Liang, Pawan Sasanka Ammanamanchi, Qinqin Zhu, Ratish Puduppully, Reno Kriz, Rifat Shahriyar, Ronald Cardenas, Saad Mahamood, Salomey Osei, Samuel Cahyawijaya, Sanja vStajner, Sébastien Montella, Shailza, Shailza Jolly, Simon Mille, Tahmid Hasan, Tianhao Shen, Tosin P. Adewumi, Vikas Raunak, Vipul Raheja, Vitaly Nikolaev, Vivian Tsai, Yacine Jernite, Yi Xu, Yisi Sang, Yixin Liu, and Yufang Hou. 2022. [Gemv2: Multilingual NLG benchmarking in a single line of code](#). *arXiv:2206.11249*.
- Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, and Rifat Shahriyar. 2021. [XL-Sum: Large-scale multilingual abstractive summarization for 44 languages](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4693–4703, Online. Association for Computational Linguistics.

- Tahmid Hasan, Abhik Bhattacharjee, Kazi Samin, Masum Hasan, Madhusudan Basak, M. Sohel Rahman, and Rifat Shahriyar. 2020. [Not low-resource anymore: Aligner ensembling, batch filtering, and new datasets for Bengali-English machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2612–2623, Online. Association for Computational Linguistics.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Jenna Kanerva, Filip Ginter, Li-Hsin Chang, Iiro Rastias, Valtteri 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.
- Kalpesh Krishna, John Wieting, and Mohit Iyyer. 2020. [Reformulating unsupervised style transfer as paraphrase generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 737–762, Online. Association for Computational Linguistics.
- Aman Kumar, Himani Shrotriya, Prachi Sahu, Raj Dabre, Ratish Puduppully, Anoop Kunchukuttan, Amogh Mishra, Mitesh M Khapra, and Pratyush Kumar. 2022. [IndicNLG suite: Multilingual datasets for diverse NLG tasks in Indic languages](#). *arXiv:2203.05437*.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Tasnim Mohiuddin, M Saiful Bari, and Shafiq Joty. 2021. [AugVic: Exploiting BiText vicinity for low-resource NMT](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3034–3045, Online. Association for Computational Linguistics.
- Tong Niu, Semih Yavuz, Yingbo Zhou, Nitish Shirish Keskar, Huan Wang, and Caiming Xiong. 2021. [Unsupervised paraphrasing with pretrained language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5136–5150, Online and Punta Cana, Dominican Republic. 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.
- Michael J. Pazzani and Carl Engelman. 1983. [Knowledge based question answering](#). In *First Conference on Applied Natural Language Processing*, pages 73–80, Santa Monica, California, USA. 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.
- Gowtham Ramesh, Sumanth Doddapaneni, Aravindh Bheemaraj, Mayank Jobanputra, Raghavan AK, Ajitesh Sharma, Sujit Sahoo, Harshita Diddee, Mahalakshmi J, Divyanshu Kakwani, Navneet Kumar, Aswin Pradeep, Srihari Nagaraj, Kumar Deepak, Vivek Raghavan, Anoop Kunchukuttan, Pratyush Kumar, and Mitesh Shantadevi Khapra. 2022. [Samanantar: The largest publicly available parallel corpora collection for 11 Indic languages](#). *Transactions of the Association for Computational Linguistics*, 10:145–162.
- Baskaran Sankaran, Kalika Bali, Monojit Choudhury, Tanmoy Bhattacharya, Pushpak Bhattacharyya, Girish Nath Jha, S. Rajendran, K. Saravanan, L. Sobha, and K.V. Subbarao. 2008. [A common parts-of-speech tagset framework for Indian languages](#). In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Yves Scherrer. 2020. [TaPaCo: A corpus of sentential paraphrases for 73 languages](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 6868–6873, Marseille, France. European Language Resources Association.
- Hong Sun and Ming Zhou. 2012. [Joint learning of a dual SMT system for paraphrase generation](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 38–42, Jeju Island, Korea. Association for Computational Linguistics.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. [BERTScore: Evaluating text generation with BERT](#). *arXiv:1904.09675*.

Shiqi Zhao, Haifeng Wang, Ting Liu, and Sheng Li. 2008. [Pivot approach for extracting paraphrase patterns from bilingual corpora](#). In *Proceedings of ACL-08: HLT*, pages 780–788, Columbus, Ohio. Association for Computational Linguistics.

Jianing Zhou and Suma Bhat. 2021. [Paraphrase generation: A survey of the state of the art](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5075–5086, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Appendix

PINC Score Details

PINC score is defined as for source sentence s and candidate sentence c as:

$$\frac{1}{N} \sum_{n=1}^N 1 - \frac{|ngram_s \cap ngram_c|}{|ngram_c|}$$

Where N is defined as the maximum n-gram we considered, and $ngram_s$ and $ngram_c$ are the lists of n-grams present in the source and candidate sentences. In all experiments, we use $N = 4$. This score can be treated as the inverse of the BLEU score since it minimizes the number of n-gram overlaps between the two sentences. We also present a PINC score vs. data amount plot in Figure 1, which we used to select the thresholds.




Figure 1: PINC Score range within [0-1] for whole BanglaParaphrase dataset

BERTScore Plot

A plot of BERTScore with BanglaBERT embeddings after the BanglaParaphrase dataset has been filtered with a PINC score of 0.76 threshold is shown in Figure 2.

Evaluation Metric Details

BLEU, METEOR, and ROUGE-L are the most common metrics used (Zhou and Bhat, 2021) for paraphrase evaluation. BLEU (Papineni et al., 2002) is a widely used metric for machine translation evaluation that ensures semantic adequacy and fluency. But it falls short for paraphrase evaluation as mentioned by (Niu et al., 2021; Zhou and Bhat, 2021). A unified metric that captures all the elements of evaluating paraphrase is still lacking




Figure 2: BERTScore with BanglaBERT embeddings within range [0.9-1.0] after whole dataset being filtered by PINC threshold of 0.76

(Zhou and Bhat, 2021), and so we present the details about different evaluation metrics we used and the criteria they measure:

Quality To ensure the quality of the generated paraphrases with respect to the target, we used sacreBLEU Score (Post, 2018) and ROUGE-L (Lin, 2004) F1-measure. Both of the scores produce a real number between the range [0 – 1], and we present the scores in percentages for our results.

Syntactic Diversity To evaluate the diversity between the generated paraphrases and the sources, we used the PINC score (Chen and Dolan, 2011). This score produces a real number between the range [0 – 1] and we report the arithmetic mean for all the sentences in the test set and present in terms of percentages for our results.

Semantic Correctness To evaluate semantic correctness, the arithmetic mean of BERTScore (Zhang et al., 2019) F1-measure between source and predictions is used. As discussed, this is a modified version of BERTScore which uses BanglaBERT embeddings to produce a real number between [0 – 1], and we present it in terms of percentages for our results.

Hybrid Score And finally, we used a modified version of a hybrid score named BERT-iBLEU introduced in (Niu et al., 2021). The formula to compute the score is:

$$\left(\frac{\beta * BERTScore^{-1} + 1.0 * (1 - selfBLEU)^{-1}}{\beta + 1.0} \right)^{-1}$$

This metric measures semantic similarity while penalizing syntactical similarity at the same time.

For the semantic similarity part, the authors used BERTScore between target and predictions, which we modified to use BERTScore with BanglaBERT embeddings. For diversity, self-BLEU was calculated between the source and the prediction. The more dissimilar the source is to the candidate, the higher will be the value of 1-selfBLEU. The final score is a weighted harmonic mean between these two scores. We used the value of β to be 4.0, as chosen by the authors. The score produces a real number between the range $[0 - 1]$, and as our modified BERTScore gives us scores in a high range (> 0.9), the scores produced by this metric is also in high range. We present the score in terms of percentages for our results.

Diverse Dataset Generation Experiment Details

We trained BanglaBERT with a token classification head with (Sankaran et al., 2008) dataset containing 30 POS tags and the entire corpus consists of 7393 sentences corresponding to 102937 tokens. We trained for 20 epochs, with a batch size of 32 and a learning rate of 0.00002 with a linear learning rate scheduler. The dataset was split into an 80:10:10 ratio into a train, test, and validation sets. We obtained close to 90% F1-Score on the test set. The test set metrics are showed in Table 3.

Dataset	Accuracy	Precision	Recall	F1-score
Test	0.924	0.896	0.900	0.898

Table 3: Validation and Test metrics for POS tagging experiment

1. VM(Main Verb): Denotes the eventual information in a sentence Example: যুমে দু চোখ জড়িয়ে(VM) আসে, আমি দোকানে গিয়েছিলাম(VM), বইটা ধরুন (VM)
2. VA(Auxiliary Verb): Helping Verbs Example: যেতে(VA) ভালোবাসে, দেখতে গিয়েছিলাম (VA), বিপদ ঘটিয়ে থাকে(VA)
3. JJ(Adjective): POS that modifies a Noun Example: আবার সেই কাজে জড়িত(JJ) হবে, চমকপ্রদ(JJ) সাফল্য, সে দ্রুত(JJ) হাটা দিলো
4. NV(Verbal Noun): Gerund and Gerundival constructs in Bangla Example: বর্ণনা করার(NV) জন্য, ঢাকা দেওয়া(NV) ভাত, সকালে উঠে জুসিং করা(NV) ভালো
5. AMN(Adverb of Manner): Adverbs modifying the way actions are described in the verb Example: আর(AMN) তর সইছিল না, আবার(AMN) জড়িত হতে হবে, কিভাবে(AMN) গাড়ি চালাতে হয়?
6. ALC(Adverb of Location): POS that denotes time and space that modifies the verb Example: আজ(ALC) পুলিশের বড় বাহিনী, এখানটার (ALC) বসে, আজও(ALC) করাছি কালও (ALC) করবো
7. NST(Spatio Temporal Noun): These are the nouns that denote space, time, direction etc Example: ওপরে(NST) দাড় করিয়ে দেয়, বারান্দায় দাড়াতেই সামনে(NST), কাজটি করার আগেই(NST)

Figure 3: Selected POS Details

After training the POS tagger, we tagged 7 carefully chosen parts of speeches namely VM (Main verb), VA (Auxiliary Verb), JJ (Adjective), NV (Verbal Noun), AMN (Adverb of Manner), ALC (Adverb of location), and NST(Spatio Temporal Noun). These POS were masked and filled in the order as

mentioned here. The parts of speeches with minimal description are shown in Figure 3. A demonstration for mask filling is shown in Figure 4.

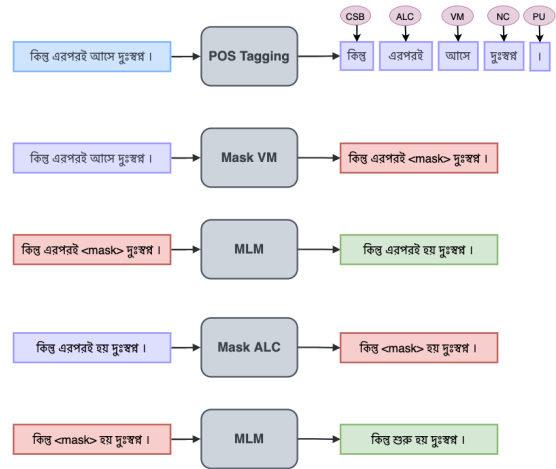


Figure 4: Diverse Sentence Generation by Mask Filling

Examples of Generated Paraphrase


We show some examples of generated paraphrases by mT5 small model on BanglaParaphrase dataset in Figure 5.

Source: সেক্ষেত্রে দেখতে হবে যে কার কোন খাবারের ক্ষেত্রে সমস্যা হচ্ছে। Target: সেই ক্ষেত্রে আমাদের দেখতে হবে যে, কিছু খাবার নিয়ে কার সমস্যা রয়েছে। Prediction: জাহলে আমাদের দেখতে হবে যে, কোনো খাবারের ক্ষেত্রে কোনও সমস্যা হয়।
Source: ধারণা করা হচ্ছে, বোতলটি জার্মানির একটি জাহাজ থেকে ছুঁতে ফেলা হয়েছিলো ভারত মহাসাগরে। Target: বোতলটি একটি জার্মান জাহাজ থেকে ভারত মহাসাগরে নিক্ষেপ করা হয়েছে বলে মনে করা হয়। Prediction: বোতলটি ভারত মহাসাগরে জার্মান জাহাজ থেকে ছুঁতে ফেলা হয়েছিল বলে ধারণা করা হয়।
Source: হোঁজ খবর রাখতেন বিজ্ঞানের অগ্রগতি নিয়ে। Target: বিজ্ঞানের অগ্রগতির দিকে তিনি নজর রেখেছিলেন। Prediction: বিজ্ঞানের অগ্রগতি সম্পর্কে তিনি খবর রাখেন।
Source: খুব দ্রুত এই টিকা তৈরি হয় আর কাজ করে চমৎকারভাবে। Target: টিকাটি দ্রুত বিকশিত হয় এবং খুব ভালভাবে কাজ করে। Prediction: ভ্যাকসিনটি খুব দ্রুত নির্মিত হয় এবং চমৎকারভাবে কাজ করে।
Source: সেটা খুবই একটা অশুভ লক্ষণ। Target: এটা একটা খারাপ লক্ষণ। Prediction: এটা খুবই মন্দ লক্ষণ।


Figure 5: Examples of Generated Paraphrase by mT5 small on released test set (trained with released training set)

BERTScore Distribution Analysis

BERTScore with mBERT gives us a value in a much more comprehensive range, $[0.7 - 0.1]$, and most scores are centered around $[0.8 - 0.9]$ as we can see from the histogram in Figure 6a whereas BERTScore with BanglaBERT embeddings gives us a score in a much higher range, $[0.8 - 0.1]$ and most of the scores are centered around $[0.9 - 0.95]$ as seen in Figure 6b. So BERTScore with BanglaBERT embeddings score above 0.8 for sentences with lesser semantic simi-



(a) BERTScore Histogram



(b) BERTScore Histogram (BanglaBERT embeddings)

Figure 6: Histograms for original dataset


larity but above 0.9 for sentences with good semantic similarity.

Comparison with IndicNLG Paraphrasing Dataset


The IndicNLG Suite (Kumar et al., 2022) has data for eleven languages: Assamese, Bangla, Gujarati, Hindi, Marathi, Odiya, Punjabi, Kannada, Malayalam, Tamil, and Telugu. The dataset has 5.57M in size overall. For Bangla Paraphrase, there are 890,445 sentences in the train set, 10,000 in the validation set, and 10,000 in the test set, with each source sentence having 5 references. The dataset uses Samanantar corpus (Ramesh et al., 2022) to generate the paraphrases by a back-translation mechanism. Then the authors filtered the sentences by removing noise and duplicates and evaluated the diversity by a scheme developed by them. They screened the sentences in a way to ensure enough diversity among the source and the references. They reported 5 references for each source sentence, which are ordered from most to least diverse. The dataset ensures diversity by a filtering mechanism developed by the authors, but they

did not include any filtering mechanism to ensure semantic similarity between the sources or the references. As the initial set of sources and the references were generated by pivoting, there are a lot of changes and variations and thus, it is vital to ensure both diversity and meaning.

To analyze, we plot the scores for the reference with most diversity in terms of PINC score. We started with the PINC score vs. data amount plot in Figure 7a. The shape of the plot looks a lot similar to the PINC plot for our whole dataset in Figure 1. We also observe that above or equal to the 0.7 threshold, there are about 0.72M sentences. And for thresholds 0.74 and 0.76, there are about close to 0.7M sentences (about 77% of the total sentences) and close to 0.66M sentences (about 73% of the total sentences), respectively. Compared to our filtering, where we chose the PINC filter to be 0.76 and ended up with about 0.86M sentences (about 63.05% of our total corpus size), the dataset ensured more diverse paraphrases.



(a) PINC Score for range [0-1.0]



(b) PINC Score for range [0.7-0.8]

Figure 7: PINC Score for IndicParaphrase dataset

We see a different scenario for the case of BERTScore (calculated with BanglaBERT embed-

ding) vs. the data amount plot for the whole dataset. In Figure 8, we observe by taking a closer look at BERTScore for the range of [0.9 - 1.0] that the amount of sentences for threshold of 0.92 is about 0.31M (35% of the whole dataset) and for 0.93 about 0.23M sentences (about 25% of the whole dataset). Compared to our dataset, for a threshold above 0.92 for BERTScore, we have a little more than 0.5M (about 37% of our dataset), and for 0.93, we have about 0.367M sentences (about 27% of our whole dataset), as seen in Figure 2. This indicates that semantic meaning is more preserved in our dataset as we only took the sentences that ensured high semantics in the whole corpus for constructing our final BanglaParaphrase dataset.




Figure 8: BERTScore with BanglaBERT embeddings for IndicParaphrase Dataset for Range [0.90-1.0]




Figure 9: LaBSE Similarity Score for range [0-1.0]

We also observe an analysis with LaBSE similarity score for IndicParaphrase dataset¹³ where we follow from Figure 9 that above 0.6 there are about more than 0.8M sentences which drastically reduces as the threshold rises. We also observe that

¹³only scores above 0 are shown in the plots

above 0.7, there are close to 0.8M sentences. If we look above 0.8, we find that the value drastically reduces to a little more than 0.5M sentences, which is just about 57% of the total data. If we look above 0.85, we only find about 0.35M sentences, which is about 38% of the total data available, and it corresponds closely to the amount of 0.31M for BERTScore of 0.92 or above that we discussed.

So the analysis leads us to the inference that the IndicParaphrase dataset is diverse, but it falls short in terms of semantics between the source and the references.