

BanglaNLG and BanglaT5: Benchmarks and Resources for Evaluating Low-Resource Natural Language Generation in Bangla

Abhik Bhattacharjee¹, Tahmid Hasan¹, Wasi Uddin Ahmad², Rifat Shahriyar¹

Bangladesh University of Engineering and Technology (BUET)¹,

University of California, Los Angeles²

abhik@ra.cse.buet.ac.bd, {tahmidhasan,rifat}@cse.buet.ac.bd

Abstract

This work presents ‘BanglaNLG,’ a comprehensive benchmark for evaluating natural language generation (NLG) models in Bangla, a widely spoken yet low-resource language. We aggregate six challenging conditional text generation tasks under the BanglaNLG benchmark, introducing a new dataset on dialogue generation in the process. Furthermore, using a clean corpus of 27.5 GB of Bangla data, we pretrain ‘BanglaT5’, a sequence-to-sequence Transformer language model for Bangla. BanglaT5 achieves state-of-the-art performance in all of these tasks, outperforming several multilingual models by up to 9% absolute gain and 32% relative gain. We are making the new dialogue dataset and the BanglaT5 model publicly available at <https://github.com/csebuethlp/BanglaNLG> in the hope of advancing future research on Bangla NLG.

1 Introduction

The emergence of pretrained language models (Devlin et al., 2019; Radford et al., 2019; Liu et al., 2019) has brought about a revolutionary change in natural language processing (NLP). With little task-specific fine-tuning, these models have achieved state-of-the-art results on many NLP tasks (Wang et al., 2018; Rajpurkar et al., 2016; Tjong Kim Sang and De Meulder, 2003). However, the focus of these models has predominantly been on natural language understanding (NLU). Even models pretrained with generative objectives (Raffel et al., 2020) concern themselves with NLU tasks more than natural language generation (NLG) tasks. Although there have been recent efforts to uplift NLG (Gehrmann et al., 2021), they are primarily geared towards high- and mid-resource languages. For example, despite being the sixth most spoken language in the world with over 230 million native speakers comprising 3% of the world’s total population,¹ Bangla has remained an underrepresented

¹<https://w.wiki/Psq>

language in the NLP literature (Joshi et al., 2020). There have been only a handful of benchmark studies on Bangla NLG (Dabre et al., 2022; Kumar et al., 2022), and that too without Bangla being the main focus. This can be attributed to the lack of diverse NLG tasks under a single benchmark and strong pretrained Bangla NLG models.

To this end, we present ‘BanglaNLG,’ a comprehensive benchmark for Bangla language generation comprising six representative tasks on machine translation, text summarization, question answering, dialogue generation, headline generation, and cross-lingual summarization. To our knowledge, BanglaNLG is the first NLG benchmark exclusively for a low-resource language.

To establish a strong baseline for this benchmark, we pretrain **BanglaT5** – a sequence-to-sequence Transformer model (Vaswani et al., 2017) pretrained on a 27.5 GB clean Bangla text corpus covering a broad range of domains. In summary:

- We develop the BanglaNLG benchmark bringing together six NLG tasks.
- We introduce a Multi-turn Dialogue dataset.
- We pretrain BanglaT5 and evaluate it on the six NLG tasks, showing strong results.

BanglaT5 outperforms similar-sized multilingual models, achieving new state-of-the-art results on three tasks with a 4% gain on average. We are releasing the BanglaT5 model and a live leaderboard to promote future research on Bangla NLG.

2 The Bangla Natural Language Generation (BanglaNLG) Benchmark

There have been sporadic works on Bangla NLG, mostly catered to machine translation (Hasan et al., 2020; Mumin et al., 2019a,b) and text summarization (Bhattacharjee et al., 2021b; Dhar et al., 2021). However, Bangla NLG lacks a unified study comprising diverse and challenging tasks. Motivated by the popular benchmarks like GLUE (Wang

Task	Corpus	Train	Dev	Test	Metric	Domain
Machine Translation	BanglaNMT, FLoRes	2,751,315	997	1,012	SacreBLEU	Misc.
Text Summarization	XL-Sum	8,102	1,012	1,012	ROUGE-2	BBC
Question Answering	BQA	127,771	2,502	2,504	EM/F1	Wikipedia
Multi-turn Dialogue	DailyDialog	76,052	7,069	6,640	BLEU-1	Misc.
News Headline Generation	XL-Sum	8,102	1,012	1,012	ROUGE-2	BBC
Cross-lingual Summarization	CrossSum	1241	153	155	ROUGE-2	BBC

Table 1: Dataset statistics and basic characteristics of BanglaNLG. Machine translation and cross-lingual summarization datasets include examples of Bangla ↔ English.

et al., 2018), XTREME (Hu et al., 2020), GEM (Gehrmann et al., 2021), that have facilitated the training/evaluation of NLP models, we establish the first-ever Bangla Natural Language Generation (BanglaNLG) Benchmark.

2.1 Task Selection Criteria

We consider the following factors while choosing the evaluation tasks:

1. Diversity: The tasks should focus on evaluating the model’s generalization capabilities. Therefore, they should vary in task nature – the input and output length, the type of generated text, the target domain, and the dataset size.

2. Practical Applicability: The choice of tasks should be driven by their practical implications. Rather than being used in abstract situations, NLG models trained on these tasks should be able to aid/reduce human effort in real-world scenarios.

3. Difficulty: The tasks should be challenging while not being unsolvable. There should be clear room for improvement to foster future research.

4. Accessibility: The selected datasets for these tasks should be openly accessible to encourage researchers to design better NLG models.

5. Evaluation: The selected tasks should have reliable automated metrics for evaluating the focused abilities of an NLG model.

2.2 Selected Tasks

Considering the criteria mentioned above, we design BanglaNLG as an aggregation of six tasks:

1. Machine Translation (MT): MT is perhaps the most studied NLG task in Bangla and the most commonly benchmarked NLG task in general. We use the BanglaNMT parallel corpus (Hasan et al., 2020), the largest Bangla-English MT dataset curated, with 2.75 million parallel pairs for training. The sentence pairs originate from various domains such as Wikipedia, news articles, religious and law

documents, etc. We evaluate the NLG models using FLoRes-100 (Goyal et al., 2022) in both directions on this dataset, i.e., Bangla to English and English to Bangla. This task is particularly challenging since it assesses an NLG model’s bilingual generation capabilities. Following standard practice, we use detokenized SacreBLEU (Post, 2018) as the evaluation metric for this task.

2. Text Summarization (TS): This task aims to generate a short and fluent summary given a long text document. We chose the Bangla portion of XL-Sum (Hasan et al., 2021) for this task. XL-Sum is a large comprehensive dataset for abstractive TS where the article and summaries are written by professional editors of BBC News. The articles cover various topics such as entertainment, politics, science, sports, etc. For this task, we use ROUGE-2² (Lin, 2004) as the evaluation metric.

3. Question Answering (QA): This is a fundamental NLP task that can be modeled as both an NLU and NLG task. We use the BQA (Bhattacharjee et al., 2022) dataset for this task. The training data is machine translated from SQuAD 2.0 (Rajpurkar et al., 2018), while the evaluation data come from the human-annotated question-answer pairs of the TyDi-QA (Clark et al., 2020) secondary gold passage task. Although TyDi-QA only contains answerable questions, BQA introduced unanswerable questions to make the task more challenging. Following SQuAD 2.0, we use Exact Match (EM) and F1 as the evaluation metrics.

4. Multi-turn Dialogue (MTD): Conversational AI is a crucial task for NLG (Chen et al., 2017). However, there is no public dataset for dialogue generation in Bangla. As such, we curate a new multi-turn dialogue dataset by translating the DailyDialog (Li et al., 2017) dataset using the English to Bangla translation model introduced by Hasan

²We use Bangla stemming supported ROUGE implementation from https://github.com/csebuetnlp/xl-sum/tree/master/multilingual_rouge_scoring.

et al. (2020). Unlike standard QA-style conversation datasets, DailyDialog reflects real-life conversations in various social situations rich in emotion, making it a perfect candidate for our benchmark. We automatically translate the training data following the same procedure described in [Bhattacharjee et al. \(2022\)](#) and have the evaluation sets manually translated by expert human translators. We use BLEU-1 as the evaluation metric for this task to properly differentiate between models since averaged BLEU scores of up to 4-gram tend to be quite low in dialogue evaluation ([Zhang et al., 2020](#)).

5. News Headline Generation (NHG): Automating headline generation can help news editors write compelling headlines to draw readers’ attention. We consider NHG as a complementary task to TS. Given an article, the objective is to generate an appropriate headline that accurately depicts the article. We repurpose the XL-Sum ([Hasan et al., 2021](#)) dataset for this task since it also includes the titles of the articles. Like TS, we use ROUGE-2 as the evaluation metric.

6. Cross-lingual Summarization (XLS): As another task for evaluating models’ bilingual generation capabilities, we consider XLS. In this task, given a piece of text in a source language, we have to generate the corresponding summary in a target language. This is potentially harder than both MT and TS considering it combines both in a single task. We consider the English-Bengali portion of the CrossSum ([Bhattacharjee et al., 2021a](#)) dataset for this task. It is curated by aligning identical articles written in different languages from the XL-Sum dataset. For evaluation, we use ROUGE-2.

We present detailed statistics of the BanglaNLG benchmark in Table 1.

3 BanglaT5

We introduce BanglaT5, a sequence-to-sequence Transformer model ([Vaswani et al., 2017](#)), to establish a strong baseline for BanglaNLG benchmark. In this section, we describe the pretraining data, objectives, and model architecture of BanglaT5.

3.1 Pretraining Data

We chose Bangla2B+ ([Bhattacharjee et al., 2022](#)) as the pretraining corpus for BanglaT5. This is a 27.5 GB dataset containing 5.25 million documents collected from a meticulously selected list of web sources. While larger sources like CCNet ([Wenzek et al., 2020](#)) and mC4 ([Xue et al., 2021](#)) are

available, these contain a lot of noise and offensive texts that are difficult to remove. For a generative model, even small amounts of unwanted texts in pretraining could lead to potentially dangerous biases in generated text ([Luccioni and Viviano, 2021](#)). Therefore, we decided not to use them.

3.2 Data Pre-processing

Following [Hasan et al. \(2020\)](#), we preprocessed the texts using their normalization pipeline³. We trained a SentencePiece ([Kudo and Richardson, 2018](#)) vocabulary of 32k subword tokens on the normalized corpus with a character coverage of 0.99995. While creating a training sample, we limited the maximum sequence length to 512 tokens for both input and output and discarded documents with a token count below 7. After tokenization, we had 4.8 million data points with an average sequence length of 402.32 tokens.

3.3 Pretraining Objective

For generative language modeling, two standard choices are decoder-only models ([Mikolov et al., 2010](#)) and encoder-decoder models ([Sutskever et al., 2014](#)). [Radford et al. \(2019\)](#) trained a decoder-only Transformer ([Vaswani et al., 2017](#)) pretrained on the conditional continuation objective. However, to provide more flexibility on generation and possible usage on understanding tasks, we only consider encoder-decoder models following the original design of the Transformer. They are generally trained with different denoising objectives to increase the encoder’s and decoder’s capacity. For instance, BART ([Lewis et al., 2020b](#)), and mBART ([Liu et al., 2020](#)) use a text-infilling-based objective. In contrast, MARGE ([Lewis et al., 2020a](#)) is a multilingual encoder-decoder model trained to reconstruct a document in one language by retrieving documents in other languages. Following [Raffel et al. \(2020\)](#), we pretrained BanglaT5 using a "span-correction" objective, empirically shown to be an optimal choice for encoder-decoder models. In this objective, consecutive spans of input tokens are replaced with a mask token, and the model is trained to reconstruct them.

3.4 Model Architecture & Hyperparameters

We pretrained the base variant of the T5 model: 12 layers, 12 attention heads, 768 hidden size, 2048 feed-forward size with GeGLU activation ([Shazeer,](#)

³<https://github.com/csebuetnlp/normalizer>

Model	Parameters	MT	TS	QA	MTD	NHG	XLS
mT5 (base)	582M	30.1/ 17.2	10.3	59.0/65.3	17.5	9.6	2.7/0.7
XLM-ProphetNet	616M	27.5/15.4	7.8	53.0/57.3	20.0	9.5	6.2/2.7
mBART-50	611M	29.7/15.5	10.4	53.4/58.9	18.5	11.2	5.4/ 3.7
IndicBART (unified)	244M	28.1/16.6	8.9	59.6/65.6	14.8	7.9	6.3/2.5
IndicBART (separate)	244M	27.5/15.7	9.2	55.3/61.2	14.1	9.1	5.3/2.4
BanglaT5	247M	31.3/17.4	13.7	68.5/74.8	19.0	13.8	6.4/4.0

Table 2: Performance comparison of the pretrained models on different BanglaNLG tasks. Scores in bold texts have statistically significant ($p < 0.05$) difference from others with bootstrap sampling (Koehn, 2004).

2020) with a batch size of 65536 tokens for 3 million steps on a v3-8 TPU instance on GCP. We used the Adam (Kingma and Ba, 2015) optimizer with a $3e-4$ learning rate, linear warmup of 10k steps, and ‘inverse square root’ learning rate decay.

4 Experiments & Results

We compared BanglaT5 it with four multilingual models: mT5 (base) (Xue et al., 2021), mBART-50 (Tang et al., 2020), XLM-ProphetNet (Qi et al., 2021), and IndicBART (both unified and separate script variants) (Dabre et al., 2022).⁴ All pretrained models were fine-tuned for 3-15 epochs with batch size 32 (128 for MT). We used linear warmup with a ratio of 0.1, label smoothing of 0.1 (Szegedy et al., 2016), and weight decay of $1e-6$ with the Adam optimizer (Kingma and Ba, 2015). The learning rate was tuned from the set $\{5e-5, 1e-4, 5e-4\}$. The best model was evaluated based on the validation performance after each epoch.

During inference, we used beam-search (Hayes-Roth et al., 1976) with beam size 5 (on all tasks except QA), removed duplicated trigrams during beam search (Fan et al., 2018), and used a length penalty (Wu et al., 2016) of 0.6. For QA, we used greedy decoding, i.e., picking the most probable token during each decoding step.

The evaluation results are presented in Table 2. In all the tasks, BanglaT5 outperformed all multilingual models by a considerable margin, on average 4% over the second-best, mT5. In all monolingual tasks except MTD, BanglaT5 achieves a big performance gain over others (up to 9.54% in QA), which can be attributed to the quality of the pretraining data. In MD, BanglaT5 lags marginally behind XLM-ProphetNet. We hypothesize this is due to the lack of colloquial data in Bangla2B+ since Bhat-tacharjee et al. (2022) left out such sources to avoid

⁴Due to computational budget limitations, we do not benchmark on billion-parameter models like large mT5 variants.

toxic and biased conversations.

We find the MT results particularly interesting, where BanglaT5 outperforms larger multilingual models in both directions. This suggests that despite having very little English data in the pretraining corpus, BanglaT5 can generalize well to a new translation language, given high-quality fine-tuning data. We explore this more in the Appendix. Conspicuously, all the models achieve relatively poor scores on the XLS task. This can be attributed to the smaller amount of training data.

BanglaT5 proves its superiority in compute and memory efficiency along with its performance due to its smaller size (less than half the parameters of all multilingual models except IndicBART). In practice, we observe 2-2.5x faster training and inference times with BanglaT5 than these larger multilingual models.

5 Related Works

Pretrained models NLP has witnessed a sea of change with the advent of pretrained language models like ULMfit (Howard and Ruder, 2018), ELMo (Peters et al., 2018), and most notably BERT (Devlin et al., 2019), achieving state-of-the-art results in many NLU benchmarks. Besides these NLU models, more and more pretrained models designed for NLG tasks have been proposed. Rothe et al. (2020) adopted pretrained NLU model checkpoints for generative tasks. GPT-2 (Radford et al., 2019), and later GPT-3 (Brown et al., 2020) showed that pretrained generative language models can perform remarkably well in zero-shot transfer tasks. More recently, Qi et al. (2020) proposed ProphetNet, which introduces the future n-gram prediction mechanism for language generation. Dabre et al. (2022) introduced IndicBART, which is pretrained on 11 Indic languages, including Bangla.

NLG Benchmarks Recently, many multi-task benchmarks have been proposed to drive the

progress of NLG models. [Moussallem et al. \(2020\)](#) proposed the BENG benchmark for NLG and knowledge extraction. GLGE ([Liu et al., 2021](#)) is a similar benchmark with a different set of tasks and difficulty levels. However, these benchmarks are limited to English only. [Gehrmann et al. \(2021\)](#) introduced the GEM benchmark for various tasks such as summarization ([Narayan et al., 2018](#)), data-to-text generation ([Nan et al., 2021](#)) across different languages. [Cahyawijaya et al. \(2021\)](#) introduced different tasks and baselines for 3 Indonesian languages. More recently, [Kumar et al. \(2022\)](#) introduced IndicNLG, a benchmark with five tasks in 11 Indic languages, including Bangla.

6 Conclusion & Future Works

NLP research in low-resource languages is lagging behind due to the lack of reliable benchmarks and datasets. To facilitate the development, evaluation, and comparison of new NLG models, we introduced a multi-task evaluation benchmark for Bangla NLG, a widely spoken yet low-resource language. We presented BanglaT5, a pretrained NLG model in Bangla, setting new state-of-the-art results with BanglaT5. We strongly believe that our contributions in this work will help the Bangla NLP community benchmark NLG tasks more easily under a unified setup.

In future work, we plan to introduce new tasks to BanglaNLG, such as personalized dialogue generation ([Zhang et al., 2018](#)), conversational question-answering ([Reddy et al., 2019](#)). We will also add more recent multilingual models to our comparison to BanglaT5, e.g., DeltaLM ([Ma et al., 2021](#)).

Limitations

Although [Bhattacharjee et al. \(2022\)](#) claimed that Bangla2B+, the pretraining corpus for BanglaT5, had been carefully filtered for offensive or unwanted texts, they alerted that there might be small amounts of these contents may be present, which can result in bias or toxicity in the pretrained model. We, therefore, recommend using BanglaT5 with caution, especially for real-world deployment.

Ethics Statement

License The TyDiQA dataset ([Clark et al., 2020](#)) is released under the Apache License 2.0, allowing modifications and distribution. All other pretraining and fine-tuning datasets are released under the

Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0), which allows modifications and distributions for non-commercial research purposes. We strictly adhere to these licenses and will release BanglaT5 and BanglaNLG benchmark resources under CC BY-NC-SA 4.0.

Annotation Expert translators who provide translation services for renowned Bangla newspapers were hired to translate the evaluation sets of the dialogue dataset. Each translated sentence was further assessed for quality by another expert. It was again translated by the original translator if found to be of low quality. If the re-translation was found to be of low quality, it was then translated by the other expert. The experts were paid hourly as per standard rates in local currency.

Hallucinated Text It is well-known that text generation models can hallucinate outputs that may not necessarily be faithful to the original input ([Maynez et al., 2020](#)). Though the texts may be fluent and human-like, the hallucinations may be factually inconsistent and impact the outputs negatively. BanglaT5 may be susceptible to the same kinds of hallucinations.

Carbon Footprint We avoided using large models for pretraining and fine-tuning, reducing their environmental impacts. BanglaT5 was trained for about 30 days on Google v3 TPUs. Google’s TPUs are specifically designed for machine learning, which makes them up to five times more efficient than GPUs. Assuming 0.080kg carbon emission per kWh,⁵ the pretraining would emit fewer than 100kg carbon into the environment, far below most computationally demanding models. All fine-tuning experiments were done on a desktop machine with an 8-core Intel Core-i7 11700k CPU and NVIDIA RTX 3090 GPU, and no single run except machine translation took more than 12 hours, which amounts to fewer than 0.5kg carbon emission. On average, machine translation runs took three days each, emitting less than 3kg of carbon.

Acknowledgements

We would like to thank the Research and Innovation Centre for Science and Engineering (RISE), BUET, for funding the project and Google TPU Research Cloud (TRC) program for providing cloud support.

⁵<https://blog.google/technology/ai/minimizing-carbon-footprint/>

References

- Abhik Bhattacharjee, Tahmid Hasan, Wasi Uddin Ahmad, Yuan-Fang Li, Yong bin Kang, and Rifat Shahriyar. 2021a. [Crosssum: Beyond english-centric cross-lingual abstractive text summarization for 1500+ language pairs](#). *CoRR*, abs/2112.08804.
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Ahmad Uddin, Kazi Mubasshir, Md. Saiful Islam, Anindya Iqbal, M. Sohel Rahman, and Rifat Shahriyar. 2022. [BanglaBERT: Language model pretraining and benchmarks for low-resource language understanding evaluation in bangla](#). In *Findings of the North American Chapter of the Association for Computational Linguistics: NAACL 2022*.
- Prithwiraj Bhattacharjee, Avi Mallick, Md. Saiful Islam, and Marium-E-Jannat. 2021b. [Bengali abstractive news summarization \(bans\): A neural attention approach](#). In *Proceedings of International Conference on Trends in Computational and Cognitive Engineering*, pages 41–51, Singapore. Springer Singapore.
- Terra Blevins and Luke Zettlemoyer. 2022. Language contamination explains the cross-lingual capabilities of english pretrained models. *arXiv preprint arXiv:2204.08110*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Samuel Cahyawijaya, Genta Indra Winata, Bryan Wilie, Karissa Vincentio, Xiaohong Li, Adhiguna Kuncoro, Sebastian Ruder, Zhi Yuan Lim, Syafri Bahar, Masayu Khodra, Ayu Purwarianti, and Pascale Fung. 2021. [IndoNLG: Benchmark and resources for evaluating Indonesian natural language generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8875–8898, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hongshen Chen, Xiaorui Liu, Dawei Yin, and Jiliang Tang. 2017. A survey on dialogue systems: Recent advances and new frontiers. *Acm Sigkdd Explorations Newsletter*, 19(2):25–35.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. [TyDi QA: A benchmark for information-seeking question answering in typologically diverse languages](#). *Transactions of the Association for Computational Linguistics*, 8:454–470.
- Raj Dabre, Himani Shrotriya, Anoop Kunchukuttan, Ratish Puduppully, Mitesh Khapra, and Pratyush Kumar. 2022. [IndicBART: A pre-trained model for indic natural language generation](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1849–1863, Dublin, Ireland. Association for Computational Linguistics.
- 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.
- Nobel Dhar, Gaurob Saha, Prithwiraj Bhattacharjee, Avi Mallick, and Md Saiful Islam. 2021. [Pointer over attention: An improved bangla text summarization approach using hybrid pointer generator network](#). In *2021 24th International Conference on Computer and Information Technology (ICCIT)*, pages 1–5.
- Angela Fan, David Grangier, and Michael Auli. 2018. [Controllable abstractive summarization](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 45–54, Melbourne, Australia. Association for Computational Linguistics.
- Sebastian Gehrmann, Tosin Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Anuoluwapo Aremu, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna-Adriana Clinciu, Dipanjan Das, Kaustubh Dhole, Wanyu Du, Esin Durmus, Ondřej Dušek, Chris Chinenye Emezue, Varun Gangal, Cristina Garbacea, Tatsunori Hashimoto, Yufang Hou, Yacine Jernite, Harsh Jhamtani, Yangfeng Ji, Shailza Jolly, Mihir Kale, Dhruv Kumar, Faisal Ladhak, Aman Madaan, Mounica Maddela, Khyati Mahajan, Saad Mahamood, Bodhisattwa Prasad Majumder, Pedro Henrique Martins, Angelina McMillan-Major, Simon Mille, Emiel van Miltenburg, Moin Nadeem, Shashi Narayan, Vitaly Nikolaev, Andre Niyongabo Rubungo, Salomey Osei, Ankur Parikh, Laura Perez-Beltrachini, Niranjana Ramesh Rao, Vikas Raunak, Juan Diego Rodriguez, Sashank Santhanam, João Sedoc, Thibault Sellam, Samira Shaikh, Anastasia Shimorina, Marco Antonio Sobrevilla Cabezero, Hendrik Strobelt, Nishant Subramani, Wei Xu, Diyi Yang, Akhila Yerukola, and Jiawei Zhou. 2021. [The GEM benchmark: Natural language generation, its evaluation and metrics](#). In *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM 2021)*, pages 96–120, Online. Association for Computational Linguistics.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. [The Flores-101 evaluation](#)

- benchmark for low-resource and multilingual machine translation. *Transactions of the Association for Computational Linguistics*, 10:522–538.
- 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.
- P Hayes-Roth, M Fox, G Gill, DJ Mostow, and R Reddy. 1976. Speech understanding systems: Summary of results of the five-year research effort. *Carnegie-Mellon University, Computer Science Department Interim Report*.
- Jeremy Howard and Sebastian Ruder. 2018. **Universal language model fine-tuning for text classification**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. **XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation**. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.
- 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.
- Diederik P. Kingma and Jimmy Ba. 2015. **Adam: A method for stochastic optimization**. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Philipp Koehn. 2004. **Statistical significance tests for machine translation evaluation**. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. **SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. 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**.
- Mike Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. 2020a. **Pre-training via paraphrasing**. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20*, Red Hook, NY, USA. Curran Associates Inc.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020b. **BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. **DailyDialog: A manually labelled multi-turn dialogue dataset**. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- 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.
- Dayiheng Liu, Yu Yan, Yeyun Gong, Weizhen Qi, Hang Zhang, Jian Jiao, Weizhu Chen, Jie Fu, Linjun Shou, Ming Gong, Pengcheng Wang, Jiusheng Chen, Daxin Jiang, Jiancheng Lv, Ruofei Zhang, Winnie Wu, Ming Zhou, and Nan Duan. 2021. **GLGE: A new general language generation evaluation benchmark**. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 408–420, Online. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. **Multilingual denoising pre-training for neural machine translation**. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- 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*.

- Alexandra Luccioni and Joseph Viviano. 2021. [What’s in the box? an analysis of undesirable content in the Common Crawl corpus](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 182–189, Online. Association for Computational Linguistics.
- Shuming Ma, Li Dong, Shaohan Huang, Dongdong Zhang, Alexandre Muzio, Saksham Singhal, Hany Hassan Awadalla, Xia Song, and Furu Wei. 2021. [Deltalm: Encoder-decoder pre-training for language generation and translation by augmenting pretrained multilingual encoders](#). *CoRR*, abs/2106.13736.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. *Interspeech*, 2(3):1045–1048.
- Diego Moussallem, Paramjit Kaur, Thiago Ferreira, Chris van der Lee, Anastasia Shimorina, Felix Conrads, Michael Röder, René Speck, Claire Gardent, Simon Mille, Nikolai Ilinykh, and Axel-Cyrille Ngonga Ngomo. 2020. [A general benchmarking framework for text generation](#). In *Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+)*, pages 27–33, Dublin, Ireland (Virtual). Association for Computational Linguistics.
- Mohammad Abdullah Al Mumin, Md Hanif Seddiqui, Muhammed Zafar Iqbal, and Mohammed Jahirul Islam. 2019a. [Neural machine translation for low-resource english-bangla](#). *Journal of Computer Science*, 15(11):1627–1637.
- Mohammad Abdullah Al Mumin, Md Hanif Seddiqui, Muhammed Zafar Iqbal, and Mohammed Jahirul Islam. 2019b. [shu-torjoma: An english-bangla statistical machine translation system](#). *Journal of Computer Science*, 15(7):1022–1039.
- Linyong Nan, Dragomir Radev, Rui Zhang, Amrit Rau, Abhinand Sivaprasad, Chiachun Hsieh, Xiangu Tang, Aadit Vyas, Neha Verma, Pranav Krishna, Yangxiaokang Liu, Nadia Irwanto, Jessica Pan, Faiaz Rahman, Ahmad Zaidi, Mutethia Mutuma, Yasin Tarabar, Ankit Gupta, Tao Yu, Yi Chern Tan, Xi Victoria Lin, Caiming Xiong, Richard Socher, and Nazneen Fatema Rajani. 2021. [DART: Open-domain structured data record to text generation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 432–447, Online. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. 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.
- Weizhen Qi, Yeyun Gong, Yu Yan, Can Xu, Bolun Yao, Bartuer Zhou, Biao Cheng, Daxin Jiang, Jiusheng Chen, Ruofei Zhang, Houqiang Li, and Nan Duan. 2021. [ProphetNet-X: Large-scale pre-training models for English, Chinese, multi-lingual, dialog, and code generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 232–239, Online. Association for Computational Linguistics.
- Weizhen Qi, Yu Yan, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. 2020. [ProphetNet: Predicting future n-gram for sequence-to-SequencePre-training](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2401–2410, Online. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*, 1(8).
- 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*, 21:1–67.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for](#)

- machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. **CoQA: A conversational question answering challenge**. *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Sascha Rothe, Shashi Narayan, and Aliaksei Severyn. 2020. **Leveraging pre-trained checkpoints for sequence generation tasks**. *Transactions of the Association for Computational Linguistics*, 8:264–280.
- Noam Shazeer. 2020. **Glu variants improve transformer**.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. **Sequence to sequence learning with neural networks**. In *Proceedings of the 27th International Conference on Neural Information Processing Systems (NIPS 2014)*, pages 3104–3112, Montreal, Canada.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. **Rethinking the inception architecture for computer vision**. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. **Multilingual translation with extensible multilingual pretraining and finetuning**. *arXiv preprint arXiv:2008.00401*.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. **Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition**. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. **Attention is all you need**. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS 2017)*, page 6000–6010, Long Beach, California, USA.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. **GLUE: A multi-task benchmark and analysis platform for natural language understanding**. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. 2020. **CCNet: Extracting high quality monolingual datasets from web crawl data**. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4003–4012, Marseille, France. European Language Resources Association.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. **Google’s neural machine translation system: Bridging the gap between human and machine translation**. *CoRR*, abs/1609.08144.
- 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.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. **Personalizing dialogue agents: I have a dog, do you have pets too?** In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. **DIALOGPT: Large-scale generative pre-training for conversational response generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.

Supplementary Material: Appendices

A Multi-turn Dialogue Scores

In Table 3, we mention BLEU-1, BLEU-2, BLEU-3, and BLEU-4 scores for different models in the multi-turn dialogue generation task.

Model	B-1	B-2	B-3	B-4
mT5 (base)	17.54	3.67	1.25	0.43
XLNet-ProphetNet	19.98	6.06	2.98	1.86
mBART-50	18.54	5.56	2.97	2.09
IndicBART (unified)	14.75	3.18	1.06	0.37
IndicBART (separate)	14.05	3.23	1.18	0.49
BanglaT5	19.00	5.02	2.04	0.92

Table 3: Performance comparison of the pretrained models on the dialogue generation task. Scores in bold texts have statistically significant ($p < 0.05$) difference from others with bootstrap sampling (Koehn, 2004).

B Cross-lingual Capabilities of BanglaT5

Despite being a monolingual model pretrained on heavily filtered Bangla data, BanglaT5 exhibits strong cross-lingual abilities, particularly in the machine translation (MT) task. In addition to the quality and size of the fine-tuning dataset, this performance can also be attributed to the presence of a significant amount of non-Bangla tokens ($\sim 10.3\%$) in the BanglaT5 vocabulary.

Since Bhattacharjee et al. (2022) curated the Bangla2B+ corpus by document-level language filtering, these documents preserve foreign text sequences occurring in the Bangla documents. We deliberately maintain these tokens while training the vocabulary of BanglaT5, using a relatively high character coverage. Our rationale behind doing this was to capture code-switching and allow better generalization across languages co-occurring with Bangla, as well as romanized forms of Bangla texts during fine-tuning, which is reflected in the MT results. However, it should be noted that the quality and size of fine-tuning data are essential for a strong cross-lingual performance since the mere existence of foreign tokens in the vocabulary is not enough to produce meaningful generation performance, as demonstrated by the poor performance in the cross-lingual summarization (XLS) task.

This phenomenon has been studied in-depth by Blevins and Zettlemoyer (2022) in the context of pretrained language models in English, where they showed that these models develop strong

cross-lingual transfer capabilities due to the non-negligible amount of foreign text present in the pretraining data and robustness to UNK tokens during fine-tuning.