

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

Abhik Bhattacharjee<sup>1</sup>, Tahmid Hasan<sup>1</sup>, Wasi Uddin Ahmad<sup>2</sup>, Rifat Shahriyar<sup>1</sup>

Bangladesh University of Engineering and Technology (BUET)<sup>1</sup>,

University of California, Los Angeles<sup>2</sup>

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 in the web domain. We aggregate three challenging conditional text generation tasks under the BanglaNLG benchmark. Then, using a clean corpus of 27.5 GB of Bangla data, we pretrain BanglaT5, a sequence-to-sequence Transformer model for Bangla. BanglaT5 achieves state-of-the-art performance in all of these tasks, outperforming mT5 (base) by up to 5.4%. We are making the BanglaT5 language model and a leaderboard publicly available in the hope of advancing future research and evaluation on Bangla NLG. The resources can be found at <https://github.com/csebuetnlp/BanglaNLG>.

## 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. With little task-specific fine-tuning, these models can achieve state-of-the-art results on many natural language processing tasks (Wang et al., 2018; Rajpurkar et al., 2016; Tjong Kim Sang and De Meulder, 2003). However, most of the focus of these pretrained models has 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 300 million native speakers comprising 4% of the world’s total population,<sup>1</sup> there has not been any comprehensive

study on Bangla NLG. This can be attributed to the lack of a diverse set of NLG tasks under a single benchmark and strong pretrained Bangla NLG models.

To overcome these deficiencies, we present ‘BanglaNLG’, a benchmark for Bangla language generation comprising three representative tasks on machine translation, abstractive text summarization and question answering. BanglaNLG is the first-ever NLG benchmark for a low-resource, to the best of our knowledge.

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 three NLG tasks.
- We pretrain BanglaT5 and evaluate it on NLG tasks, showing strong results.

BanglaT5 outperforms its multilingual counterpart mT5 (base), achieving new state-of-the-art results on three tasks with a 5.4% average gain over mT5. We are releasing the BanglaT5 model and a live leaderboard to promote future research on Bangla NLG.

## 2 The Bangla Natural Language Generation (BNLG) Benchmark

There have been prior sporadic works on Bangla NLG, mostly catered to Machine Translation (Hasan et al., 2020a; Mumin et al., 2019a,b) and Text Summarization (Bhattacharjee et al., 2020; Dhar et al., 2021). However, Bangla NLG still lacks a unified study comprising diverse and challenging tasks. To this end, we establish the first-ever Bangla Natural Language Generation Benchmark (BNLG). When selecting the evaluation tasks for BNLG, we take the following considerations into account:

<sup>1</sup><https://w.wiki/Psq>

Task	Corpus	Train	Dev	Test	Metric	Domain
Machine Translation (bn ↔ en)	BanglaNMT	2,751,315	598	1,000	SacreBLEU	Misc.
Abstractive Text Summarization	XLSum	8,102	1,012	1,012	ROUGE-2	BBC
Question Answering	TyDiQA	127,771	2,502	2,504	F1/EM	Wiki.

Table 1: Dataset statistics and basic characteristics of BNLG

**Task diversity** The tasks should focus on evaluating the generalization capabilities of an NLG model. Therefore, they should vary in task nature – the length of input and output, the type of generated text, target domain, and the size of the dataset.

**Task difficulty** The tasks should be reasonably challenging while not being unsolvable. In addition, models evaluated on these tasks should be able to compete with human performance.

**Reliable evaluation** The selected tasks should have reliable automated evaluation metrics to assess the quality of the generated text. Therefore, we do not consider weakly conditioned tasks such as answer-agnostic question generation, where there are many plausible outputs for a given input.

**Quality and availability** The selected datasets for these tasks should meet the minimum quality standards and be accessible to encourage researchers to design better NLG models.

Considering the above, we design BNLG as an aggregation of three tasks, namely, *Machine Translation*, *Abstractive Text Summarization*, and *Question Answering*. We briefly describe them below:

**1. Machine Translation.** Machine translation is perhaps the most studied NLG task in Bangla and the most commonly benchmarked NLG task in general. For this task, we use the BanglaNMT parallel corpus introduced by (Hasan et al., 2020a). It is the largest Bangla-English machine translation dataset curated, with 2.75 million parallel pairs as the training data. The sentence pairs originate from various domains such as Wikipedia, news articles, religious and law documents, etc. We evaluate the NLG models in both directions on this dataset, i.e., English to Bangla and Bangla to English. This 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. Abstractive Text Summarization.** 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., 2021b) for this task.

XL-Sum is a large comprehensive dataset for abstractive text summarization where the article and summaries are annotated by professional annotators of BBC. The news articles cover various topics such as entertainment, science, technology, sports, etc. For this task, we use ROUGE-2 (Lin, 2004)<sup>2</sup> as the evaluation metric.

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

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

### 3 BanglaT5

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 data 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 of Bangla data dumps 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 opted not to use them.

<sup>2</sup>We use the rouge implementation of [https://github.com/csebuetnlp/xl-sum/tree/master/multilingual\\_rouge\\_scoring](https://github.com/csebuetnlp/xl-sum/tree/master/multilingual_rouge_scoring), since it supports Bangla stemming

Model	Params.	MT	ATS	QA	BNLG Score
mT5 (base)	582M	36.6/22.5	10.27	58.95/65.32	38.73
BanglaT5	247M	<b>38.8/25.2</b>	<b>13.66</b>	<b>68.49/74.77</b>	<b>44.18</b>

Table 2: Performance comparison on different downstream tasks

### 3.2 Pre-processing

Following Hasan et al. (2020b), we preprocessed the text using their normalization pipeline. We trained a SentencePiece (Kudo and Richardson, 2018) vocabulary of 32000 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. 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 language models (Mikolov et al., 2010) and encoder-decoder models (Sutskever et al., 2014). (Radford et al., 2019) train a decoder-only left-to-right Transformer language model 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 Transformer (Vaswani et al., 2017). These models are trained typically by maximizing the likelihood of the target output given an input. To increase the capacity of both the encoder and decoder, they are generally trained with different denoising objectives. For instance, BART (Lewis et al., 2020b), and mBART (Liu et al., 2020) use a text infilling based denoising objective, whereas MARGE (Lewis et al., 2020a) is a multilingual encoder-decoder model that is trained to reconstruct a document in one language by retrieving documents in other languages.

Following (Raffel et al., 2020), we pretrain BanglaT5 using a masked language modeling "span-correction" denoising objective, which has been 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 the masked-out tokens.

### 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, 2020) in the feed-forward layer for both the encoder and decoder) 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 and linear warmup of 10k steps, and 'inverse square root' learning rate decay.

## 4 Experiments & Results

We fine-tuned BanglaT5 on the selected tasks of BNLG and compared it with its multilingual counterpart mT5 (base) (Xue et al., 2021) trained on 101 languages.

All pretrained models were fine-tuned for 3-15 epochs on each task with batch size 32-128. We used linear warmup with a ratio of 0.1, label smoothing of 0.1, and weight decay of  $1e-6$  with the Adam optimizer. The best model was evaluated based on the validation performance after each epoch.

During inference, we used beam-search (Och and Ney, 2004) 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 use greedy decoding, i.e., picking the most probable token during each decoding step.

The results on different downstream tasks are detailed in Table 2. In all the tasks, BanglaT5 outperformed mT5 by a considerable margin of 5.45%, achieving an average score of 44.18. In monolingual tasks, as expected, BanglaT5 achieves a big performance gain over mT5 (up to 9.54% in the QA task), which can be attributed to the quality of the pretraining data. However, we find the Machine Translation results particularly interesting, where BanglaT5 outperforms the multilingual mT5 in both (bn→en) and (en→bn) by 2.2% and 2.7%, respectively. This suggests that despite having very little English data in the pretraining corpus, BanglaT5 can generalize well to a new translation language, given quality fine-tuning data. Furthermore, BanglaT5 is superior in performance and substantially compute- and memory-efficient due to its smaller size (less than half the parameters of mT5). In practice, we observe 2-2.5x faster train-

ing and inference times with BanglaT5 compared to mT5.

## 5 Related Works

**Pretrained models** Natural language processing 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). These models achieved state-of-the-art results in many NLU benchmarks. Besides these NLU models, more and more pretraining-based models designed for NLG tasks have been proposed. (Rothe et al., 2020) adopt pretrained NLU model checkpoints for generative tasks. GPT-2 (Radford et al., 2019) and later GPT-3 (Brown et al., 2020) show 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 and English.

**NLG Benchmarks** Recently, there have been a lot of multi-task NLG benchmarks proposed to drive the progress of generalizable models. Mousallem et al. (2020) proposed BENG as a benchmarking platform for natural language generation and knowledge extraction system. 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 data only. Gehrmann et al. (2021) introduced the GEM benchmark for various tasks such as summarization (Hasan et al., 2021b), data-to-text generation (Nan et al., 2021) across different languages. Cahyawijaya et al. (2021) introduced different tasks and baseline models for 3 Indonesian languages. More recently, (Kumar et al., 2022) introduced Indic-NLG, a benchmark with five tasks spanning 11 Indic languages, including Bangla.

## 6 Conclusion & Future Works

A detailed review of recent literature reveals that NLP research in low-resource 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 generalizable NLG model

in Bangla, setting new state-of-the-art results on all tasks with BanglaT5. We strongly believe that our contributions in this work will help Bangla NLP community benchmark NLG tasks more easily under a unified setup.

In future work, we plan to introduce new tasks to BNLG, such as personalized dialogue generation (Zhang et al., 2018), conversational question-answering (Reddy et al., 2019), cross-lingual summarization (Hasan et al., 2021a). We will also add more recent multilingual models to our comparison to BanglaT5 such as mBART-50 (Tang et al., 2020) and DeltaLM (Ma et al., 2021).

## References

- 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. 2020. [Bengali abstractive news summarization\(bans\): A neural attention approach](#). *CoRR*, abs/2012.01747.
- 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.
- 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**. *CoRR*, abs/2111.10269.
- 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, Niranjan 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.
- Tahmid Hasan, Abhik Bhattacharjee, 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.
- Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, and Rifat Shahriyar. 2021b. **XLsum: 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. 2020a. **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.
- Tahmid Hasan, Abhik Bhattacharjee, Kazi Samin, Masum Hasan, Madhusudan Basak, M. Sohel Rahman, and Rifat Shahriyar. 2020b. **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.
- 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.
- 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*.
- 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.
- 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.
- Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *Interspeech*, volume 2, pages 1045–1048. Makuhari.
- Diego Moussallem, Paramjot 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, Xiangru 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.
- Franz Josef Och and Hermann Ney. 2004. The alignment template approach to statistical machine translation. *Computational linguistics*, 30(4):417–449.
- 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, 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.
- 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](#).
- 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.