

# GLGE: A New General Language Generation Evaluation Benchmark

Dayiheng Liu<sup>♦\*</sup> Yu Yan<sup>†♥</sup> Yeyun Gong<sup>†♦</sup> Weizhen Qi<sup>♦</sup> Hang Zhang<sup>♦</sup>  
Jian Jiao<sup>♥</sup> Weizhu Chen<sup>‡</sup> Jie Fu<sup>◇</sup> Linjun Shou<sup>♦</sup> Ming Gong<sup>♦</sup>  
Pengcheng Wang<sup>♥</sup> Jiusheng Chen<sup>♥</sup> Daxin Jiang<sup>♦</sup> Jiancheng Lv<sup>♦</sup>  
Ruofei Zhang<sup>♥</sup> Winnie Wu<sup>♥</sup> Ming Zhou<sup>♦</sup> Nan Duan<sup>♦</sup>

<sup>♦</sup>College of Computer Science, Sichuan University <sup>♦</sup>Microsoft Research Asia <sup>◇</sup>Mila

<sup>‡</sup>Microsoft Azure AI <sup>♦</sup>Microsoft Search Technology Center Asia <sup>♥</sup>Microsoft

glge@microsoft.com <sup>†</sup>Equal contribution

## Abstract

Multi-task benchmarks such as GLUE and SuperGLUE have driven great progress of pre-training and transfer learning in Natural Language Processing (NLP). These benchmarks mostly focus on a range of Natural Language Understanding (NLU) tasks, without considering the Natural Language Generation (NLG) models. In this paper, we present the **General Language Generation Evaluation (GLGE)**, a new multi-task benchmark for evaluating the generalization capabilities of NLG models across eight language generation tasks. For each task, we continue to design three sub-tasks in terms of task difficulty (**GLGE-Easy**, **GLGE-Medium**, and **GLGE-Hard**). This introduces 24 subtasks to comprehensively compare model performance. To encourage research on pretraining and transfer learning on NLG models, we make GLGE publicly available and build a leaderboard with strong baselines including MASS, BART, and ProphetNet<sup>1</sup>.

## 1 Introduction

Pretrained language models, such as BERT (Devlin et al., 2019) and other advanced pretrained models (Raffel et al., 2020; Yang et al., 2019; Liu et al., 2019; Alberti et al., 2019; Brown et al., 2020; Clark et al., 2020) have made great progress in a host of Natural Language Understanding (NLU) tasks. Meanwhile, the development of general evaluation benchmarks has also helped drive the progress of these models. These benchmarks usually use an overall score to evaluate the performance of models across a wide range of NLU tasks. In addition to GLUE (Wang et al., 2019b) and SuperGLUE (Wang et al., 2019a) which are general

language understanding evaluation benchmarks for English, several general language understanding evaluation benchmarks for other languages are proposed, such as CLUE (Xu et al., 2020) for Chinese, FLUE (Le et al., 2020) for French, and IndoNLU (Wilie et al., 2020) for Indonesian. Furthermore, the multilingual multi-task benchmarks such as XTREME (Hu et al., 2020) and XGLUE (Liang et al., 2020) are proposed for cross-lingual evaluation.

In addition to NLU tasks, an increasing number of pretrained language models designed for Natural Language Generation (NLG) tasks have recently been proposed, such as MASS (Song et al., 2019), BERT-share (Rothe et al., 2020), BART (Lewis et al., 2020), ProphetNet (Qi et al., 2020), and ERINE-GEN (Xiao et al., 2020). However, the generalization capabilities of the language generation of these models are usually evaluated with different tasks, datasets, and metrics, which cannot provide a coherent and comprehensive evaluation. Although there are several general evaluation benchmarks as we mentioned above, none of them are particularly designed for general language generation evaluation.

To fill the gap of the NLG evaluation benchmark, we introduce the **General Language Generation Evaluation (GLGE)** benchmark, a new multi-task benchmark for evaluating the generalization capabilities of NLG in English language. It contains eight English language generation tasks, covering text summarization, question generation, generative question answering, and dialogue. We select six pre-existing popular datasets and introduce two new datasets selected from real-world scenarios. Moreover, in order to provide more diversified difficulty challenges, we employ two simple but effective strategies to build three NLG evaluation benchmarks (called **GLGE-Easy**, **GLGE-Medium**, and **GLGE-Hard**) in terms of task difficulty.

\* Work is done during internship at Microsoft Research Asia.

<sup>1</sup>The source code and dataset are publicly available at <https://github.com/microsoft/glge>.

To better understand the challenges posed by GLGE, we conduct experiments with existing widely used non-pretrained models (e.g., vanilla LSTM Seq2Seq (Bahdanau et al., 2015), vanilla Transformer (Vaswani et al., 2017)), and pre-trained models (e.g., MASS (Song et al., 2019), BART (Lewis et al., 2020), and ProphetNet (Qi et al., 2020)). We further analyze the n-gram diversity of the output samples. The experimental results show that there is a large performance gap between the pretrained models and the non-pretrained models. However, on the GLGE-hard task, the performance of the pretrained models still has great room for improvement.

In summary, the contributions of this work are five-fold: (1) a new multi-task NLG evaluation benchmark consisting of eight distinct datasets across four kinds of typical NLG tasks, (2) three NLG evaluation benchmarks of different difficulty levels, (3) standardized evaluation metrics and scripts for model evaluation and comparison, (4) open-sourced baselines and a public leaderboard<sup>2</sup> for the benchmark, (5) a thorough comparative study on existing widely used non-pretrained models and pretrained models with a detailed analysis of the results.

## 2 GLGE Benchmark

### 2.1 Design Principles

For the GLGE benchmark, we design and select the NLG tasks based on the following principles:

#### 2.1.1 Task Diversity

The tasks in GLGE focus on evaluating the generalization capabilities of a NLG model, varying the task, the length of the input text, the length of the output text, the type of generated text, and the size of the dataset.

#### 2.1.2 Task Difficulty

The tasks in GLGE should be challenging but solvable, which can encourage researchers to design better NLG models. Furthermore, we aim to provide benchmarks of different difficulty levels like GLUE (Wang et al., 2019b) and SuperGLUE (Wang et al., 2019a), which allows researchers to comprehensively evaluate the models. Researchers can also select the benchmark with moderate difficulty according to the size of the

<sup>2</sup><https://microsoft.github.io/glge/>.

model and the scale of the used pretraining corpus for comparison.

#### 2.1.3 Ease of Evaluation

The tasks in GLGE should be easily evaluated automatically. For some unconditional, open-ended, and weak conditional language generation tasks (e.g., answer-agnostic question generation, single-turn chit-chat response generation, and story generation), reasonable generation results are diverse. Due to the limited number of references in the automatic evaluation of text generation tasks, it is more difficult for automatic evaluation of those tasks. Therefore, instead of selecting unconditional and weak conditional language generation tasks, we tend to select language generation tasks with stronger conditions (e.g., answer-aware question generation), which makes the automatic evaluation more convincing.

#### 2.1.4 Task Popularity

Most tasks in GLGE should use widely-used NLG datasets, which have been implicitly agreed upon by the NLG community as challenging and meaningful. Since GLGE is mainly designed for the generalization capabilities evaluation of the English NLG pretrained model, the choice of task also refers to several related works of NLG pre-training model, such as MASS (Song et al., 2019), BART (Lewis et al., 2020), ProphetNet (Qi et al., 2020), and ERNIE-GEN (Xiao et al., 2020).

Based on the above principles, we invite 10 NLG experts<sup>3</sup> to discuss and vote on existing widely-used NLG datasets. Note that since the GLGE is designed for evaluating the generalization capabilities of NLG in English language, we do not include the cross-lingual NLG tasks, such as machine translation and cross-lingual text summarization (Liang et al., 2020). Finally, we select 6 existing popular NLG datasets. Besides, we also introduce two new datasets selected from real-world scenarios for the GLGE benchmark, which makes GLGE have more practical values. Unlike the existing datasets, the test sets of these new two datasets are hidden, which further ensures the fairness of the evaluation results. The input sequence and output sequence of the selected tasks are all well-defined. We preprocess them and provide the input and output sequence pairs directly, which benefits researchers to focus on model improvements.

<sup>3</sup>Each expert has experience in publishing multiple papers at top conferences in the NLG field.

Corpus	Train	Dev	Test	Src.	Tgt.	Input	Output	Metric
<b>Abstractive Text Summarization</b>								
CNN/DailyMail	287,113	13,368	11,490	822.3	57.9	article	summary	R-1/R-2/R-L
Gigaword	3,803,957	189,651	1,951	33.7	8.7	passage	headline	R-1/R-2/R-L
XSUM	204,017	11,327	11,333	358.5	21.1	article	summary	R-1/R-2/R-L
MSNews	136,082	7,496	7,562	310.7	9.7	article	headline	R-1/R-2/R-L
<b>Answer-aware Question Generation</b>								
SQuAD 1.1	75,722	10,570	11,877	149.4	11.5	answer/passage	question	R-L/B-4/MTR
MSQG	198,058	11,008	11,022	45.9	5.9	highlight/passage	question	R-L/B-4/MTR
<b>Conversational Question Answering</b>								
CoQA	108,647	3,935	4,048	354.4	2.6	history/passage	answer	F1-Score
<b>Personalizing Dialogue</b>								
PersonaChat	122,499	14,602	14,056	120.8	11.8	persona/history	response	B-1/B-2/D-1/D-2

Table 1: GLGE task descriptions and statistics. |Train|: the number of examples in train set. |Src.|: the average number of words in source inputs. R-L: ROUGE-L. B-4: BLUE-4. MTR: METEOR. D-2: Distinct-2.

## 2.2 Tasks and Datasets

GLGE contains eight English NLG tasks, covering text summarization, question generation, generative question answering, and dialogue. Descriptions and statistics of these tasks are shown in Table 1, with concrete examples shown in Appendix.

### 2.2.1 Abstractive Text Summarization

As a typical NLG task, abstractive text summarization aims to generate a short and fluent summary of a long text document. GLGE contains four abstractive text summarization tasks. As discussed in Bhandari et al. (2020), we use ROUGE-1, ROUGE-2, and ROUGE-L (Lin, 2004) as the metrics for these tasks.

**CNN/DailyMail** (Hermann et al., 2015) dataset contains 220K articles from the Daily Mail newspapers and 93K articles from the CNN. Each article contains a bullet point summary. GLGE uses the non-anonymized variant See et al. (2017). After the pre-processing, there are 311,971 ⟨article, summary⟩ pairs, where the source input is the article, and the target output is the summary which consists of multiple sentences.

**Gigaword** (Rush et al., 2015) contains 4M examples extracted from the news articles of the Gigaword corpus (Graff et al., 2003). After the pre-processing, there are 3,995,559 ⟨passage, summary⟩ data pairs, where the source input is the first sentence of the article, and the target output is the headline that usually contains a single sentence.

**XSum** (Narayan et al., 2018) consists of 227K online articles from the British Broadcasting Corporation (BBC), which contains professionally writ-

ten single-sentence summaries. After the pre-processing, there are 226,677 ⟨article, summary⟩ data pairs, where the source input is the news article, and the target output is a single-sentence summary.

**MSNews** **MicroSoft News** headline generation (MSNews) is a new News headline generation dataset we collected for GLGE. We random select 151K online news articles from 2012-01-01 to 2020-09-01 from a real-world news search engine. Each article contains a professionally written single-sentence headline. After the pre-processing, there are 151,140 ⟨article, headline⟩ data pairs, where the source input is the news article, and the target output is a news headline.

### 2.2.2 Answer-aware Question Generation

The question generation task is another typical NLG task, which aims to generate a question based on a given text passage or document. Compared with answer-agnostic question generation tasks that can generate lots of reasonable questions, answer-aware question generation (Zhou et al., 2017) is asked to generate a question asks towards the given answer span based on a given text passage or document. In order to facilitate automatic evaluation, GLGE selects two answer-aware question generation tasks:

**SQuAD 1.1** (Rajpurkar et al., 2016) dataset contains over 100K crowd-worker created questions with the corresponding answer spans in 536 Wikipedia articles. Since the original hidden test set of the SQuAD 1.1 is hidden, we follow (Du et al., 2017; Zhao et al., 2018) to re-split the dataset

with the examples from the original training set and development set. After the pre-processing, there are 98,169 ⟨answer, passage, question⟩ data triples, in which the source input is a Wikipedia passage along with an answer span, and the target output is a question. ROUGE-L, BLEU-4 (Papineni et al., 2002), and METEOR (Banerjee and Lavie, 2005) are used as the metrics.

**MSQG** MicroSoft Question Generation (MSQG) is another dataset we collected, which is a new challenge dataset, the questions in this dataset are freely edited by daily users. For MSQG, we collect 220K passages from a real world search engine. Each passage contains a highlight span and a related query, we regard the queries as questions in this dataset. After the pre-processing, there are 220,088 ⟨highlight span, passage, question⟩ data triples, where the source input is a news passage along with highlight span, and the target output is a user question. ROUGE-L, BLEU-4, and METEOR are used as the metrics.

### 2.2.3 Conversational Question Answering

Conversational question answering is a classic and popular generative question answering task. Compared with the extractive question answering, such as SQuAD (Rajpurkar et al., 2016), conversational question answering requires the model to answer the question based on a running conversation history and the given passage.

**CoQA** (Reddy et al., 2019) dataset contains 127K questions with answers, obtained from 8K conversations about text passages from seven diverse domains. After the pre-processing, there are 116,630 ⟨conversation history, passage, question, answer⟩ data 4-tuples, where the source input is a sequence of conversation history along with a given question and a given passage, and the target output is a free-form answer text. F1-Score (Rajpurkar et al., 2016) is used as the metric.

### 2.2.4 Personalized Dialogue

Conversational AI is an important topic in NLG. Compared with text summarization, the responses of single-turn conversations are diverse and might lack of specification, and thus it is hard to use the single ground-truth for automatic evaluation. We select the personalizing dialogue task, which is a challenging multi-turn conversation task. In addition to the conversation history, this task gives the profile information as an additional condition to facilitate specific response generation.

**PersonaChat** (Zhang et al., 2018) dataset consists of about 160K utterances, which requires the model to generate responses according to given multi-turn conversations and persona profile. After pre-processing, there are 151,157 ⟨persona profile description text, conversation history, response⟩ data triples, where the source input is a sequence of conversation history along with several sentences of persona profile description information, and the target output is a response. BLEU-1, BLEU-2, Distinct-1, and Distinct-2 (Li et al., 2016) are used as the evaluation metrics.

## 2.3 Overall Score

Similar to GLUE (Wang et al., 2019b) and SuperGLUE (Wang et al., 2019a), we seek to give an overall system performance over all GLGE tasks by aggregating the scores of all tasks. We follow GLUE to adopt a simple approach that weighs each task equally. For the tasks with multiple metrics, we firstly average those metrics to get a task score. Besides, because the values of the original Distinct-1 (D-1) and Distinct-2 (D-2) (Li et al., 2016) scores which are used as the metrics for dialogue task are usually quite small (less than 0.01), we re-scale them by 100.0 so that these score values are in the same order of magnitude as other scores.

## 2.4 Challenges of Three Difficulty Levels

As discussed in § 2.1.2, GLGE provides three levels of difficulty for each task, called **GLGE-Easy**, **GLGE-Medium**, and **GLGE-Hard**. The original 8 task datasets as described in § 2.2 constitute the GLGE-Easy. Based on GLGE-Easy, we employ two strategies to further increase the task difficulty. **Low-resource**. We increase the difficulty of GLGE by simulating low-resource scenarios. For each task, we keep the test and development sets of GLGE-Easy and randomly reduce the scale of the training data to 50% of the original train set. The dataset of 8 tasks under this setting is regarded as GLGE-Medium.

**Low-frequency**. In order to further evaluate the generalization capability of the NLG model, we increase the difficulty of GLGE by reducing the word overlap rate between the output of the training set and the output of the test set. The motivation is that a good NLG model should be able to generate a fluent target output based on the input information, even if the target output may contain some low-frequency words. For the test set and development sets of GLGE-Hard, we still use those in GLGE-



Easy. For the training set of GLGE-Hard, we first count the frequency of each token in the target sentence of the test set. Then we remove the stop words of the target sentence of each training sample in GLGE-Easy, ranking them by calculating their word frequency score of the test set. This can be formulated as

$$\text{Score}(y) = \frac{\sum_{w_y \in y} \text{TF}(w_y)}{|y|}, \quad (1)$$

where  $y$  is a target sentence without stop words,  $w_y$  is a token in  $y$ ,  $\text{TF}(w_y)$  denotes the word frequency of the token  $w_y$  in the target sentences of the whole test set, and  $|y|$  denotes the token length of  $y$ . Instead of reducing the training data scale randomly as in GLGE-Medium, we select the top 25% training data with minimum word frequency score of the test set from the original training set as the training set of each dataset. The dataset of 8 tasks under this setting is regarded as GLGE-Hard.

### 3 Experiments

#### 3.1 Baselines

For the baselines, we first evaluate two widely-used non-pretrained models: vanilla LSTM based Seq2Seq (Bahdanau et al., 2015) and vanilla Transformer (Vaswani et al., 2017). Besides, we evaluate several widely used pretrained NLG models, including MASS (Song et al., 2019), BART (Lewis et al., 2020), and ProphetNet (Qi et al., 2020). To further evaluate the performance of the pretrained NLG models of different model sizes and the different scales of the pretraining corpus, we compare the MASS<sub>base</sub>, ProphetNet<sub>base</sub>, MASS<sub>middle</sub>, BART<sub>large</sub>, and ProphetNet<sub>large</sub> on GLGE.

#### 3.2 Implementation Details

**Vanilla LSTM** (Bahdanau et al., 2015). The hyper-parameters and implementation of LSTM-Seq2Seq are based on the *LSTM* register model of Fairseq<sup>4</sup>, where the word embedding dimension, the hidden size, the number of the encoder layer, and the number of the decoder layer are 512, 512, 1, and 1, respectively. For each task in GLGE, we use Adam (Kingma and Ba, 2015) with an initial learning rate of between 0.0001 and 0.0003, and train the LSTM-Seq2Seq for a maximum of 100 epochs.

<sup>4</sup><https://github.com/pytorch/fairseq/blob/master/fairseq/models/lstm.py>.

**Vanilla Transformer** (Vaswani et al., 2017). The hyper-parameters and implementation of Transformer are based on the *transformer\_vaswani\_wmt\_en\_de\_big* register model of Fairseq<sup>5</sup>, which contains a 6-layer encoder and a 6-layer decoder with 1024 embedding/hidden size and 4096 feed-forward filter size. For each task in GLGE, we use Adam with the initial learning rate of between 0.0003 and 0.001, and train the Transformer for a maximum of 20 epochs.

**MASS<sub>base</sub>** (Song et al., 2019). The hyper-parameters and implementation of MASS are based on their source code<sup>6</sup>. MASS<sub>base</sub> contains a 6-layer encoder and a 6-layer decoder with 768 embedding/hidden size and 3072 feed-forward filter size. The MASS<sub>base</sub> is pretrained on BookCorpus (Zhu et al., 2015) and English Wikipedia (16GB in total). For each task in GLGE, we fine-tune MASS<sub>base</sub> with the same hyper-parameters used in their source code<sup>7</sup> for a maximum of 25 epochs.

**MASS<sub>middle</sub>** (Song et al., 2019) which contains a 6-layer encoder and a 6-layer decoder with 1024 embedding/hidden size and 4096 feed-forward filter size. The MASS<sub>middle</sub> is also pretrained on BookCorpus and English Wikipedia (16GB in total). For each task in GLGE, we use the same hyper-parameters as used in MASS<sub>base</sub>.

**ProphetNet<sub>base</sub>** (Qi et al., 2020). The hyper-parameters and implementation of ProphetNet are based on their source code<sup>8</sup>. ProphetNet<sub>base</sub> contains a 6-layer encoder and a 6-layer decoder with 768 embedding/hidden size and 3072 feed-forward filter size. Similar to MASS, the ProphetNet<sub>base</sub> is pretrained on BookCorpus and English Wikipedia (16GB in total) with 125K steps. For each task in GLGE, we fine-tune ProphetNet<sub>base</sub> with the same hyper-parameters used in their source code for a maximum of 10 epochs.

**ProphetNet<sub>large</sub>** (Qi et al., 2020). It contains a 12-layer encoder and 12-layer decoder with 1024 embedding/hidden size and 4096 feed-forward filter size. The ProphetNet<sub>large</sub> is pretrained on the 160GB English language corpora of news, books, stories, and web text for 14 epochs. For each task in

<sup>5</sup><https://github.com/pytorch/fairseq/blob/master/fairseq/models/transformer.py>.

<sup>6</sup><https://github.com/microsoft/MASS>

<sup>7</sup><https://github.com/microsoft/MASS/tree/master/MASS-summarization>

<sup>8</sup><https://github.com/microsoft/ProphetNet>

Models	Avg.	Text Summarization				Question Generation		QA	Dialogue
		CNN/DM	Gigaword	XSUM	MSNews	SQuAD 1.1	MSQG	CoQA	PesonaChat
Metrics		R-1/R-2/R-L				R-L/B-4/MTR		F1	B-1/B-2/D-1/D-2
<b>GLGE-Easy</b>									
LSTM	20.0	37.3/15.7/34.4	34.2/16.0/31.8	25.1/6.9/19.9	30.0/14.6/27.7	27.2/3.8/8.9	25.3/3.5/14.1	15.1	42.2/35.9/0.2/0.7
Transformer	21.9	39.5/16.7/36.7	37.1/18.4/34.5	30.5/10.4/24.2	33.0/15.4/30.0	30.7/4.8/10.9	29.3/5.1/16.6	15.7	38.3/33.6/0.2/0.7
MASS <sub>base</sub>	33.6	42.1/19.5/39.0	38.7/19.7/35.9	39.7/17.2/31.9	39.4/21.0/36.1	49.4/20.1/24.4	38.9/10.2/23.3	65.4	41.0/35.7/1.4/6.9
ProphetNet <sub>base</sub>	33.8	42.5/19.7/39.5	38.9/19.9/36.0	39.8/17.1/32.0	40.6/21.6/37.0	48.0/19.5/23.9	37.1/9.3/22.7	65.3	46.0/38.4/1.3/7.3
MASS <sub>middle</sub>	34.3	42.9/19.8/39.8	38.9/20.2/36.2	39.1/16.5/31.4	40.4/21.5/36.8	49.9/21.3/25.2	38.9/9.5/23.5	67.6	46.0/38.2/1.2/6.2
BART <sub>large</sub>	35.8	44.1/21.2/40.9	38.1/18.4/34.9	45.1/22.2/37.2	43.8/24.0/39.2	50.3/22.0/26.4	38.8/9.2/24.3	68.6	49.9/40.0/1.3/8.0
ProphetNet <sub>large</sub>	36.5	44.2/21.1/41.3	39.5/20.4/36.6	44.4/21.3/36.4	44.1/24.4/40.2	51.5/22.5/26.0	38.3/9.6/23.3	73.0	46.7/39.0/1.3/7.5
<b>GLGE-Medium</b>									
LSTM	18.1	35.3/14.1/32.8	33.3/15.2/31.1	21.5/4.6/17.1	27.0/12.1/24.9	26.6/3.5/8.2	18.6/1.7/9.5	12.9	41.3/35.3/0.1/0.5
Transformer	19.5	35.0/11.0/32.4	36.7/18.1/34.1	27.5/8.3/21.8	26.8/9.7/24.3	28.3/4.1/9.8	27.0/4.2/15.0	14.2	37.7/29.6/0.2/0.7
MASS <sub>base</sub>	33.0	41.2/18.8/38.2	37.9/19.1/35.2	37.4/14.9/29.8	38.9/20.5/35.6	48.9/20.0/24.3	38.2/9.5/22.8	65.0	42.8/36.7/1.3/6.2
ProphetNet <sub>base</sub>	32.6	41.6/19.2/38.7	38.6/19.6/35.7	37.8/15.3/30.4	39.0/20.4/35.7	46.4/17.9/22.5	37.0/8.7/22.3	62.5	45.4/37.7/1.4/7.3
MASS <sub>middle</sub>	33.6	41.5/19.0/38.5	38.3/19.1/35.4	38.4/15.8/30.7	39.6/20.9/36.0	49.3/20.4/24.4	38.3/9.9/22.7	67.2	44.0/37.3/1.3/6.1
BART <sub>large</sub>	35.3	42.8/20.1/39.1	38.0/18.3/34.7	43.1/19.5/34.1	43.4/23.6/38.9	49.7/21.6/25.9	38.4/9.5/24.0	69.4	50.4/39.1/1.2/7.4
ProphetNet <sub>large</sub>	35.5	43.1/20.3/40.1	39.1/19.8/36.1	41.8/18.7/33.8	43.3/23.5/39.4	50.4/21.9/25.8	39.3/10.0/23.7	72.3	42.0/36.4/1.4/7.8
<b>GLGE-Hard</b>									
LSTM	12.6	26.2/6.8/24.2	26.3/9.2/24.6	17.8/2.4/14.3	8.2/0.9/7.6	27.3/1.0/6.7	12.5/0.4/5.0	10.3	36.8/28.7/0.1/0.4
Transformer	14.4	28.3/6.2/25.8	28.6/10.8/26.5	23.0/5.3/18.3	18.0/3.5/16.2	25.9/1.1/7.0	17.0/1.3/8.2	9.9	30.0/29.7/0.1/0.2
MASS <sub>base</sub>	28.2	40.4/18.0/37.3	32.2/13.6/29.5	33.7/11.6/26.7	35.4/17.0/32.4	42.8/13.4/19.0	34.1/7.5/18.6	50.2	40.1/34.9/1.6/7.8
ProphetNet <sub>base</sub>	28.0	40.9/18.4/37.7	32.0/13.5/29.5	34.2/11.6/26.8	35.2/17.0/32.1	41.6/13.4/18.9	32.3/7.2/18.0	48.5	41.6/35.5/1.6/8.3
MASS <sub>middle</sub>	29.1	41.1/18.5/38.0	32.2/13.5/29.9	34.9/12.5/27.6	36.6/18.0/33.4	45.1/16.0/21.3	34.3/8.0/19.0	51.2	41.4/35.4/1.5/7.6
BART <sub>large</sub>	31.0	41.7/19.1/37.9	33.0/13.6/30.0	39.7/16.1/30.9	40.8/20.8/36.4	45.9/18.1/23.7	35.1/8.5/20.7	53.5	48.3/37.3/1.3/7.2
ProphetNet <sub>large</sub>	30.5	41.2/18.7/38.0	32.4/13.7/29.9	39.4/16.1/31.6	40.3/20.5/36.4	46.4/17.0/22.1	34.0/8.2/19.0	54.1	40.5/35.2/1.8/9.2

Table 2: Overall results of baselines on all GLGE tasks. We use the color to highlight the overall score. R-1: ROUGE-1. R-2: ROUGE-2. R-L: ROUGE-L. B-4: BLUE-4. MTR: METEOR. D-1: Distinct-1. D-2: Distinct-2. Note that as discussed in § 2.3, the values of Distinct-1 and Distinct-2 are multiplied by 100.

GLGE, we fine-tune ProphetNet<sub>large</sub> with the same hyper-parameters used in their source code for a maximum of 10 epochs.

**BART<sub>large</sub>** (Lewis et al., 2020). The hyper-parameters and implementation of BART<sub>large</sub> are based on the source code<sup>9</sup>. BART<sub>large</sub> contains a 12-layer encoder and 12-layer decoder with 1024 embedding/hidden size and 4096 feed-forward filter size. The pretraining of BART<sub>large</sub> uses the same pretraining data as Liu et al. (2019), consisting of 160GB of news, books, stories, and web text. For each task in GLGE, we fine-tune BART<sub>large</sub> with the same hyper-parameters used in their source code<sup>10</sup> for a maximum of 20000 iterations.

Except BART, all the baselines adopt BERT-uncased tokenizer. We fine-tune all baselines on each individual task with  $4 \times 16$ GB NVIDIA V100 GPUs. We evaluate the best model checkpoint based on the loss on the development set. During inference, we use beam search (Och and Ney, 2004) with beam size 4 or 5 and remove the duplicated trigrams in beam search (Fan et al., 2018) to obtain

<sup>9</sup><https://github.com/pytorch/fairseq/tree/master/examples/bart>

<sup>10</sup><https://github.com/pytorch/fairseq/blob/master/examples/bart/README.summarization.md>

the generated results.

### 3.3 Results and Analysis

**Overall results.** The main results are presented in Table 2. From the overall scores (highlighted in color), we can observe the fairly consistent gains moving from LSTM to Transformer, and then to pretrained base models and pretrained large models, such as ProphetNet<sub>base</sub> and ProphetNet<sub>large</sub>. The performance gap between the pretrained model and non-pretrained model is obvious. The difference in terms of overall score is about absolute 15% on the three levels GLGE benchmarks (GLGE-Easy, GLGE-Medium, and GLGE-Hard). As expected, the pretrained large models (ProphetNet<sub>large</sub> and BART<sub>large</sub>) achieve the best overall scores.

From the results of each model on the three levels GLGE benchmarks, we can see that each model has a significant drop in performance from GLGE-Easy to GLGE-Medium and GLGE-Hard benchmarks. For both the non-pretrained model and pretrained models, there is a nearly 2% drop in terms of overall score from GLGE-Easy to GLGE-Medium, and about 4%-8% drop from GLGE-Easy to GLGE-Hard. These results illustrate the diversified difficulty of GLGE. We recommend that researchers choose a GLGE benchmark with moder-

Models	Avg.	Text Summarization			Question Generation			QA	Dialogue
		CNN/DM	Gigaword	XSUM	MSNews	SQuAD 1.1	MSQG	CoQA	PesonaChat
Metrics		R-1/R-2/R-L			R-L/B-4/MTR			F1	B-1/B-2/D-1/D-2
<b>GLGE-Medium + Low-frequency Strategy</b>									
LSTM	16.4	35.4/13.9/32.7	31.3/13.2/29.3	20.8/4.0/16.5	25.3/10.4/23.3	24.8/1.2/6.8	15.3/1.1/7.1	10.5	34.8/31.6/0.2/0.6
Transformer	18.1	36.8/13.7/34.0	33.2/14.6/30.9	25.9/7.3/20.7	25.2/7.6/22.7	26.4/1.9/7.8	22.8/2.7/11.9	10.8	39.6/34.3/0.1/0.2
MASS <sub>base</sub>	29.9	41.3/18.8/38.3	34.4/15.3/31.7	36.1/13.6/28.5	36.4/17.9/33.1	45.3/16.0/21.4	34.6/8.5/19.9	54.6	40.3/34.6/1.5/7.5
ProphetNet <sub>base</sub>	29.9	41.7/19.1/38.7	35.1/16.1/32.4	36.5/14.0/28.9	37.4/18.6/34.0	43.7/15.1/20.5	33.8/7.6/20.0	53.1	42.3/35.9/1.4/7.6
MASS <sub>middle</sub>	30.8	41.5/18.9/38.4	35.1/16.0/32.5	37.0/14.4/29.4	37.9/19.1/34.6	46.3/17.0/22.1	35.5/8.9/20.5	56.6	41.0/35.5/1.6/8.0
BART <sub>large</sub>	32.0	42.2/19.6/38.4	35.7/16.1/32.6	41.5/17.9/32.3	41.5/21.4/37.0	46.5/18.1/24.5	35.9/8.4/21.7	54.6	49.1/38.1/1.3/8.4
ProphetNet <sub>large</sub>	32.3	42.7/20.0/39.5	35.3/16.2/32.6	40.6/17.4/32.6	40.9/21.4/37.1	47.8/19.2/23.9	36.4/9.4/20.9	57.6	44.4/36.9/1.4/8.3

Table 3: Overall results of baselines across the tasks of GLGE-Medium + Low-frequency strategy. We use the color to highlight the overall score. R-1: ROUGE-1. R-2: ROUGE-2. R-L: ROUGE-L. B-4: BLUE-4. MTR: METEOR. D-1: Distinct-1. D-2: Distinct-2.

ate difficulty based on the model size of the pre-trained model and the scale size of the pretrained corpus. At the same time, researchers can also perform more comprehensive evaluation of model performance on GLGE benchmarks at all difficulty levels.

**Low-frequency strategy analysis.** To further verify the effectiveness of the low-frequency strategy as described in § 2.4. We build the GLGE-Medium + low-frequency benchmark and evaluate the models on it. Both GLGE-Medium and GLGE-Medium + low-frequency retain 50% of the training samples in the GLGE-Easy training set. The only difference between them is that GLGE-Medium uses random sampling to retain 50% training samples, while GLGE-Medium + low-frequency uses the low-frequency strategy as used in GLGE-Hard to select 50% training samples. We use the same baselines with the same settings in § 3.2 to compare the model performance on GLGE-Medium and GLGE-Medium + low-frequency. The results are shown in Table 3. We can see that after the introduction of the low-frequency strategy, the performance of the models has dropped significantly. These results demonstrate that the low-frequency strategy can effectively improve the difficulty of the benchmark.

**Output diversity analysis.** We further compare the output diversity of each model on all the tasks of GLGE-Easy. We report the mean of the Distinct bigram (Distinct-2) (Li et al., 2016) ratios of the generated samples to the golden references. Note that if the bigram diversity of the generated samples is close to that of the real samples, the distinct bigram ratio is close to 1. The results are shown in Figure 1. In general, output bigram diversity of the pre-trained model is higher than non-pretrained models. For the tasks of CNN/DailyMail (CNN/DM), Gigaword, and MSNews, the bigram diversity of the

generated samples are close to that of the real samples. However, for the tasks of XSUM, SQuAD 1.1, MSQG, CoQA, and PersonaChat, the bigram diversity of the non-pretrained model is significantly lower than that of the pretrained model. For these tasks, the non-pretrained model tends to generate universal responses (Li et al., 2016) or outputs. Moreover, there is still a huge gap between the bigram diversity of pretrained models and real samples (golden) on the task of XSUM, MSQG, and PersonaChat. Obviously, there exists great room for future improvement of the pretrained models in terms of output diversity.

## 4 Related Works

**Benchmarks** Recently, the development of general natural language understanding (NLU) evaluation benchmarks has helped drive the progress of pretraining and transfer learning in NLP. Conneau and Kiela (2018) propose a toolkit, SentEval, for evaluating the quality of universal sentence representations. DecaNLP (McCann et al., 2018) casts ten diversified NLP tasks as a general question-answering format for evaluation. Wang et al. (2019b) propose a widely-used multi-task benchmark, GLUE, for NLU in the English language. There are nine NLU tasks in GLUE, including two single-sentence tasks, three similarity and paraphrase tasks, and four natural language inference tasks. After that, SuperGLUE (Wang et al., 2019a) is proposed as a harder counterpart of GLUE. Besides sentence- and sentence-pair classification tasks used in GLUE, SuperGLUE extends the task formats by introducing coreference resolution and question answering tasks. More recently, A new NLU benchmark called DialoGLUE (Mehri et al., 2020) for task-oriented dialogue is proposed. It consists of seven task-oriented dialogue datasets




Figure 1: Distinct bigram ratios of the generated samples to the golden references on GLGE-Easy.

covering four kinds of NLU tasks: intent prediction, slot tagging, semantic parsing, and dialogue state tracking.

In addition to the English NLU evaluation benchmark, there has been an increasing amount of new benchmarks in other languages. For example, CLUE (Xu et al., 2020) is a Chinese NLU benchmark that consists of eight diverse Chinese NLU tasks, including single-sentence, sentence-pair, and machine reading comprehension tasks. FLUE (Le et al., 2020) is proposed for the French language, which is a French NLU benchmark that includes several NLU tasks, such as text classification, paraphrasing, language inference, parsing, POS tagging, and word sense disambiguation. IndoNLU (Wilie et al., 2020) is a new benchmark for evaluating Indonesian language understanding, which introduces twelve tasks, ranging from single sentence classification to pair-sentences sequence labeling. Furthermore, the multilingual multi-task benchmarks are proposed for cross-lingual evaluating. Hu et al. (2020) introduce XTREME benchmark which is a multi-task benchmark for evaluating the cross-lingual generalization capabilities of multilingual representations across forty languages and nine tasks. Almost at the same time, XGLUE (Liang et al., 2020) is proposed which is a new multilingual multi-task benchmark for cross-lingual pretraining, understanding, and generation. There are eleven cross-lingual tasks including nine NLU tasks and two NLG tasks in XGLUE, each task provides labeled data in multiple languages.

The above benchmarks mostly focus on NLU which provides a range of language understanding

tasks. However, to our best knowledge, there is no benchmark designed specifically for general NLG evaluation. To fill this gap, we introduce GLGE, a new multi-task benchmark for evaluating the generalization capabilities of NLG across eight language generation tasks.

**Pretrained NLG Models** In recent years, pretrained language models (Devlin et al., 2019; Raffel et al., 2020; Yang et al., 2019; Liu et al., 2019; Alberti et al., 2019; Brown et al., 2020; Clark et al., 2020) have achieved state-of-the-art results in several NLU benchmarks. Besides, more and more pretraining based models which are designed for NLG tasks are proposed. Rothe et al. (2020) adopt the Transformer-based sequence-to-sequence (seq2seq) model and leverage the checkpoints of the pretrained NLU models for sequence generation tasks. MASS (Song et al., 2019) pretrains the seq2seq model by dropping a continuous token span to corrupt the text and learns to reconstruct it. Raffel et al. (2020) investigate several model structures and pretraining tasks, and further propose a unified text-to-text transformer called T5. Similarly, BART (Lewis et al., 2020) adopts the encoder-decoder structure and is pretrained with randomly sentence order reconstruction and text in-filling tasks.

More recently, Qi et al. (2020) propose, ProphetNet, which introduces the future n-gram prediction mechanism for language generation. ENRINE-GEN (Xiao et al., 2020) introduces the infilling generation mechanism, noise-aware generation, and span-by-span generation task for NLG model



pretraining. In addition to general NLG tasks, some task-specific pretrained NLG models are proposed. For dialogue and conversation, a dialogue generative pretrained transformer called DialoGPT (Zhang et al., 2020c) is proposed for conversational response generation, which is pretrained on a large-scale conversation-like exchange corpus. Furthermore, PLATO (Bao et al., 2020) is a dialogue generation pretraining framework for chit-chat, knowledge grounded dialogues, and conversational question answering. PLATO introduces discrete latent variables to tackle the one-to-many mapping problem in response generation. For text summarization, Zhang et al. (2020a) propose PE-GASUS, which design the pretraining objectives called gap sentence generation tailored for abstract text summarization.

## 5 Conclusion

To facilitate the development, evaluation, and comparison of new NLG models, we introduce GLGE, a multi-task evaluation benchmark for NLG with three difficulty levels. To the best of our knowledge, GLGE is the first comprehensive NLG evaluation benchmark. We evaluate several baselines on GLGE and analyze their results. The GLGE benchmark is hosted publicly and we invite the research community to submit to the leaderboard.

In future work, we will try to introduce other automatic evaluation metrics, such as BERTscore (Zhang et al., 2020b) and BLEURT (Sellam et al., 2020). Besides, we will compare the correlation between these metrics and human judgment.

## Acknowledgments

This work was supported in part by the National Key R&D Program of China under Grant 2020YFB1406702, in part by NFSC under Grant 61625204, 61836006, and the Science and Technology Major Project of Sichuan province under Grant 2020YFG0478. We thank STC-A NLP, Bing Answers, Bing Ads, Bing Relevance and Microsoft News for providing the datasets.

## References

- Chris Alberti, Kenton Lee, and Michael Collins. 2019. A bert baseline for the natural questions. *arXiv preprint arXiv:1901.08634*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2020. PLATO: Pre-trained dialogue generation model with discrete latent variable. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 85–96, Online. Association for Computational Linguistics.
- Manik Bhandari, Pranav Gour, Atabak Ashfaq, Pengfei Liu, and Graham Neubig. 2020. Re-evaluating evaluation in text summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 9347–9359.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher 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 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. ELECTRA: pre-training text encoders as discriminators rather than generators. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Alexis Conneau and Douwe Kiela. 2018. SentEval: An evaluation toolkit for universal sentence representations. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- 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.
- Kinya Du, Junru Shao, and Claire Cardie. 2017. [Learning to ask: Neural question generation for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1342–1352, Vancouver, Canada. Association for Computational Linguistics.
- 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.
- David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2003. English gigaword. *Linguistic Data Consortium, Philadelphia*, 4(1):34.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1693–1701.
- 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 generalization](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.
- 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*.
- Hang Le, Loïc Vial, Jibril Frej, Vincent Segonne, Maximin Coavoux, Benjamin Lecouteux, Alexandre Alauzen, Benoit Crabbé, Laurent Besacier, and Didier Schwab. 2020. [FlauBERT: Unsupervised language model pre-training for French](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2479–2490, Marseille, France. European Language Resources Association.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [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.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fengei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, 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.
- 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*.
- Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. 2018. [The natural language deathlon: Multitask learning as question answering](#). *arXiv preprint arXiv:1806.08730*.
- Shikib Mehri, Mihail Eric, and Dilek Hakkani-Tur. 2020. [Dialoglue: A natural language understanding benchmark for task-oriented dialogue](#). *arXiv preprint arXiv:2009.13570*.
- 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.
- Franz Josef Och and Hermann Ney. 2004. [The alignment template approach to statistical machine translation](#). *Computational Linguistics*, 30(4):417–449.
- 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.
- 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.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- 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.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. [A neural attention model for abstractive sentence summarization](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 379–389, Lisbon, Portugal. Association for Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur P Parikh. 2020. [BLEURT: learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7881–7892. Association for Computational Linguistics.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. [Mass: Masked sequence to sequence pre-training for language generation](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 5926–5936. PMLR.
- 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 *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019a. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). In *Advances in Neural Information Processing Systems*, pages 3266–3280.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019b. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Bryan Wilie, Karissa Vincentio, Genta Indra Winata, Samuel Cahyawijaya, Xiaohong Li, Zhi Yuan Lim, Sidik Soleman, Rahmad Mahendra, Pascale Fung, Syafri Bahar, and Ayu Purwarianti. 2020. [IndoNLU: Benchmark and resources for evaluating Indonesian natural language understanding](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 843–857, Suzhou, China. Association for Computational Linguistics.
- Dongling Xiao, Han Zhang, Yukun Li, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2020. [Ernie-gen: An enhanced multi-flow pre-training and fine-tuning framework for natural language generation](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 3997–4003.
- Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, Yin Tian, Qianqian Dong, Weitang Liu, Bo Shi, Yiming Cui, Junyi Li, Jun Zeng, Rongzhao Wang, Weijian Xie, Yanting Li, Yina Patterson, Zuoyu Tian, Yiwen Zhang, He Zhou, Shaowei Hua Liu, Zhe Zhao, Qipeng Zhao, Cong Yue, Xinrui Zhang, Zhengliang Yang, Kyle Richardson, and Zhenzhong Lan. 2020. [CLUE: A Chinese language understanding evaluation benchmark](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4762–4772, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5754–5764.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J Liu. 2020a. [Pegasus: Pre-training with](#)

- extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 11328–11339. PMLR.
- 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.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2020b. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020c. [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.
- Yao Zhao, Xiaochuan Ni, Yuanyuan Ding, and Qifa Ke. 2018. [Paragraph-level neural question generation with maxout pointer and gated self-attention networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3901–3910, Brussels, Belgium. Association for Computational Linguistics.
- Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. 2017. [Neural question generation from text: A preliminary study](#). In *National CCF Conference on Natural Language Processing and Chinese Computing*, pages 662–671.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 19–27. IEEE Computer Society.



CNN/DailyMail	<p><b>Article:</b> <i>The only thing crazier than a guy in snowbound Massachusetts boxing up the powdery white stuff and offering it for sale online? People are actually buying it. For \$ 89, self-styled entrepreneur Kyle Waring will ship you 6 pounds of Boston-area snow in an insulated Styrofoam box – enough for 10 to 15 snowballs, he says. [...], a coastal suburb north of Boston. He joked about shipping the stuff to friends and family in warmer states, and an idea was born [...]</i></p> <p><b>Target:</b> <u>A man in suburban Boston is selling snow online to customers in warmer states. For \$ 89, he will ship 6 pounds of snow in an insulated Styrofoam box.</u></p>
Gigaword	<p><b>Passage:</b> <i>U.S. business leaders lashed out Wednesday at legislation that would penalize companies for employing illegal immigrants.</i></p> <p><b>Target:</b> <u>U.S. business attacks tough immigration law.</u></p>
XSUM	<p><b>Article:</b> <i>Burberry reported pre-tax profits of £166m for the year to March. [...], Sales rose 7% to £1.28bn, with the company recording double-digit sales growth in Europe and Asia Pacific. Adjusted profit rose 23% to £215m, taking into account one-off items and a favourable exchange rate. Stores in London in particular benefited from favourable currency movements and increased tourism. [...], Burberry shares were up 7.6% at 659 pence in afternoon trading.</i></p> <p><b>Target:</b> <u>Luxury fashion designer Burberry has returned to profit after opening new stores and spending more on online marketing.</u></p>
MSNews	<p><b>Article:</b> <i>Los Angeles : Actor Chadwick Boseman, who played Black icons Jackie Robinson and James Brown before finding fame as the regal Black Panther in the Marvel cinematic universe, died Friday of cancer, his representative said. He was 43. Boseman died at his home in the Los Angeles area with his wife and family by his side, his publicist Nicki Fioravante told The Associated Press. [...]</i></p> <p><b>Target:</b> <u>Black Panther actor Chadwick Boseman dies of cancer at 43.</u></p>
SQuAD 1.1	<p><b>Passage:</b> <i>Super Bowl 50 was an American football game to determine the champion of the National Football League (NFL) for the 2015 season. The American Football Conference (AFC) champion Denver Broncos defeated the National Football Conference (NFC) champion Carolina Panthers 24–10 to earn their third Super Bowl title. The game was played on February 7, 2016, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California. As this was the 50th Super Bowl, [...].</i></p> <p><b>Answer:</b> <i>Santa Clara, California</i></p> <p><b>Target:</b> <u>Where did Super Bowl 50 take place?</u></p>
MSQG	<p><b>Passage:</b> <i>On March 28, 1830, Congress passed the Indian Removal Act, beginning the forced relocation of thousands of Native Americans in what became known as the Trail of Tears. Not all members of Congress supported the Indian Removal Act.</i></p> <p><b>Highlight Span:</b> <i>Indian Removal Act</i></p> <p><b>Target:</b> <u>What act was passed to relocate the native Americans?</u></p>
CoQA	<p><b>Passage:</b> <i>John was in the third grade, and nine years old. Every day he had to walk home from school. There were some kids in his class who were mean to him, and during the winter they would throw snowballs at him. [...], John thought it was a good deal, and ended up being much better at math.</i></p> <p><b>Conversation History:</b> <i>Q1: Who is in third grade? A1: John. Q2: How old is he? A2: Nine.</i></p> <p><b>Question:</b> <i>What did kids do to him?</i></p> <p><b>Target:</b> <u>Throw snowballs at him.</u></p>
PersonaChat	<p><b>Persona Profile description text:</b> <i>My wife left me and took my children. I don't believe in god. I'm overweight and unhappy. I work at a nursing home. I spend most of my time on Facebook when I'm not working.</i></p> <p><b>Conversation History:</b> <i>Q1: I got a big house with 7 rooms. R1: Nice, majority of my time I am on Facebook one time. Q2: I saw a man fly to the moon.</i></p> <p><b>Target:</b> <u>Like on TV? I'm obese and unhappy.</u></p>

Table 4: Development-set examples from the tasks in GLGE. **Bold** text denotes part of the example format for each task. Text in *italics* is part of the source input text. Underlined text is the target output text. [...] denotes the omitted texts.