

# MVP: MULTI-TASK SUPERVISED PRE-TRAINING FOR NATURAL LANGUAGE GENERATION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Pre-trained language models (PLMs) have achieved remarkable success in natural language generation (NLG) tasks. Up to now, most NLG-oriented PLMs are pre-trained in an unsupervised manner using the large-scale general corpus. In the meanwhile, an increasing number of models pre-trained with labeled data (*i.e.*, “*supervised pre-training*”) showcase superior performance compared to unsupervised pre-trained models. Motivated by the success of supervised pre-training, we propose **M**ulti-task **s**uper**V**ised **P**re-training (**MVP**) for natural language generation. We collect a large-scale natural language generation corpus, MVPCorpus, from 77 datasets over 11 diverse NLG tasks. Then we unify these examples into a general text-to-text format to pre-train the text generation model MVP in a supervised manner. For each task, we further pre-train specific soft prompts to stimulate the model’s capacity to perform a specific task. Extensive experiments have demonstrated the effectiveness and generalizability of our MVP model in a number of NLG tasks, which achieves state-of-the-art performance on 13 out of 17 datasets.

## 1 INTRODUCTION

Natural language generation (NLG, also known as text generation) is a crucial capacity for language intelligence, which aims to generate human-like texts on demand (Garbacea & Mei, 2020). Since the emergence of the pre-training and fine-tuning paradigm, pre-trained language models (PLMs) have dominated the mainstream approaches for NLG tasks (Lewis et al., 2020; Brown et al., 2020). With a large-scale general corpus, the majority of PLMs are pre-trained in an unsupervised (self-supervised) manner by leveraging intrinsic data correlations as supervision signals. However, unsupervised pre-training is likely to incorporate noise that affects the performance of downstream tasks (Feng et al., 2022), also leading to a slower rate of acquiring knowledge (Zhang et al., 2021).

In the meanwhile, more and more large-scale labeled datasets have become easily accessible (Deng et al., 2009; Liu et al., 2020). There is growing evidence that pre-training with labeled data can further improve the performance of PLMs, both in the fields of computer vision (He et al., 2016; Dosovitskiy et al., 2021) and natural language processing (Lin et al., 2020b; Su et al., 2022). These promising developments motivate us to consider pre-training text generation models with labeled data, which is called “*supervised pre-training*” (Feng et al., 2022). Existing work has shown that supervised pre-training can explicitly learn task-specific characteristics and alleviate the discrepancy between unsupervised pre-training and supervised fine-tuning (Sanh et al., 2022; Lin et al., 2020b).

Furthermore, most NLG systems are often trained in a supervised way, requiring supervision signals to learn the input-to-output transformation. For example, dialogue systems learn to generate appropriate responses based on historical utterances, and text summarization systems learn to extract essential information from long documents according to human-written summaries. Therefore, we suspect that supervised pre-training is more suited for NLG-oriented PLMs in essence since it can provide task-related instructions early in the *pre-training stage* instead of a later *fine-tuning stage*.

Inspired by the recent success of supervised pre-training, we propose **M**ulti-task **s**uper**V**ised **P**re-training (**MVP**) for natural language generation by leveraging a variety of labeled text generation datasets. Specially, we collect a large-scale labeled corpus, MVPCorpus, consisting of 77 datasets over 11 text generation tasks. Since recent research shows that an extensive scale of multi-task pre-training (Aribandi et al., 2022) is the key to generalizing to new tasks for large PLMs, we combine these labeled datasets for multi-task pre-training. Existing popular works, as shown in Table 1, mainly

Table 1: Representative PLMs for NLG and NLU tasks using (un)supervised pre-training. We present a more detailed comparison and discussion about supervised pre-training in Section 6.

Settings	Supervised Pre-training	Unsupervised Pre-training
NLG	MVP (ours)	GPT-2, GPT-3, BART, T5, UniLM, MASS, PEGASUS
NLU	FLAN, T0, Muppet, ExT5	BERT, RoBERTa, T5, UniLM, XLNet, ELECTRA

focus on NLU tasks (Sanh et al., 2022; Aribandi et al., 2022) or use unsupervised pre-training (Lewis et al., 2020; Raffel et al., 2020), with no consideration of supervised pre-training on NLG tasks. To fill this gap, we explore supervised pre-training and multi-task learning for deriving both *effective* and *general* NLG models.

To develop our approach, we adopt a Transformer-based (Vaswani et al., 2017) sequence-to-sequence model as the pre-training backbone. In multi-task training, different tasks may “neutralize” the ability learned through other tasks (He & Choi, 2021). To mitigate this potential issue, we propose to learn task-specific prompts based on the MVP model, following the structure of prefix-tuning (Li & Liang, 2021). Task-specific pre-training enables prompts to “store” specialized knowledge for each corresponding task. Integrating MVP with task-specific prompts can further stimulate the model’s capacity to perform some specific tasks.

To summarize, our main contributions center around the following research questions:

- *How to train an NLG-oriented PLM in a supervised pre-training way?* In order to prepare the supervised corpus, we collect a massive labeled MVPCorpus, consisting of 77 datasets over 11 NLG tasks across various domains and specific objectives. To the best of our knowledge, MVPCorpus is the largest collection of NLG datasets. Firstly, we formulate different NLG tasks as a general text-to-text form so that the supervised corpus can be used in a unified way for pre-training an NLG model. Our work presents a simple yet general approach for pre-training a more capable NLG model by leveraging various labeled NLG datasets.
- *Can supervised pre-trained NLG models be both effective and general?* Extensive experiments show that the supervised pre-trained MVP outperforms its unsupervised pre-trained counterpart BART in both full tuning (+7.0% on *avarege*) and parameter-efficient tuning (+4.3% on *avarege*) settings. Our MVP model achieves state-of-the-art performance on 13 out of 17 datasets. Furthermore, the experiments on unseen NLG and NLU tasks demonstrate that our supervised MVP model has a strong generalization ability for unseen tasks.

For reproducing and reusing our work, we release the collection MVPCorpus, the models (*e.g.*, MVP, task-specific prompts, and multi-task variants), and codes for pre-training and fine-tuning at the link: <https://anonymous.4open.science/r/ICLR-2023-Paper3518/>.

## 2 RELATED WORK

**Pre-trained Language Models.** Pre-trained language models have achieved exceptional success in a wide range of tasks, and the majority of them are pre-trained in an unsupervised manner (Brown et al., 2020; Devlin et al., 2019; Lewis et al., 2020; Raffel et al., 2020). For example, with large-scale plain texts as the unsupervised pre-training corpus, GPT-3 (Brown et al., 2020) employ language modeling as the pre-training task, *i.e.*, predicting the next token conditioned on previous tokens; BART (Lewis et al., 2020) learns to recover the original text from corrupted text which has been altered by arbitrary noise transformations. GPT-3 and BART use 570GB and 160GB of unlabeled text as the pre-training corpora, respectively. In the meanwhile, the computer vision community benefits a lot from the labeled dataset ImageNet (Deng et al., 2009). Influential models, such as ResNet (He et al., 2016) and ViT (Dosovitskiy et al., 2021), leverage ImageNet for pre-training. Inspired by the success of pre-training with labeled data, machine translation researchers explore supervised pre-training (McCann et al., 2017; Lin et al., 2020b). Lin et al. (2020b) attempt to pre-train a translation model with parallel data in multiple languages. Despite using much less pre-trained data, mRASP still achieves better performance than translation models pre-trained in an unsupervised manner (Lample & Conneau, 2019; Liu et al., 2020). In this paper, we propose to pre-train a universal NLG model in a supervised manner with collections of labeled datasets (23GB).

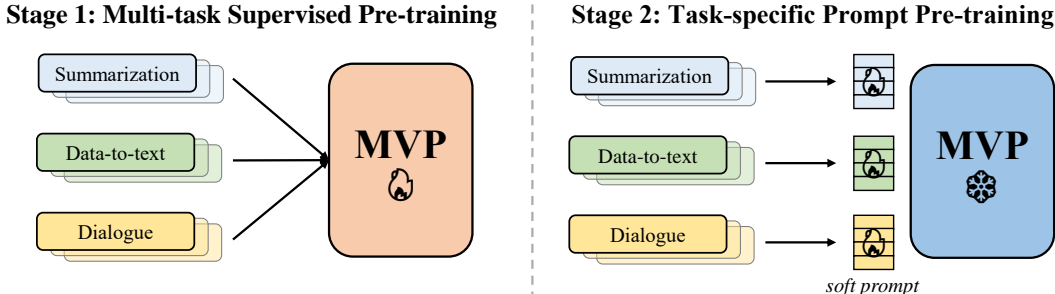


Figure 1: The overview of the pre-training process of our MVP model and task-specific prompts.

**Multi-task Learning.** Our pre-training process is also related to multi-task learning (MTL), a method of mixing multiple tasks into a single training process (Collobert & Weston, 2008). A model trained with MTL can benefit from helpful knowledge of relevant tasks, resulting in improved performance (McCann et al., 2018; Subramanian et al., 2018). Recently, MT-DNN (Liu et al., 2019a) and Muppet (Aghajanyan et al., 2021) collect tens of datasets in the multi-task procedure and achieve better performance in downstream tasks. The *pre-finetuning* schema proposed in Muppet shares a similar idea with our study. Aribandi et al. (2022) further combine the denoising pre-training task of T5 (Raffel et al., 2020) and multi-task learning to pre-train a new model, ExT5. MTL has also contributed to sub-fields of text generation, such as open-ended dialogue system (Zhang et al., 2020), task-oriented dialogue system (Su et al., 2022), text style transfer (Bujnowski et al., 2020), and question answering (Khashabi et al., 2020). At the same time, researchers explore the transferability of models trained on multi-task datasets (Mishra et al., 2022). FLAN (Wei et al., 2022), T0 (Sanh et al., 2022), and ZeroPrompt (Xu et al., 2022) investigate the zero-shot generalization abilities of large PLMs trained on numerous task datasets with well-designed prompts. Ye et al. (2021) develop a benchmark CrossFit to study the few-shot learning ability of models. Compared with these works, we aim to explore multi-task learning to derive both *effective* and *general* NLG models.

**Prompt Learning.** Prompt learning is a thriving method in the field of NLP. Prompt learning converts fine-tuning text into a format similar to pre-training to leverage implicit pre-training knowledge and alleviate the discrepancy between pre-training and fine-tuning (Liu et al., 2021b). GPT-2 (Radford et al., 2019) and T5 (Raffel et al., 2020) add human-written task prompts to the input text. For instance, T5 prepends “*Summarize:*” to the input document for summarization tasks; GPT-3 (Brown et al., 2020) further combines several demonstrations to input to learn task patterns, which is called *in-context learning*. Some researchers also design elaborate prompts or demonstrations for each task and dataset and investigate their effectiveness and robustness (Wei et al., 2022; Sanh et al., 2022; Xu et al., 2022; Mishra et al., 2022). To overcome the constraints of manually constructed prompts, researchers develop continuous (soft) prompts that can be optimized in the continuous space (Lester et al., 2021; Qin & Eisner, 2021). Prefix-tuning (Li & Liang, 2021) increases the number of parameters in prompts and employs prompting in each Transformer layer. Gu et al. (2022) propose PPT to pre-train continuous prompts using unlabeled data. SPoT (Vu et al., 2022) and UnifiedSKG (Xie et al., 2022) learn the prompts on related tasks and transfer them to new tasks.

### 3 THE MVP MODEL

This section introduces our **MVP** model: a **M**ulti-task **s**uper**V**ised **P**re-trained model for natural language generation. We first collect a large-scale NLG corpus, MVPCorpus, from 77 datasets over 11 diverse NLG tasks. After that, we pre-train our MVP model using a mixture of labeled data from MVPCorpus. We further learn the task-specific prompts to stimulate the MVP model to perform a certain task. The overview of our model is illustrated in Figure 1.

### 3.1 DATA COLLECTION

Formally, the natural language generation (NLG) task aims to generate a sequence of tokens  $\mathcal{Y} = (y_1, y_2, \dots, y_n)$  conditioned on input data  $\mathcal{X}$  (*e.g.*, a piece of text or structured data) (Li et al., 2022). Typically, NLG tasks are categorized according to the data format of  $\mathcal{X}$  and  $\mathcal{Y}$ . For example, text summarization condenses a long document into a brief text containing essential information; data-to-text generation produces descriptive text about structured input; and a dialogue system creates pertinent responses given multiple dialog utterances.

In this paper, we collect a large-scale labeled MVPCorpus consisting of 77 labeled datasets from 11 representative NLG tasks<sup>1</sup>, including commonsense generation, data-to-text generation, open-ended dialogue system, paraphrase generation, question answering, question generation, story generation, task-oriented dialogue system, text simplification, text style transfer, and text summarization. These datasets come from various domains and are of different sizes. Some datasets are elaborately hand-crafted and thus relatively small in size, while others are created for large-scale weak supervision. Despite originating from various tasks, these diverse labeled datasets contain rich task-specific supervision signals for establishing global sequence-to-sequence mapping relations. The detailed descriptions of these tasks can be found in Appendix B.1.

Next, we transform all tasks into a unified text-to-text format and convert different input data  $\mathcal{X}$  into a text format. For instance, we linearize structured data (*e.g.*, knowledge graph or table) by concatenating triples or key-value pairs using the special token “[SEP]” for data-to-text generation, and we utilize the special token “[X\_SEP]” to separate answer and paragraph for question generation. The transformed input format for each task can be found in Appendix E.

We divide MVPCorpus into two parts, which are used for pre-training and fine-tuning (evaluation), respectively. For supervised pre-training, we utilize 50 datasets from 7 tasks, including data-to-text generation, open-ended dialogue system, question answering, question generation, story generation, task-oriented dialogue system, and text summarization. To enrich our pre-training corpus, we reverse the input and output of some tasks for obtaining new datasets (*e.g.*, story generation and summarization, question generation and question answering). We also eliminate pre-training examples overlapping with evaluation data to avoid data leakage (more details in Appendix B.2). Finally, we have a 25GB supervised pre-training corpus containing 32M examples. The statistics of datasets for pre-training are listed in Table 7.

For evaluation, we utilize the rest 27 datasets which are more commonly used in the literature. Among these datasets, 23 datasets are from the 7 tasks used in pre-training. We refer to them as *seen* tasks and use them to test the effectiveness of our model. The remaining 4 datasets are from the tasks of commonsense generation, paraphrase generation, simplification, and style transfer, respectively. We call them *unseen* tasks and use them to examine the generalization ability of our model.

### 3.2 MODEL ARCHITECTURE

We pre-train our MVP model and task-specific prompts in two stages. In the first stage, we pre-train the MVP backbone using a mixture of labeled datasets from seven tasks to learn general text-to-text relationships and transferable semantic information across tasks. To indicate each task, we apply human-written prompts to each task instance. For example, we write “*Summarize:*” as the prompt for summarization tasks. The manual prompts for each task are shown in Appendix E.

In the second stage, we freeze our MVP backbone and pre-train a set of task-specific soft prompts (*i.e.*, continuous vectors) to stimulate the model’s capacity to perform some specific tasks. We learn them using a mixture of corresponding intra-task datasets (*i.e.*, datasets under the same task<sup>2</sup>). These soft prompts, which are not shared between tasks, encode the task-specific semantic knowledge to alleviate the blurring-out problem induced by multi-task learning (He & Choi, 2021).

Specifically, we employ the standard Transformer encoder-decoder (Vaswani et al., 2017) as our backbone. Compared to decoder-only architectures such as GPT-3 (Brown et al., 2020) and prefix LMs such as UniLM (Dong et al., 2019), the encoder-decoder architecture is more effective for

<sup>1</sup>We do not consider machine translation tasks but only focusing on English tasks in this work.

<sup>2</sup>For instance, we train summarization-specific prompts using summarization datasets (*e.g.*, Newsroom (Grusky et al., 2018), WikiHow (Koupaee & Wang, 2018), and MSNews (Liu et al., 2021a)).

text generation tasks (Raffel et al., 2020). As for task-specific soft prompts, we insert continuous vectors at each Transformer layer, following prefix-tuning (Li & Liang, 2021). Compared to prompt tuning (Lester et al., 2021), which only adds trainable embeddings to the input layer, the layer-wise prompting of prefix-tuning is more effective and stable (Liu et al., 2022), especially for NLG tasks.

### 3.3 TRAINING DETAILS

Our MVP model adopts a Transformer with 12 layers in both encoder and decoder (406M parameters), the same as the model size of BART<sub>LARGE</sub> (Lewis et al., 2020). The hidden size is 1,024 and the inner hidden size of the feed-forward network is 4,096. We employ the byte-pair-encoding (BPE) tokenizer, and the vocabulary size is 50,267. We initialize the backbone with the BART parameters to provide a good starting point for NLG tasks following previous work (Dong et al., 2019; Zhang et al., 2020). We pre-train the model with a batch size of 8,192 and adopt a temperature-scaled mixing strategy (Raffel et al., 2020) with a rate of  $T = 2$  to mitigate the disparity in tasks and datasets.

We follow prefix-tuning (Li & Liang, 2021) to pre-train task-specific prompts by prepending trainable continuous vectors to the keys and values of the multi-head attention module at each layer. The prompt length is set to 100, and we utilize the MLP reparameterization function with a hidden size of 800 to improve the training robustness and performance (Li & Liang, 2021). Hence, every task prompts have approximately 62M parameters. Then, we freeze the MVP model and train seven groups of task-specific prompts, each of which corresponds to a different task. The batch size is set to 8,192, and we leverage the mixing strategy with a rate of  $T = 2$ .

In the two stages, the maximum length of both input and output sequences is set to 1,024 for supporting examples to contain more tokens. We optimize the model with a constant learning rate of  $3 \times 10^{-5}$  using standard sequence-to-sequence cross-entropy loss. We apply the AdamW optimizer (Loshchilov & Hutter, 2019) with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.98$ ,  $\epsilon = 1 \times 10^{-6}$  to improve training stability (Liu et al., 2019b). The weight decay coefficient is 0.1. For testing, we select the checkpoint with the highest validation performance. All the experiments are conducted on 32 NVIDIA Tesla V100 32GB GPUs. We implement our model using the library `Hugging Face` (Wolf et al., 2020).

In summary, we pre-train a 406M text generation model MVP and seven groups of 62M task-specific prompts. For each downstream task, users can either utilize the MVP backbone (406M) directly or further combine MVP with task-specific prompts (468M).

## 4 EXPERIMENT RESULTS

In this section, we mainly investigate the effectiveness of our proposed supervised pre-training for NLG. Specifically, we fine-tune our MVP model on new datasets for pre-trained (seen) generation tasks under *full tuning* and *parameter-efficient tuning* settings.

For the full tuning setting, we fine-tune the entire model (including the backbone MVP and prompts), while for the parameter-efficient tuning, we only fine-tune prompts but freeze the parameter weights of MVP. We optimize the model via the seq2seq loss with label smoothing (Szegedy et al., 2016) factor of 0.1 and the AdamW optimizer with default hyper-parameters. We sweep over the batch size in  $\{16, 64, 256\}$  and the learning rate in  $\{5 \times 10^{-6}, 1 \times 10^{-5}, 3 \times 10^{-5}\}$  to find the optimal hyper-parameters for each evaluation task. We utilize the checkpoint with the best validation performance for test set inference. During inference, we set the beam size to 5 and the no-repetitive ngram size to 3. For evaluation, we leverage the automatic generation metrics BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), and METEOR (Banerjee & Lavie, 2005) to measure the quality of the generated text and employ Distinct (Li et al., 2016) to evaluate its diversity. Details regarding fine-tuning and evaluation can be found in Appendix C.

We conduct extensive experiments with in different settings. Under full tuning scenarios, we employ the 23 datasets from 7 seen tasks for evaluation. Section 4.1 and Appendix D analyze the performance of our methods on these datasets. To better compare with ExT5 (Aribandi et al., 2022), we conduct experiments on the GEM benchmark (Gehrmann et al., 2021) in Appendix D.2. Under parameter-efficient tuning settings, we utilize the same datasets as in Section 4.1 and the results can be found in Section 4.2. Furthermore, we evaluate our models without fine-tuning and compare them with T0 (Sanh et al., 2022) in Appendix D.3. These extensive results show that our MVP model consistently

Table 2: The main results on seven seen tasks under full tuning settings. The best and second-best results among all the methods are marked in **bold** and underlined, respectively. The SQuAD dataset here is used for the question generation task. The letters B, R, D, and ME denote BLEU, ROUGE, Distinct, and METEOR, respectively. “-” means the work does not compute the corresponding result. These setups and abbreviations are the same below. <sup>a</sup> (Ravaut et al., 2022) <sup>b</sup> (Ke et al., 2021) <sup>c</sup> (Bao et al., 2021) <sup>d</sup> (Xiao et al., 2020) <sup>e</sup> (Lewis et al., 2020) <sup>f</sup> (Liu et al., 2021a) <sup>g</sup> (Guan et al., 2021) <sup>h</sup> (Chen et al., 2022) <sup>i</sup> (He et al., 2022) <sup>j</sup> (Lin et al., 2020c)

Methods	CNN/DailyMail			WebNLG			SQuAD (QG)			CoQA	
	R-1	R-2	R-L	B-4	ME	R-L	B-4	ME	R-L	F1	EM
SOTA	<b>47.16<sup>a</sup></b>	<b>22.55</b>	<b>43.87</b>	66.14 <sup>b</sup>	47.25	76.10	25.97 <sup>c</sup>	<b>27.33</b>	53.43	84.50 <sup>d</sup>	-
<b>MVP</b>	<u>44.44</u>	<u>21.61</u>	<u>40.99</u>	<u>67.76</u>	<u>47.74</u>	<u>77.04</u>	<b>26.21</b>	<u>27.20</u>	<b>53.46</b>	<u>86.51</u>	<u>77.62</u>
BART	44.16 <sup>e</sup>	21.28	40.90	64.55 <sup>b</sup>	46.51	75.13	22.00 <sup>f</sup>	26.40	52.55	68.60 <sup>f</sup>	-
Single	44.36	21.54	40.88	67.74	46.89	76.94	26.09	27.15	53.29	86.20	77.26
<b>MVP+S</b>	44.30	21.53	40.83	<b>68.21</b>	<b>47.77</b>	<b>77.09</b>	<u>26.13</u>	27.18	<u>53.44</u>	<b>86.75</b>	<b>77.97</b>
MVP+R	44.14	21.45	40.72	67.61	47.65	76.70	25.71	27.03	53.09	85.95	77.22
MVP+M	43.97	21.16	40.46	67.45	47.57	76.81	25.46	26.79	52.95	86.28	77.26

Methods	ROCStories				PersonaChat				MultiWOZ		
	B-1	B-2	D-1	D-4	B-1	B-2	D-1	D-2	B-4	Success	Inform
SOTA	33.40 <sup>g</sup>	15.40	-	69.30	49.90 <sup>f</sup>	40.00	<u>1.50<sup>h</sup></u>	9.40	<b>20.50<sup>i</sup></b>	<b>85.30</b>	<b>94.40</b>
<b>MVP</b>	<u>33.42</u>	<u>15.54</u>	<u>2.92</u>	<u>75.06</u>	<b>50.07</b>	<b>40.54</b>	<b>1.54</b>	<b>9.86</b>	<u>20.04</u>	79.20	87.20
BART	30.70 <sup>g</sup>	13.30	-	69.90	49.90 <sup>f</sup>	40.00	1.30	8.00	17.89 <sup>j</sup>	74.91	84.88
Single	32.67	15.29	2.72	72.97	<u>49.96</u>	<u>40.53</u>	1.27	7.63	19.73	75.60	83.70
<b>MVP+S</b>	<b>33.92</b>	<b>15.60</b>	<b>3.44</b>	<b>80.58</b>	47.64	39.84	1.40	8.10	19.91	77.60	84.60
MVP+R	32.93	15.32	2.88	73.83	48.45	40.09	1.30	7.95	19.02	73.30	81.80
MVP+M	33.30	15.51	2.71	74.24	46.26	39.30	1.36	8.07	19.93	72.70	79.70

outperforms various baselines in different scenarios, which demonstrates the effectiveness of our supervised pre-training method for NLG.

#### 4.1 FULL TUNING PERFORMANCE

We design several model variants to verify the effectiveness of our two-stage pre-training method proposed in Section 3.2. For the first-stage model **MVP** that uses multi-task supervised pre-training, we compare it with two competitive backbones using different pre-training strategies:

- **BART<sub>LARGE</sub>** (Lewis et al., 2020): BART is a widely-used PLM for natural language generation using unsupervised pre-training task, *i.e.*, denoising auto encoder.
- **Single-task pre-training (Single)**: We individually train a single model for each task using intra-task datasets under the same pre-training settings in multi-task training. For instance, we pre-train a summarization model using summarization datasets (*e.g.*, Newsroom, WikiHow, and MSNews). Therefore, we have seven single-task pre-trained models in total.

For the second-stage model that integrates pre-trained task-specific prompts (denoted by **MVP+S**), we compare it with two variants using different prompts:

- **Randomly initialized prompts (MVP+R)**: The layer-wise prompts for the MVP model are randomly initialized without pre-training.
- **Multi-Task pre-trained prompts (MVP+M)**: We only pre-train one group of prompts for all tasks, using the same mixed datasets as in the backbone pre-training.

Besides these variants, we further include the best-reported results from original papers in the literature for comparison (denoted as **SOTA**). From the results in Table 2, we can see that:

First, supervised pre-training models (*i.e.*, MVP and Single) achieve better performance than the unsupervised pre-trained model BART, yielding an average improvement of 7.0% and 4.4% (in ratio), respectively. This finding demonstrates the effectiveness of our supervised pre-training method. With

Table 3: The main results on seven seen tasks under parameter-efficient settings. We also include the results of BART and MVP under the full tuning setting (denoted as FT) for comparison.

Methods	CNN/DailyMail			WebNLG			SQuAD (QG)			CoQA	
	R-1	R-2	R-L	B-4	ME	R-L	B-4	ME	R-L	F1	EM
FT BART	44.16	21.28	40.90	64.55	46.51	75.13	22.00	26.40	52.55	68.60	–
FT MVP	44.44	21.61	40.99	67.76	47.74	77.04	26.21	27.20	53.46	86.51	77.62
<b>MVP+S</b>	<b>43.03</b>	<b>20.27</b>	<b>39.72</b>	<b>66.73</b>	<b>47.42</b>	<b>76.36</b>	<b>25.28</b>	<b>26.66</b>	<b>52.69</b>	<b>86.44</b>	<b>76.84</b>
BART+R	42.47	19.82	39.15	65.54	46.86	75.24	24.27	26.07	52.03	82.22	71.92
MVP+R	42.84	20.21	39.61	66.12	47.12	75.83	25.05	26.34	52.57	85.51	75.56
MVP+M	42.99	<b>20.36</b>	<b>39.70</b>	<b>66.40</b>	<b>47.16</b>	<b>75.89</b>	<b>25.24</b>	<b>26.49</b>	<b>52.88</b>	<b>85.90</b>	<b>76.34</b>
Methods	ROCStories				PersonaChat				MultiWOZ		
	B-1	B-2	D-1	D-4	B-1	B-2	D-1	D-2	B-4	Success	Inform
FT BART	30.70	13.30	–	69.90	49.90	40.00	1.30	8.00	17.89	74.91	84.88
FT MVP	33.42	15.54	2.92	75.06	50.07	40.54	1.54	9.86	20.04	79.20	87.20
<b>MVP+S</b>	<b>32.94</b>	<b>15.12</b>	<b>2.98</b>	<b>71.09</b>	<b>47.11</b>	<b>39.51</b>	<b>1.39</b>	<b>7.28</b>	<b>19.24</b>	<b>71.40</b>	<b>77.80</b>
BART+R	32.14	14.71	2.85	68.94	46.23	38.98	1.30	6.82	17.94	62.20	69.20
MVP+R	32.28	14.85	<b>2.97</b>	<b>70.29</b>	46.70	39.23	1.31	6.98	18.86	64.40	71.40
MVP+M	<b>32.62</b>	<b>15.28</b>	2.95	69.58	46.78	<b>39.40</b>	<b>1.33</b>	<b>7.13</b>	<b>19.13</b>	<b>67.20</b>	<b>72.90</b>

labeled datasets, supervised pre-training enables the model to acquire more task-specific information, thus leading to improved results on downstream tasks. Regarding multi-task pre-training (MVP) and single-task (Single), our MVP model outperforms its single-task counterparts by 2.7%. This result indicates that the proposed multi-task learning approach can enhance single-task performance by learning transferable semantic information across tasks.

Second, task-specific prompt learning is effective to alleviate the “blurring-out” issue of multi-task learning. For tasks such as data-to-text generation and question answering, MVP with single-task prompt pre-training (MVP+S) consistently outperforms the other two variants (MVP+R, MVP+M). This verifies that task-specific prompts can acquire specialized knowledge of each task and stimulate the capacity of the MVP model to perform certain tasks.

Finally, our supervised pre-training approach achieves five new SOTA results on data-to-text generation, question generation, question answering, story generation, and open-ended dialogue tasks in Table 2. We also achieve SOTA performance in six out of eight datasets in Table 9, which shows the strong text generation capability of our MVP model. As for the remaining tasks, the SOTA models incorporate specific techniques tailored to the tasks, *e.g.*, the re-ranking framework (Ravaut et al., 2022) and various task-specific objectives (He et al., 2022), which yield better performance than our models. In contrast, the results of our models are very competitive, which is developed based on a general architecture and a unified learning objective.

#### 4.2 PARAMETER-EFFICIENT TUNING PERFORMANCE

In the lightweight fine-tuning setting, we only tune the prompts while freezing the backbone MVP model. Besides our MVP+S model, we consider comparing the following methods:

- **Prefix-tuning** (Li & Liang, 2021): Prefix-tuning is a popular prompt-based lightweight tuning method for text generation. We employ BART<sub>LARGE</sub> as its backbone, denoted as **BART+R**.
- **Only tuning randomly initialized prompts (MVP+R)**: This variant only tunes the randomly initialized prompts of MVP+R, and it shares a similar idea with prefix-tuning.
- **Only tuning multi-task pre-trained prompts (MVP+M)**: This variant only tunes the multi-task pre-trained prompts of MVP+M. Such an idea has been used in SPoT (Vu et al., 2022).

From the experimental results in Table 3, we can see that: the good performance of the MVP model in lightweight settings further demonstrates the effectiveness of supervised pre-training. By comparing two randomly initialized prompting methods (BART+R and MVP+R), we can see that MVP+R achieves superior performance to BART+R (+2.0%) due to its multi-task supervised backbone. Furthermore, when initialized with pre-trained prompts, MVP+S and MVP+M achieve improved

Table 4: The main results of unseen NLG tasks. We use AESOP and SC & BLEU to denote the methods proposed by Sun et al. (2021) and Lai et al. (2021), respectively. Accuracy is calculated by a pre-trained TextCNN to evaluate the style strength, and HM denotes the harmonic mean of BLEU-4 and style accuracy. <sup>a</sup> (Sun et al., 2021) <sup>b</sup> (Lai et al., 2021)

AESOP	Quora					SC & BLEU	GYAFC E&M			GYAFC F&R		
	B-4	R-1	R-2	R-L	ME		B-4	Accuracy	HM	B-4	Accuracy	HM
+BART	47.30 <sup>a</sup>	73.30	54.10	75.10	49.70	+BART	76.50 <sup>b</sup>	93.70	83.90	79.30	92.00	85.20
+MVP	<b>49.86</b>	<b>74.93</b>	<b>56.55</b>	<b>76.56</b>	<b>52.27</b>	+MVP	<b>77.01</b>	<b>94.66</b>	<b>84.92</b>	<b>79.70</b>	<b>93.07</b>	<b>85.87</b>

Table 5: The main results of NLU tasks on the GLUE benchmark. We evaluate the results on the official website <https://gluebenchmark.com/>. Matt. means the Matthews correlation coefficient. Acc. stands for the accuracy rate. P/S Corr. denote Pearson and Spearman correlation coefficients. m./mm. refer to the accuracy of the matched and mismatched domains. Avg. is a macro-average of scores defined in Wang et al. (2019).

Methods	CoLA Matt.	SST-2 Acc.	MRPC F1/Acc.	STS-B P/S Corr.	QQP F1/Acc.	MNLI m./mm.	QNLI Acc.	RTE Acc.	Average
BART	<b>60.30</b>	96.30	90.47 / 86.70	90.97 / 90.30	73.03 / 89.87	<b>90.03 / 89.27</b>	94.60	79.83	85.17
MVP	59.87	<b>96.43</b>	<b>92.07 / 89.43</b>	<b>91.37 / 90.90</b>	<b>73.20 / 90.13</b>	89.70 / 88.73	<b>95.10</b>	<b>82.87</b>	<b>85.88</b>

results over MVP+R, which is consistent with the findings of SPoT (Vu et al., 2022). When compared with MVP+M, MVP+S performs marginally better by 1.2%, indicating that task-specific prompts are useful to improve the model in specific generation tasks.

Surprisingly, our lightweight MVP+S can even outperform fully tuned BART on tasks such as question generation and question answering, showcasing the effectiveness of the proposed supervised pre-training approach. Another note is that lightweight prompting methods (Lester et al., 2021; Vu et al., 2022) that work on NLU tasks cannot achieve competitive performances when compared to full tuning methods on NLG tasks.

## 5 GENERALIZATION ABILITY

In this section, we test our MVP model on unseen NLG and NLU tasks to verify the generalizability.

**Generalization to Unseen NLG Tasks.** According to Deng et al. (2021), an NLG task can be assigned to one of the following three categories: compression (*e.g.*, summarization), transduction (*e.g.*, translation), or creation (*e.g.*, story generation). Since we do not include any transduction tasks during pre-training, we evaluate our MVP model using two unseen transduction NLG tasks: paraphrase generation and text style transfer. We select the SOTA methods for these two tasks, *i.e.*, AESOP (Sun et al., 2021) for paraphrase generation and SC & BLEU (Lai et al., 2021) for text style transfer, and replace their backbone BART with our MVP model for comparison. The experimental setup remains the same as described in Section 4, and details are reported in Appendix C. From the results in Table 4, we can see that our model outperforms BART by a ratio of 2.2% and achieves two new SOTA results, which verifies the strong generalizability of our model. This finding shows that our MVP model is more capable than BART and can serve as a general yet effective backbone to solve more specific tasks by providing superior parameter initialization.

**Generalization to Unseen NLU Tasks.** Although MVP is designed especially for NLG tasks, we also evaluate its performance on unseen NLU tasks using the widely-used GLUE benchmark (Wang et al., 2019). We compare our model to BART<sub>LARGE</sub> using its original sequence classification method (Lewis et al., 2020). The detailed settings can be found in Appendix C. According to the results presented in Table 5, our MVP model outperforms BART on 9 of 12 metrics and has a superior overall performance of 0.71%. This result indicates the strong generalization ability of our MVP model and further demonstrates that our supervised pre-training not only learns generation ability but also improves the overall semantic representations.



Table 6: Comparison of our work with existing *supervised pre-training* methods. #NLG/#NLU denote the number of NLG and NLU tasks, respectively. PT denotes pre-training, FT denotes fine-tuning and SP denotes supervised pre-training.

Methods	#NLG (PT)	#NLU (PT)	#NLG (FT)	#NLU (FT)	SP model	SP prompts	Open source
FLAN	3	9	2	9	✓	✗	✗
T0	2	6	0	4	✓	✗	✓
Muppet	1	3	1	3	✓	✗	✓
ExT5	3	8	6	8	✓	✗	✗
SPoT	1	4	0	6	✗	✓	✗
MVP (ours)	<b>7</b>	0	<b>11</b>	3	✓	✓	✓

## 6 DISCUSSION

**Differences with Existing Methods.** To the best of our knowledge, existing supervised pre-training works mainly focus on NLU tasks (Aghajanyan et al., 2021; Aribandi et al., 2022) or a small number of NLG tasks (Lin et al., 2020b; Su et al., 2022). Given the superior performance achieved by supervised pre-training approaches, it is important to explore supervised pre-training for deriving both *effective* and *general* NLG models. Our work makes a significant contribution in this direction, achieving SOTA performance with a single model on 13 of 17 datasets. Compared with its strong counterpart ExT5 (Aribandi et al., 2022), our MVP model outperforms it in 26 out of 27 metrics (detailed in Appendix D.2). In order to better understand the difference between our paper with previous supervised (multi-task) pre-training studies, we present a detailed comparison in Table 6. As we can see, our work conducts the study with the largest number of NLG tasks for both supervised pre-training and fine-tuning, incorporates task-specific prompts, and also releases all the important resources for reproducing or reusing our work.

**Applicability.** To facilitate the application of our work, we have released the collection corpus, pre-trained models, task-specific prompts, and the generated texts. Our collected MVPCorpus is the largest NLG task collection. We can use all the data to pre-train a general model or select a subset to continue pre-training a domain- or task- specific model (Gururangan et al., 2020). Our MVPCorpus can also be considered as the evaluation benchmark for different NLG tasks. Furthermore, our MVP model can be used to achieve new state-of-the-art results in various NLG tasks. Users can either fine-tune the MVP model or integrate it with task-specific prompts to achieve better results based on sufficient labeled data. Even in data-scarce domains, our MVP model can be also directly employed to obtain good performance without fine-tuning. In addition, our MVP model can provide effective parameter initialization for improving existing methods, as described in Section 5. Finally, the pre-trained task-specific prompts and the generated texts can be further used to study the task similarity and their effect on the multi-task pre-training.

## 7 CONCLUSION

In this paper, we present **M**ulti-task **s**uper**V**ised **P**re-training (**MVP**) for natural language generation. Firstly, we collect a large-scale NLG corpus, MVPCorpus, from 77 datasets over 11 diverse NLG tasks. After converting various NLG tasks into a unified text-to-text format, we propose multi-task supervised pre-training to learn an *effective* and *general* model **MVP** with task-specific prompts for NLG tasks. Extensive experiments have demonstrated that: (1) supervised pre-training is beneficial for NLG tasks as a general solution. Our MVP model outperforms the unsupervised pre-trained model BART and even achieves SOTA performance on 13 out of 17 datasets; (2) supervised pre-trained models have strong generalization ability on unseen generation or even understanding tasks.

In future work, we will explore the multilingual version of our MVP model by covering more datasets in other languages. Such a model is expected to capture language-independent task characteristics and improve the generation tasks in the minority language. Besides, it is interesting to study how different tasks relate to each other in the unified semantic space of the MVP model, which can inspire methods that incorporate task relations as prior.

## BROADER IMPACTS

In this paper, we pre-trained a language model MVP using labeled NLG datasets. According to the research (Bender et al., 2021; Bommasani et al., 2021), PLMs tend to “remember” what they have “seen” in pre-training corpus. This could result in the reproduction of undesirable biases from pre-training data on downstream tasks. Training data intervention could be a solution to alleviate this issue (Lu et al., 2020). It is also interesting to investigate whether supervised pre-training produces fewer biases than unsupervised pre-training in the future.

Environmental impact is another factor we should consider. We have attempted a more efficient pre-training strategy and released our PLM for future work. In contrast to large PLMs with tens of billions of parameters, such as T5 (Raffel et al., 2020) and GPT-3 (Brown et al., 2020), we pre-train only a small model with hundreds of millions of parameters. In addition, we utilize supervised pre-training data and initialize our model with pre-trained BART, both of which improve the convergence of our model. Ultimately, our model is pre-trained for about 20,000 steps, whereas BART of the same size is pre-trained for 500,000 steps.

## REPRODUCIBILITY

For reproducing and reusing our work, we have released the collection MVPCorpus, the models (e.g., MVP, task-specific prompts and multi-task variants), intermediate results (e.g., the generated texts), and source codes for pre-training and fine-tuning at the link: <https://anonymous.4open.science/r/ICLR-2023-Paper3518/>. The detailed settings of experiments are listed in Appendix C. We hope that these open-source resources will facilitate future work on supervised pre-training and contribute to the advancement of NLG research.

## REFERENCES

URL <https://github.com/markriedl/WikiPlots>.

Oshin Agarwal, Heming Ge, Siamak Shakeri, and Rami Al-Rfou. Knowledge graph based synthetic corpus generation for knowledge-enhanced language model pre-training. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 3554–3565, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.278. URL <https://aclanthology.org/2021.naacl-main.278>.

Armen Aghajanyan, Anchit Gupta, Akshat Shrivastava, Xilun Chen, Luke Zettlemoyer, and Sonal Gupta. Muppet: Massive multi-task representations with pre-finetuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 5799–5811, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.468. URL <https://aclanthology.org/2021.emnlp-main.468>.

Huda Alamri, Vincent Cartillier, Raphael Gontijo Lopes, Abhishek Das, Jue Wang, Irfan Essa, Dhruv Batra, Devi Parikh, Anoop Cherian, Tim K Marks, et al. Audio visual scene-aware dialog (avsd) challenge at dstc7. *arXiv preprint arXiv:1806.00525*, 2018. URL <http://arxiv.org/abs/1806.00525>.

Fernando Alva-Manchego, Louis Martin, Antoine Bordes, Carolina Scarton, Benoît Sagot, and Lucia Specia. ASSET: A dataset for tuning and evaluation of sentence simplification models with multiple rewriting transformations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4668–4679, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.424. URL <https://aclanthology.org/2020.acl-main.424>.

Vamsi Aribandi, Yi Tay, Tal Schuster, Jinfeng Rao, Huaixiu Steven Zheng, Sanket Vaibhav Mehta, Honglei Zhuang, Vinh Q. Tran, Dara Bahri, Jianmo Ni, Jai Gupta, Kai Hui, Sebastian Ruder, and Donald Metzler. Ext5: Towards extreme multi-task scaling for transfer learning. In *International*

- Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=Vzh1BFUCiIX>.
- Satanjeev Banerjee and Alon Lavie. 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*, pp. 65–72, Ann Arbor, Michigan, June 2005. Association for Computational Linguistics. URL <https://aclanthology.org/W05-0909>.
- Hangbo Bao, Li Dong, Wenhui Wang, Nan Yang, and Furu Wei. s2s-ft: Fine-tuning pretrained transformer encoders for sequence-to-sequence learning. *arXiv preprint arXiv:2110.13640*, 2021. URL <https://arxiv.org/abs/2110.13640>.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, pp. 610–623, New York, NY, USA, 2021. Association for Computing Machinery. doi: 10.1145/3442188.3445922. URL <https://doi.org/10.1145/3442188.3445922>.
- Luisa Bentivogli, Ido Dagan, Hoa Trang Dang, Danilo Giampiccolo, and Bernardo Magnini. The fifth pascal recognizing textual entailment challenge. In *In Proc Text Analysis Conference (TAC'09)*, 2009.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. Semantic parsing on Freebase from question-answer pairs. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1533–1544, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL <https://aclanthology.org/D13-1160>.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. URL <http://arxiv.org/abs/2108.07258>.
- 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. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 5016–5026, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1547. URL <https://aclanthology.org/D18-1547>.
- Paweł Bujnowski, Kseniia Ryzhova, Hyungtak Choi, Katarzyna Witkowska, Jarosław Piersa, Tymoteusz Krumholz, and Katarzyna Beksa. An empirical study on multi-task learning for text style transfer and paraphrase generation. In *Proceedings of the 28th International Conference on Computational Linguistics: Industry Track*, pp. 50–63, Online, December 2020. International Committee on Computational Linguistics. doi: 10.18653/v1/2020.coling-industry.6. URL <https://aclanthology.org/2020.coling-industry.6>.
- Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Ben Goodrich, Daniel Duckworth, Semih Yavuz, Amit Dubey, Kyu-Young Kim, and Andy Cedilnik. Taskmaster-1: Toward a realistic and diverse dialog dataset. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 4516–4525, Hong Kong, China, November

2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1459. URL <https://aclanthology.org/D19-1459>.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pp. 1–14, Vancouver, Canada, August 2017. Association for Computational Linguistics. doi: 10.18653/v1/S17-2001. URL <https://aclanthology.org/S17-2001>.
- Mingda Chen, Sam Wiseman, and Kevin Gimpel. WikiTableT: A large-scale data-to-text dataset for generating Wikipedia article sections. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 193–209, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.17. URL <https://aclanthology.org/2021.findings-acl.17>.
- Wei Chen, Yeyun Gong, Song Wang, Bolun Yao, Weizhen Qi, Zhongyu Wei, Xiaowu Hu, Bartuer Zhou, Yi Mao, Weizhu Chen, Biao Cheng, and Nan Duan. DialogVED: A pre-trained latent variable encoder-decoder model for dialog response generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 4852–4864, Dublin, Ireland, May 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.acl-long.333>.
- Wenhu Chen, Jianshu Chen, Yu Su, Zhiyu Chen, and William Yang Wang. Logical natural language generation from open-domain tables. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7929–7942, Online, July 2020a. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.708. URL <https://aclanthology.org/2020.acl-main.708>.
- Wenhu Chen, Yu Su, Xifeng Yan, and William Yang Wang. KGPT: Knowledge-grounded pre-training for data-to-text generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 8635–8648, Online, November 2020b. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.697. URL <https://aclanthology.org/2020.emnlp-main.697>.
- Liyang Cheng, Dekun Wu, Lidong Bing, Yan Zhang, Zhanming Jie, Wei Lu, and Luo Si. ENT-DESC: Entity description generation by exploring knowledge graph. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1187–1197, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.90. URL <https://aclanthology.org/2020.emnlp-main.90>.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen-tau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. QuAC: Question answering in context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2174–2184, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1241. URL <https://aclanthology.org/D18-1241>.
- Ronan Collobert and Jason Weston. A unified architecture for natural language processing: deep neural networks with multitask learning. In *Machine Learning, Proceedings of the Twenty-Fifth International Conference (ICML 2008), Helsinki, Finland, June 5-9, 2008*, volume 307 of *ACM International Conference Proceeding Series*, pp. 160–167. ACM, 2008. doi: 10.1145/1390156.1390177. URL <https://doi.org/10.1145/1390156.1390177>.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. The pascal recognising textual entailment challenge. In *Machine Learning Challenges. Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Tectual Entailment*, pp. 177–190, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops)*, pp. 248–255, Los Alamitos, CA, USA, jun 2009. IEEE Computer Society. doi: 10.1109/CVPR.2009.5206848. URL <https://doi.ieeeecomputersociety.org/10.1109/CVPR.2009.5206848>.

- Mingkai Deng, Bowen Tan, Zhengzhong Liu, Eric Xing, and Zhiting Hu. Compression, transduction, and creation: A unified framework for evaluating natural language generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 7580–7605, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.599. URL <https://aclanthology.org/2021.emnlp-main.599>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. Wizard of wikipedia: Knowledge-powered conversational agents. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=r1173iRqKm>.
- Jesse Dodge, Andreea Gane, Xiang Zhang, Antoine Bordes, Sumit Chopra, Alexander H. Miller, Arthur Szlam, and Jason Weston. Evaluating prerequisite qualities for learning end-to-end dialog systems. In *4th International Conference on Learning Representations, ICLR 2016*, 2016. URL <http://arxiv.org/abs/1511.06931>.
- William B. Dolan and Chris Brockett. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*, 2005. URL <https://aclanthology.org/I05-5002>.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. Unified language model pre-training for natural language understanding and generation. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/c20bb2d9a50d5ac1f713f8b34d9aac5a-Paper.pdf>.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- Layla El Asri, Hannes Schulz, Shikhar Sharma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul Mehrotra, and Kaheer Suleman. Frames: a corpus for adding memory to goal-oriented dialogue systems. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pp. 207–219, Saarbrücken, Germany, August 2017. Association for Computational Linguistics. doi: 10.18653/v1/W17-5526. URL <https://aclanthology.org/W17-5526>.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. Key-value retrieval networks for task-oriented dialogue. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pp. 37–49, Saarbrücken, Germany, August 2017. Association for Computational Linguistics. doi: 10.18653/v1/W17-5506. URL <https://aclanthology.org/W17-5506>.
- Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 889–898, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1082. URL <https://aclanthology.org/P18-1082>.
- Yutong Feng, Jianwen Jiang, Mingqian Tang, Rong Jin, and Yue Gao. Rethinking supervised pre-training for better downstream transferring. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=Jjcv9MTqhcq>.
- Cristina Garbacea and Qiaozhu Mei. Neural language generation: Formulation, methods, and evaluation. *arXiv preprint arXiv:2007.15780*, 2020. URL <http://arxiv.org/abs/2007.15780>.

- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. Creating training corpora for NLG micro-planners. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 179–188, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1017. URL <https://aclanthology.org/P17-1017>.
- 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 Cabezudo, Hendrik Strobelt, Nishant Subramani, Wei Xu, Diyi Yang, Akhila Yerukola, and Jiawei Zhou. 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)*, pp. 96–120, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.gem-1.10. URL <https://aclanthology.org/2021.gem-1.10>.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. The third PASCAL recognizing textual entailment challenge. In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pp. 1–9, Prague, June 2007. Association for Computational Linguistics. URL <https://aclanthology.org/W07-1401>.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pp. 70–79, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-5409. URL <https://aclanthology.org/D19-5409>.
- Karthik Gopalakrishnan, Behnam Hedayatnia, Qinglang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, and Dilek Hakkani-Tür. Topical-chat: Towards knowledge-grounded open-domain conversations. In *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association*, pp. 1891–1895. ISCA, 2019. doi: 10.21437/Interspeech.2019-3079. URL <https://doi.org/10.21437/Interspeech.2019-3079>.
- David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. English gigaword. *Linguistic Data Consortium, Philadelphia*, 4(1):34, 2003.
- Max Grusky, Mor Naaman, and Yoav Artzi. Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 708–719, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1065. URL <https://aclanthology.org/N18-1065>.
- Jing Gu, Mostafa Mirshekari, Zhou Yu, and Aaron Sisto. ChainCQG: Flow-aware conversational question generation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 2061–2070, Online, April 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.eacl-main.177. URL <https://aclanthology.org/2021.eacl-main.177>.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. PPT: Pre-trained prompt tuning for few-shot learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8410–8423, Dublin, Ireland, May 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.acl-long.576>.
- Jian Guan, Xiaoxi Mao, Changjie Fan, Zitao Liu, Wenbiao Ding, and Minlie Huang. Long text generation by modeling sentence-level and discourse-level coherence. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint*

- Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 6379–6393, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.499. URL <https://aclanthology.org/2021.acl-long.499>.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. Don’t stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 8342–8360, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.740. URL <https://aclanthology.org/2020.acl-main.740>.
- R Bar Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, Bernardo Magnini, and Idan Szpektor. The second pascal recognising textual entailment challenge. In *Proceedings of the Second PASCAL Challenges Workshop on Recognising Textual Entailment*, volume 7, 2006.
- Han He and Jinho D. Choi. The stem cell hypothesis: Dilemma behind multi-task learning with transformer encoders. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 5555–5577, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.451. URL <https://aclanthology.org/2021.emnlp-main.451>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, Los Alamitos, CA, USA, jun 2016. IEEE Computer Society. doi: 10.1109/CVPR.2016.90. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2016.90>.
- Wanwei He, Yinpei Dai, Yinhe Zheng, Yuchuan Wu, Zheng Cao, Dermot Liu, Peng Jiang, Min Yang, Fei Huang, Luo Si, Jian Sun, and Yongbin Li. Galaxy: A generative pre-trained model for task-oriented dialog with semi-supervised learning and explicit policy injection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):10749–10757, 2022. doi: 10.1609/aaai.v36i10.21320. URL <https://ojs.aaai.org/index.php/AAAI/article/view/21320>.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL <https://proceedings.neurips.cc/paper/2015/file/afdec7005cc9f14302cd0474fd0f3c96-Paper.pdf>.
- Xinyu Hua and Lu Wang. PAIR: Planning and iterative refinement in pre-trained transformers for long text generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 781–793, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.57. URL <https://aclanthology.org/2020.emnlp-main.57>.
- Chao Jiang, Mounica Maddela, Wuwei Lan, Yang Zhong, and Wei Xu. Neural CRF model for sentence alignment in text simplification. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7943–7960, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.709. URL <https://aclanthology.org/2020.acl-main.709>.
- Zhijing Jin, Qipeng Guo, Xipeng Qiu, and Zheng Zhang. GenWiki: A dataset of 1.3 million content-sharing text and graphs for unsupervised graph-to-text generation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 2398–2409, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics. doi: 10.18653/v1/2020.coling-main.217. URL <https://aclanthology.org/2020.coling-main.217>.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1601–1611, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1147. URL <https://aclanthology.org/P17-1147>.

- Pei Ke, Haozhe Ji, Yu Ran, Xin Cui, Liwei Wang, Linfeng Song, Xiaoyan Zhu, and Minlie Huang. JointGT: Graph-text joint representation learning for text generation from knowledge graphs. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 2526–2538, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.223. URL <https://aclanthology.org/2021.findings-acl.223>.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. UNIFIEDQA: Crossing format boundaries with a single QA system. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 1896–1907, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.171. URL <https://aclanthology.org/2020.findings-emnlp.171>.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. The NarrativeQA reading comprehension challenge. *Transactions of the Association for Computational Linguistics*, 6:317–328, 2018. doi: 10.1162/tacl.a.00023. URL <https://aclanthology.org/Q18-1023>.
- Rik Koncel-Kedziorski, Dhanush Bekal, Yi Luan, Mirella Lapata, and Hannaneh Hajishirzi. Text Generation from Knowledge Graphs with Graph Transformers. 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)*, pp. 2284–2293, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1238. URL <https://aclanthology.org/N19-1238>.
- Mahnaz Koupaee and William Yang Wang. Wikihow: A large scale text summarization dataset. *arXiv preprint arXiv:1810.09305*, 2018. URL <http://arxiv.org/abs/1810.09305>.
- Ashutosh Kumar, Kabir Ahuja, Raghuram Vadapalli, and Partha Talukdar. Syntax-guided controlled generation of paraphrases. *Transactions of the Association for Computational Linguistics*, 8: 329–345, 2020. doi: 10.1162/tacl.a.00318. URL <https://aclanthology.org/2020.tacl-1.22>.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466, 2019. doi: 10.1162/tacl.a.00276. URL <https://aclanthology.org/Q19-1026>.
- Faisal Ladhak, Esin Durmus, Claire Cardie, and Kathleen McKeown. WikiLingua: A new benchmark dataset for cross-lingual abstractive summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 4034–4048, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.360. URL <https://aclanthology.org/2020.findings-emnlp.360>.
- Huiyuan Lai, Antonio Toral, and Malvina Nissim. Thank you BART! rewarding pre-trained models improves formality style transfer. 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)*, pp. 484–494, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-short.62. URL <https://aclanthology.org/2021.acl-short.62>.
- Guillaume Lample and Alexis Conneau. Cross-lingual language model pretraining. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/c04c19c2c2474dbf5f7ac4372c5b9af1-Paper.pdf>.
- Rémi Lebret, David Grangier, and Michael Auli. Neural text generation from structured data with application to the biography domain. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1203–1213, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1128. URL <https://aclanthology.org/D16-1128>.



- Sungjin Lee, Hannes Schulz, Adam Atkinson, Jianfeng Gao, Kaheer Suleman, Layla El Asri, Mahmoud Adada, Minlie Huang, Shikhar Sharma, Wendy Tay, and Xiujun Li. Multi-domain task-completion dialog challenge. In *Dialog System Technology Challenges*, volume 8, 2019. URL [http://workshop.colips.org/dstc7/dstc8/DTSC8\\_multidomain\\_task\\_proposal.pdf](http://workshop.colips.org/dstc7/dstc8/DTSC8_multidomain_task_proposal.pdf).
- Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 3045–3059, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.243. URL <https://aclanthology.org/2021.emnlp-main.243>.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 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*, pp. 7871–7880, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.703. URL <https://aclanthology.org/2020.acl-main.703>.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 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*, pp. 110–119, San Diego, California, June 2016. Association for Computational Linguistics. doi: 10.18653/v1/N16-1014. URL <https://aclanthology.org/N16-1014>.
- Junyi Li, Tianyi Tang, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. A survey of pretrained language models based text generation. *arXiv preprint arXiv:2201.05273*, 2022. URL <https://arxiv.org/abs/2201.05273>.
- Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 4582–4597, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.353. URL <https://aclanthology.org/2021.acl-long.353>.
- Xiujun Li, Yu Wang, Siqi Sun, Sarah Panda, Jingjing Liu, and Jianfeng Gao. Microsoft dialogue challenge: Building end-to-end task-completion dialogue systems. *arXiv preprint arXiv:1807.11125*, 2018. URL <http://arxiv.org/abs/1807.11125>.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. DailyDialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 986–995, Taipei, Taiwan, November 2017. Asian Federation of Natural Language Processing. URL <https://aclanthology.org/I17-1099>.
- Percy Liang, Michael Jordan, and Dan Klein. Learning semantic correspondences with less supervision. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pp. 91–99, Suntec, Singapore, August 2009. Association for Computational Linguistics. URL <https://aclanthology.org/P09-1011>.
- Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. CommonGen: A constrained text generation challenge for generative commonsense reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 1823–1840, Online, November 2020a. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.165. URL <https://aclanthology.org/2020.findings-emnlp.165>.
- Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pp. 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL <https://aclanthology.org/W04-1013>.

- Zehui Lin, Xiao Pan, Mingxuan Wang, Xipeng Qiu, Jiangtao Feng, Hao Zhou, and Lei Li. Pre-training multilingual neural machine translation by leveraging alignment information. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2649–2663, Online, November 2020b. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.210. URL <https://aclanthology.org/2020.emnlp-main.210>.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, and Pascale Fung. MinTL: Minimalist transfer learning for task-oriented dialogue systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 3391–3405, Online, November 2020c. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.273. URL <https://aclanthology.org/2020.emnlp-main.273>.
- Pierre Lison, Jörg Tiedemann, and Milen Kouylekov. OpenSubtitles2018: Statistical rescoring of sentence alignments in large, noisy parallel corpora. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May 2018. European Language Resources Association (ELRA). URL <https://aclanthology.org/L18-1275>.
- 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. GLGE: A new general language generation evaluation benchmark. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 408–420, Online, August 2021a. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.36. URL <https://aclanthology.org/2021.findings-acl.36>.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*, 2021b. URL <http://arxiv.org/abs/2107.13586>.
- Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 61–68, Dublin, Ireland, May 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.acl-short.8>.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. Multi-task deep neural networks for natural language understanding. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4487–4496, Florence, Italy, July 2019a. Association for Computational Linguistics. doi: 10.18653/v1/P19-1441. URL <https://aclanthology.org/P19-1441>.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019b. URL <http://arxiv.org/abs/1910.10683>.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742, 2020. doi: 10.1162/tacl.a.00343. URL <https://aclanthology.org/2020.tacl-1.47>.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.
- Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. *Gender Bias in Neural Natural Language Processing*, pp. 189–202. Springer International Publishing, Cham, 2020. doi: 10.1007/978-3-030-62077-6\_14. URL [https://doi.org/10.1007/978-3-030-62077-6\\_14](https://doi.org/10.1007/978-3-030-62077-6_14).

- Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. Learned in translation: Contextualized word vectors. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/20c86a628232a67e7bd46f76fba7ce12-Paper.pdf>.
- Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. The natural language decathlon: Multitask learning as question answering. *arXiv preprint arXiv:1806.08730*, 2018. URL <http://arxiv.org/abs/1806.08730>.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. Cross-task generalization via natural language crowdsourcing instructions. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3470–3487, Dublin, Ireland, May 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.acl-long.244>.
- Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. OpenDialKG: Explainable conversational reasoning with attention-based walks over knowledge graphs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 845–854, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1081. URL <https://aclanthology.org/P19-1081>.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. A corpus and cloze evaluation for deeper understanding of commonsense stories. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 839–849, San Diego, California, June 2016. Association for Computational Linguistics. doi: 10.18653/v1/N16-1098. URL <https://aclanthology.org/N16-1098>.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. Neural belief tracker: Data-driven dialogue state tracking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1777–1788, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1163. URL <https://aclanthology.org/P17-1163>.
- 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, Faiyaz 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. 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*, pp. 432–447, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.37. URL <https://aclanthology.org/2021.naacl-main.37>.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 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*, pp. 1797–1807, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1206. URL <https://aclanthology.org/D18-1206>.
- Thong Nguyen, Anh Tuan Luu, Truc Lu, and Tho Quan. Enriching and controlling global semantics for text summarization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 9443–9456, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.744. URL <https://aclanthology.org/2021.emnlp-main.744>.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. MS MARCO: A human generated machine reading comprehension dataset. In *CoCo@NIPS*, volume 1773 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2016. URL [http://ceur-ws.org/Vol-1773/CoCoNIPS\\_2016\\_paper9.pdf](http://ceur-ws.org/Vol-1773/CoCoNIPS_2016_paper9.pdf).

- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. The E2E dataset: New challenges for end-to-end generation. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pp. 201–206, Saarbrücken, Germany, August 2017. Association for Computational Linguistics. doi: 10.18653/v1/W17-5525. URL <https://aclanthology.org/W17-5525>.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073135. URL <https://aclanthology.org/P02-1040>.
- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. True few-shot learning with language models. In *Advances in Neural Information Processing Systems*, volume 34, pp. 11054–11070. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/file/5c04925674920eb58467fb52ce4ef728-Paper.pdf>.
- Guanghui Qin and Jason Eisner. Learning how to ask: Querying LMs with mixtures of soft prompts. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 5203–5212, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.410. URL <https://aclanthology.org/2021.naacl-main.410>.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. URL [https://cdn.openai.com/better-language-models/language\\_models\\_are\\_unsupervised\\_multitask\\_learners.pdf](https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. URL <http://jmlr.org/papers/v21/20-074.html>.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 2383–2392, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1264. URL <https://aclanthology.org/D16-1264>.
- Sudha Rao and Joel Tetreault. Dear sir or madam, may I introduce the GYAFC dataset: Corpus, benchmarks and metrics for formality style transfer. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 129–140, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1012. URL <https://aclanthology.org/N18-1012>.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. Towards empathetic open-domain conversation models: A new benchmark and dataset. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 5370–5381, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1534. URL <https://aclanthology.org/P19-1534>.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. volume 34, pp. 8689–8696, Apr. 2020a. doi: 10.1609/aaai.v34i05.6394. URL <https://ojs.aaai.org/index.php/AAAI/article/view/6394>.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. volume 34, pp. 8689–8696, Apr. 2020b. doi: 10.1609/aaai.v34i05.6394. URL <https://ojs.aaai.org/index.php/AAAI/article/view/6394>.

- Mathieu Ravaut, Shafiq Joty, and Nancy Chen. SummaReranker: A multi-task mixture-of-experts re-ranking framework for abstractive summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 4504–4524, Dublin, Ireland, May 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.acl-long.309>.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. CoQA: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266, 2019. doi: 10.1162/tacl.a.00266. URL <https://aclanthology.org/Q19-1016>.
- Pedro Rodriguez, Paul Crook, Seungwhan Moon, and Zhiguang Wang. Information seeking in the spirit of learning: A dataset for conversational curiosity. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 8153–8172, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.655. URL <https://aclanthology.org/2020.emnlp-main.655>.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. A neural attention model for abstractive sentence summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 379–389, Lisbon, Portugal, September 2015. Association for Computational Linguistics. doi: 10.18653/v1/D15-1044. URL <https://aclanthology.org/D15-1044>.
- Ananya B. Sai, Akash Kumar Mohankumar, and Mitesh M. Khapra. A survey of evaluation metrics used for nlg systems. *ACM Comput. Surv.*, 55(2), jan 2022. ISSN 0360-0300. doi: 10.1145/3485766. URL <https://doi.org/10.1145/3485766>.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M Rush. Multitask prompted training enables zero-shot task generalization. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=9Vrb9D0WI4>.
- Maarten Sap, Eric Horvitz, Yejin Choi, Noah A. Smith, and James Pennebaker. Recollection versus imagination: Exploring human memory and cognition via neural language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 1970–1978, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.178. URL <https://aclanthology.org/2020.acl-main.178>.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 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)*, pp. 1073–1083, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1099. URL <https://aclanthology.org/P17-1099>.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1631–1642, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL <https://aclanthology.org/D13-1170>.
- Karl Stratos. Mutual information maximization for simple and accurate part-of-speech induction. 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)*, pp. 1095–1104, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1113. URL <https://aclanthology.org/N19-1113>.
- Yixuan Su, David Vandyke, Sihui Wang, Yimai Fang, and Nigel Collier. Plan-then-generate: Controlled data-to-text generation via planning. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pp. 895–909, Punta Cana, Dominican Republic, November 2021.

- Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-emnlp.76. URL <https://aclanthology.org/2021.findings-emnlp.76>.
- Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, and Yi Zhang. Multi-task pre-training for plug-and-play task-oriented dialogue system. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 4661–4676, Dublin, Ireland, May 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.acl-long.319>.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J Pal. Learning general purpose distributed sentence representations via large scale multi-task learning. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=B18WgG-CZ>.
- Jiao Sun, Xuezhe Ma, and Nanyun Peng. AESOP: Paraphrase generation with adaptive syntactic control. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 5176–5189, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.420. URL <https://aclanthology.org/2021.emnlp-main.420>.
- Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. DREAM: A challenge data set and models for dialogue-based reading comprehension. *Transactions of the Association for Computational Linguistics*, 7:217–231, 2019. doi: 10.1162/tacl.a.00264. URL <https://aclanthology.org/Q19-1014>.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2818–2826, Los Alamitos, CA, USA, jun 2016. IEEE Computer Society. doi: 10.1109/CVPR.2016.308. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2016.308>.
- Xiangru Tang, Arjun Nair, Borui Wang, Bingyao Wang, Jai Desai, Aaron Wade, Haoran Li, Asli Celikyilmaz, Yashar Mehdad, and Dragomir Radev. CONFIT: Toward faithful dialogue summarization with linguistically-informed contrastive fine-tuning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 5657–5668, Seattle, United States, July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-main.415. URL <https://aclanthology.org/2022.naacl-main.415>.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. NewsQA: A machine comprehension dataset. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pp. 191–200, Vancouver, Canada, August 2017. Association for Computational Linguistics. doi: 10.18653/v1/W17-2623. URL <https://aclanthology.org/W17-2623>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4566–4575, Los Alamitos, CA, USA, jun 2015. IEEE Computer Society. doi: 10.1109/CVPR.2015.7299087. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2015.7299087>.
- Tu Vu, Brian Lester, Noah Constant, Rami Al-Rfou’, and Daniel Cer. SPoT: Better frozen model adaptation through soft prompt transfer. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5039–5059, Dublin, Ireland, May 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.acl-long.346>.

- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJ4km2R5t7>.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641, 2019. doi: 10.1162/tacl.a.00290. URL <https://aclanthology.org/Q19-1040>.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=gEzrGCozdqR>.
- Anuradha Welivita, Yubo Xie, and Pearl Pu. A large-scale dataset for empathetic response generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 1251–1264, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.96. URL <https://aclanthology.org/2021.emnlp-main.96>.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gašić, Lina M. Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. A network-based end-to-end trainable task-oriented dialogue system. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pp. 438–449, Valencia, Spain, April 2017. Association for Computational Linguistics. URL <https://aclanthology.org/E17-1042>.
- Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1112–1122, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1101. URL <https://aclanthology.org/N18-1101>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.6. URL <https://aclanthology.org/2020.emnlp-demos.6>.
- Dongling Xiao, Han Zhang, Yukun Li, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 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-20*, pp. 3997–4003. International Joint Conferences on Artificial Intelligence Organization, 7 2020. doi: 10.24963/ijcai.2020/553. URL <https://doi.org/10.24963/ijcai.2020/553>. Main track.
- Tianbao Xie, Chen Henry Wu, Peng Shi, Ruiqi Zhong, Torsten Scholak, Michihiro Yasunaga, Chien-Sheng Wu, Ming Zhong, Pengcheng Yin, Sida I Wang, et al. Unifiedskg: Unifying and multi-tasking structured knowledge grounding with text-to-text language models. *arXiv preprint arXiv:2201.05966*, 2022. URL <http://arxiv.org/abs/2201.05966>.
- Hanwei Xu, Yujun Chen, Yulun Du, Nan Shao, Yanggang Wang, Haiyu Li, and Zhilin Yang. Zeroprompt: Scaling prompt-based pretraining to 1,000 tasks improves zero-shot generalization. *arXiv preprint arXiv:2201.06910*, 2022. URL <http://arxiv.org/abs/2201.06910>.
- Peng Xu, Davis Liang, Zhiheng Huang, and Bing Xiang. Attention-guided generative models for extractive question answering. *arXiv preprint arXiv:2110.06393*, 2021. URL <http://arxiv.org/abs/2110.06393>.

- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics*, 4:401–415, 2016. doi: 10.1162/tacl.a.00107. URL <https://aclanthology.org/Q16-1029>.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2369–2380, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1259. URL <https://aclanthology.org/D18-1259>.
- Qinyuan Ye, Bill Yuchen Lin, and Xiang Ren. CrossFit: A few-shot learning challenge for cross-task generalization in NLP. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 7163–7189, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.572. URL <https://aclanthology.org/2021.emnlp-main.572>.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 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)*, pp. 2204–2213, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1205. URL <https://aclanthology.org/P18-1205>.
- Yian Zhang, Alex Warstadt, Xiaocheng Li, and Samuel R. Bowman. When do you need billions of words of pretraining data? In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1112–1125, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.90. URL <https://aclanthology.org/2021.acl-long.90>.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 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*, pp. 270–278, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-demos.30. URL <https://aclanthology.org/2020.acl-demos.30>.
- Kangyan Zhou, Shrimai Prabhunoye, and Alan W Black. A dataset for document grounded conversations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 708–713, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1076. URL <https://aclanthology.org/D18-1076>.
- Chenguang Zhu, Yang Liu, Jie Mei, and Michael Zeng. MediaSum: A large-scale media interview dataset for dialogue summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 5927–5934, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.474. URL <https://aclanthology.org/2021.naacl-main.474>.



## A LIMITATIONS

Despite our efforts to collect as many generation tasks and datasets as possible, we only evaluate the generation quality and generalization ability of our models on a small number of tasks and datasets. The interpretability and robustness of our models require further analysis. Besides, there exists subjectivity when collecting intra-task datasets, albeit our attempts to employ widely-recognized categorizations from the literature. Due to limitation of computing power, we do not study the performance of our method at different model scales. The effectiveness of multi-task pre-training from scratch, similar to ExT5 (Aribandi et al., 2022), also merits an in-depth study. Regarding evaluation methods, we only consider basic automatic metrics such as BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004). However, there is still a certain gap between these metrics and human judgments (Sai et al., 2022).

## B TASKS AND DATASETS

### B.1 DESCRIPTION OF TASKS AND DATASETS

We provide the details of the tasks and datasets used in our paper for pre-training and fine-tuning in Tables 7 and 8. If the dataset for pre-training does not have a valid set, we divide 10% of the training set for validation.

We list the licenses for all datasets if they have. All datasets are publicly available. The majority of them can be directly downloaded from GitHub or Google Drive. ROCStories (Mostafazadeh et al., 2016) and CommonGen (Lin et al., 2020a) can be obtained after filling out a form. GYAFC (Rao & Tetreault, 2018) is accessible after requesting Yahoo and the authors of the dataset.

The tasks and datasets we use in this paper are as follows:

- **Data-to-text generation** aims to generate descriptive text about structured data, such as the knowledge graph and the table. We use the following datasets for pre-training:

1. AGENDA (Koncel-Kedziorski et al., 2019);
2. ENT-DESC (Cheng et al., 2020);
3. GenWiki (Jin et al., 2020);
4. LogicNLG (Chen et al., 2020a);
5. TEKGEN (Agarwal et al., 2021);
6. WEATHERGOV (Liang et al., 2009);
7. WikiTableT (Chen et al., 2021).

We utilize the following datasets for fine-tuning evaluation:

1. WebNLG (Gardent et al., 2017), we utilize the version 2.1;
2. WikiBio (Lebret et al., 2016).

- **Open-ended dialogue system**, also known as chatbots, is focused on daily communication. We use the following datasets for pre-training:

1. Cleaned OpenSubtitles Dialogs (Cleaned OS Dialogs) (Welivita et al., 2021), which is a cleaned variant of OpenSubtitles Dialogs (Lison et al., 2018);
2. CMU Document Grounded Conversations (CMUDog) (Zhou et al., 2018);
3. Curiosity (Rodriguez et al., 2020);
4. DREAM (Sun et al., 2019);
5. Empathetic Dialogues (Rashkin et al., 2019);
6. Movie Dialog (Dodge et al., 2016);
7. MuTual (Stratos, 2019);
8. OpenDialKG (Moon et al., 2019);
9. Topical-Chat (Gopalakrishnan et al., 2019);
10. Wizard of Wikipedia (Dinan et al., 2019).

We utilize the following datasets for fine-tuning evaluation:

1. DailyDialog (Li et al., 2017);
  2. DSTC7-AVSD (Alamri et al., 2018);
  3. PersonaChat (Zhang et al., 2018).
- **Paraphrase generation** involves rewriting a sentence with the same semantic meaning but a different syntactic or lexical form. We utilize the following datasets for fine-tuning evaluation:
    1. Quora (also known as QQP-Pos) (Kumar et al., 2020), which is a subset of Quora Question Pairs <sup>3</sup>.
  - **Question answering** requires the model to answer a question based on optional background information. Note that we conduct this task in a generative way in our paper. We use the following datasets for pre-training:
    1. HotpotQA (Yang et al., 2018);
    2. MS MARCO (Nguyen et al., 2016);
    3. MSQG (Liu et al., 2021a), since it is designed for QG, we reverse the question and answer to enrich QA examples;
    4. NarrativeQA (Kočíský et al., 2018);
    5. Natural Questions (Kwiatkowski et al., 2019);
    6. NewsQA (Trischler et al., 2017);
    7. QuAC (Choi et al., 2018);
    8. TriviaQA (Joshi et al., 2017);
    9. WebQuestions (Berant et al., 2013).

We utilize the following datasets for fine-tuning evaluation:

1. CoQA (Reddy et al., 2019);
  2. SQuAD (Rajpurkar et al., 2016), we utilize the version 1.1.
- **Question generation** generates a coherent question given a passage and its corresponding answer. We use the following datasets for pre-training:
    1. HotpotQA (Yang et al., 2018);
    2. MS MARCO (Nguyen et al., 2016);
    3. MSQG (Liu et al., 2021a);
    4. NarrativeQA (Kočíský et al., 2018);
    5. NewsQA (Trischler et al., 2017);
    6. QuAC (Choi et al., 2018);

Most of them are QA tasks, and we invert the question and answer to enrich QG examples.

We utilize the following datasets for fine-tuning evaluation:

1. CoQA (Reddy et al., 2019);
  2. SQuAD (Rajpurkar et al., 2016), we utilize the version 1.1.
- **Story generation** creates a long and informative text with a short title. We use the following datasets for pre-training:
    1. ChangeMyView (Hua & Wang, 2020);
    2. English Gigaword (Rush et al., 2015);
    3. Hippocorpus (Sap et al., 2020);
    4. WikiPlots (wik);
    5. WritingPrompts (Fan et al., 2018), we split the original training set for pre-training and corresponding validation.

Considering English Gigaword is a large summarization dataset, we use the summary as the title to generate the passage in turn to enrich the examples of story generation.

We utilize the following datasets for fine-tuning evaluation:

1. ROCStories (Mostafazadeh et al., 2016);

---

<sup>3</sup><https://www.kaggle.com/c/quora-question-pairs>

2. WritingPrompts (Fan et al., 2018), we use the sets created by Guan et al. (2021) (who split the original valid and test sets for training, validation, and testing) to fine-tune our model for a fair comparison.
- **Task-oriented dialogue system** meets real-life needs of users, such as restaurant reservations and airplane bookings. We use the datasets for pre-training, following Su et al. (2022):
    1. CamRest676 (Wen et al., 2017);
    2. Frames (El Asri et al., 2017);
    3. KVRET (Eric et al., 2017);
    4. MetaLWOZ (Lee et al., 2019);
    5. MSR-E2E (Li et al., 2018);
    6. MultiWOZ (Budzianowski et al., 2018);
    7. Schema-Guided (Rastogi et al., 2020a);
    8. TaskMaster (Byrne et al., 2019);
    9. WOZ (Mrkšić et al., 2017).

We utilize the following datasets for fine-tuning evaluation:

1. MultiWOZ (Budzianowski et al., 2018), we utilize the version 2.0;
- **Text style transfer** modifies the style (*e.g.*, sentiment and formality) of given texts while retaining their style-independent content. We utilize the following datasets for fine-tuning evaluation:
    1. GYAFC (Rao & Tetreault, 2018), which has two sub-domains “Entertainment and Music” (E&M) and “Family and Relationships” (F&R).
  - **Text summarization** condenses a long document into a brief text while retaining the essential details. We use the following datasets for pre-training:
    1. English Gigaword (Graff et al., 2003), we use the variant provided by Rush et al. (2015);
    2. MediaSum (Zhu et al., 2021);
    3. MSNews (Liu et al., 2021a);
    4. Newsroom (Grusky et al., 2018);
    5. WikiHow (Koupae & Wang, 2018).

We utilize the following datasets for fine-tuning evaluation:

1. CNN/DailyMail (Hermann et al., 2015), we use the variant provided by See et al. (2017);
2. SAMSum (Gliwa et al., 2019);
3. XSum (Narayan et al., 2018).

To better compare with Ext5 (Aribandi et al., 2022), we utilize the language generation benchmark GEM (Gehrmann et al., 2021) for fine-tuning evaluation. GEM includes five tasks:

- **Commonsense generation:**
  1. CommonGen (CG) (Lin et al., 2020a).
- **Data-to-text generation:**
  1. DART (Nan et al., 2021);
  2. E2E NLG cleaned (Novikova et al., 2017);
  3. ToTTo (Su et al., 2021);
  4. WebNLG (Gardent et al., 2017).
- **Dialogue system:**
  1. Schema-Guided Dialog (SGD) (Rastogi et al., 2020b).
- **Text simplification:**
  1. WikiAuto + Turk/ASSET (WiA-T/A) (Jiang et al., 2020; Xu et al., 2016; Alva-Manchego et al., 2020).
- **Text summarization:**
  1. Wiki-Lingua (WLE) (Ladhak et al., 2020).

To test the generalization ability of our model, we also utilize the natural language standing benchmark GLUE (Wang et al., 2019), which is composed of three tasks:

- **Natural language inference:**

1. MNLI (Williams et al., 2018);
2. QNLI (Rajpurkar et al., 2016; Wang et al., 2019);
3. RTE (Dagan et al., 2006; Haim et al., 2006; Giampiccolo et al., 2007; Bentivogli et al., 2009).

- **Paraphrase detection:**

1. MRPC (Dolan & Brockett, 2005);
2. QQP<sup>3</sup>;
3. STS-B (Cer et al., 2017).

- **Text classification:**

1. CoLA (Warstadt et al., 2019);
2. SST-2 (Socher et al., 2013).

## B.2 DATA LEAKAGE

Since our model is pre-trained on a large number of labeled datasets, it may have “seen” examples from fine-tuning test sets during pre-training, which leads to an unfair comparison with other methods. Hence, we eliminate the pre-training examples that share  $n$ -gram overlap with either of the test datasets. Following Brown et al. (2020),  $n$  is the 5<sup>th</sup> percentile example length in words, and the maximum value of  $n$  is set to 13. Finally, we have removed 17,848 examples from the pre-training datasets. The number of “cleaned” examples for each dataset can be found in Table 7.

Table 7: The statistics and licenses of datasets for pre-training our MVP model. The #Train, #Valid, and #Test denote the number of examples in the train, valid, and test sets, respectively. Cleaned #Train represents the number of training examples after filtering. Input and Output are the average number of words (split by space) in the input and output sequences, respectively. These setups and abbreviations are the same below.

Dataset	#Train	Cleaned #Train	#Valid	#Test	Input	Output	License
AGENDA	38,720	38,720	1,000	1,000	52.1	141.2	N/A
ENT-DESC	88,652	88,652	11,081	11,081	279.9	31.0	N/A
GenWiki	681,436	681,436	75,716	1,000	21.4	29.5	MIT
LogicNLG	28,450	28,450	4,260	4,305	178.4	14.2	MIT
TEKGEN	6,310,061	6,307,995	788,746	796,982	17.0	21.2	CC BY-SA 2.0
WEATHERGOV	25,000	25,000	1,000	3,528	148.7	30.6	N/A
WikiTableT	1,453,794	1,452,778	4,533	4,351	81.0	99.7	MIT
Cleaned OS Dialogs	13,355,487	13,355,368	1,483,944	-	75.5	16.7	N/A
CMUDoG	82,818	82,818	5,555	14,510	433.0	12.2	N/A
Curiosity	64,930	64,551	8,539	8,495	144.4	20.2	CC BY-NC 4.0
DREAM	14,264	14,242	4,709	4,766	75.6	13.6	N/A
Empathetic Dialogues	64,636	64,636	9,308	8,426	52.7	12.9	CC BY-NC 4.0
Movie Dialog	762,751	762,711	8,216	8,066	126.9	44.0	N/A
MuTual	33,691	33,691	4,090	3,248	53.6	14.5	N/A
OpenDialKG	69,680	69,680	7,743	-	54.2	12.4	CC BY-NC 4.0
Topical-Chat	179,750	179,750	22,295	22,452	223.3	20.0	CDLA-Sharing-1.0
Wizard of Wikipedia	148,357	147,702	15,767	15,564	297.0	16.7	MIT
HotpotQA	90,447	87,815	7,405	-	187.9	2.2	CC BY-SA 4.0
MS MARCO	681,445	681,226	77,580	-	68.7	13.3	N/A
MSQG	198,058	198,029	11,008	-	48.1	3.7	CC BY-SA 4.0
NarrativeQA	65,494	65,494	6,922	21,114	584.1	4.2	Apache 2.0
Natural Questions	96,676	96,676	10,693	6,490	9.0	2.1	CC BY-SA 3.0
NewsQA	97,850	97,700	5,486	5,396	726.8	5.0	MIT
QuAC	83,568	83,485	31,906	-	487.9	12.5	CC BY-SA 4.0
TriviaQA	78,785	78,785	8,837	11,313	14.0	2.0	Apache 2.0
WebQuestions	8,933	8,933	4,863	4,863	6.7	2.4	CC BY 4.0
HotpotQA	90,440	87,808	6,972	-	79.6	19.8	CC BY-SA 4.0
MS MARCO	681,445	681,226	77,580	-	75.9	6.0	N/A
MSQG	198,058	198,029	11,008	11,022	45.9	6.0	CC BY-SA 4.0
NarrativeQA	65,494	65,494	6,922	21,114	579.7	8.6	Apache 2.0
NewsQA	97,850	97,700	5,486	5,396	724.2	7.6	MIT
QuAC	69,109	69,026	26,301	-	496.7	6.5	CC BY-SA 4.0
ChangeMyView	42,462	42,459	6,480	7,562	17.9	104.1	MIT
English Gigaword	3,803,957	3,802,620	189,651	1,951	8.8	33.3	MIT
Hippocorpus	6,168	6,168	686	-	34.1	262.6	CDLA-Permissive 2.0
WikiPlots	101,642	101,641	11,294	-	3.4	338.5	N/A
WritingPrompts	272,600	272,518	15,620	15,138	28.4	630.8	MIT
CamRest676	4,872	4,872	616	-	55.3	9.4	N/A
Frames	26,631	26,631	2,106	-	116.1	13.0	MIT
KVRET	14,136	14,136	1,616	-	30.5	9.3	N/A
MetalWOZ	176,073	176,073	17,912	-	45.6	8.0	N/A
MSR-E2E	103,362	103,362	5,235	-	51.3	12.8	Microsoft
Schema-Guided	494,946	494,933	73,089	-	120.8	12.5	CC BY-SA 4.0
TaskMaster	249,664	249,662	20,680	-	95.6	12.0	CC BY 4.0
WOZ	6,364	6,359	1,260	-	47.0	10.6	N/A
English Gigaword	3,803,957	3,802,620	189,651	1,951	33.3	8.8	MIT
MediaSum	443,596	442,021	10,000	10,000	1641.0	14.4	N/A
MSNews	136,082	135,937	7,496	7,562	309.9	9.8	CC BY-SA 4.0
Newsroom	995,041	989,351	108,837	108,862	642.4	26.7	N/A
WikiHow	157,252	157,247	5,599	5,577	502.6	45.6	CC BY-NC-SA

Table 8: The statistics and licenses of datasets for evaluating our MVP model. The license of the MNLi dataset is composed of OANC, CC BY-SA 3.0, and CC BY 3.0. The license of the CoQA dataset is composed of CC BY-SA 4.0, MSR-LA, and Apache 2.0. The license of the WiA-A/T datasets is composed of CC BY-NC 3.0, CC BY-NC 4.0, and GNU General Public License v3.0.

Task	Dataset	#Train	#Valid	#Test	Input	Output	License
<b>Commonsense generation</b>	CommonGen	67,389	993	–	5.5	11.6	MIT
<b>Data-to-text generation</b>	DART	62,659	2,768	–	27.5	21.5	MIT
	E2E	33,525	4,299	–	9.5	20.6	CC BY-SA 4.0
	ToTTo	120,761	7,700	–	37.8	18.0	CC BY-SA 3.0
	WebNLG	34,338	4,313	4,222	18.0	19.9	CC BY-NA-SA 4.0
	WebNLG (GEM)	35,426	1,667	–	17.7	22.7	CC BY-NA-SA 4.0
	WikiBio	582,659	72,831	72,831	81.6	26.1	CC BY-SA 3.0
<b>Open-ended dialogue</b>	DailyDialog	76,052	7,069	6,740	72.5	13.9	CC BY-NC-SA 4.0
	DSTC7-AVSD	76,590	17,870	1,710	148.2	11.5	MIT
	PersonaChat	122,499	14,602	14,056	132.1	11.9	MIT
	SGD	164,982	10,000	–	134.7	11.3	CC BY-SA 4.0
<b>Natural language inference</b>	MNLi-m	–	9,815	9,796	–	–	Mixed
	MNLi-mm	392,702	9,832	9,847	29.8	–	Mixed
	QNLI	104,743	5,463	5,463	36.6	–	CC BY-SA 4.0
	RTE	2,490	277	3,000	51.0	–	N/A
<b>Paraphrase generation</b>	Quora	137,185	3,000	3,000	10.9	10.8	N/A
<b>Paraphrase detection</b>	MRPC	3,668	408	1,725	43.8	–	N/A
	QQP	363,846	40,430	390,965	22.3	–	N/A
	STS-B	5,749	1,500	1,379	20.3	–	N/A
<b>Question answering</b>	CoQA	107,286	31,621	–	349.4	2.6	Mixed
	SQuAD	75,722	10,570	11,877	156.2	3.6	CC BY-SA 4.0
<b>Question generation</b>	CoQA	107,286	31,621	–	346.6	5.5	Mixed
	SQuAD	75,722	10,570	11,877	148.3	11.6	CC BY-SA 4.0
<b>Story generation</b>	ROCStories	176,688	9,816	4,909	9.0	40.7	N/A
	WritingPrompts	53,516	4,000	2,000	25.5	150.4	MIT
<b>Task-oriented dialogue</b>	MultiWOZ	170,220	22,074	22,116	128.3	11.3	MIT
<b>Text classification</b>	CoLA	8,551	1,043	1,063	7.7	–	N/A
	SST-2	67,349	872	1,821	9.8	–	N/A
<b>Text simplification</b>	WiA-A WiA-T	483,801	20,000	359 359	26.2	21.5	Mixed
<b>Text style transfer</b>	GYAFC-E&M	52,595	11,508	1,416	9.9	10.6	N/A
	GYAFC-F&R	51,967	11,152	1,332	10.7	11.3	N/A
<b>Text summarization</b>	CNN/DailyMail	287,227	13,368	11,490	679.8	48.3	MIT
	SAMSum	14,732	818	819	103.4	20.3	CC BY-NC-ND 4.0
	WLE	99,020	28,614	–	367.6	33.4	CC0 1.0
	XSum	204,045	11,332	11,334	373.7	21.1	MIT

## C FINE-TUNING AND EVALUATION DETAILS

In this section, we introduce the details for fine-tuning and evaluating each downstream task.

For the experiments in Section 4 (Tables 2 and 3), and Appendix D.1 (Table 9), the fine-tuning details are introduced in Section 4, and the evaluation details are presented as follows:

- For data-to-text generation tasks, we use BLEU(-4), ROUGE-L, and METEOR for evaluation. We use the script provided by Chen et al. (2020b) <sup>4</sup>;
- For open-ended dialogue system tasks, we use BLEU-1, BLEU-2, Distinct-1, and Distinct-2 for evaluation. For DSTC7-AVSD we also utilize CIDEr (Vedantam et al., 2015). We employ NLTK 3.5 with smoothing function 7 to compute BLEU for PersonaChat and DailyDialog, and utilize the script <sup>5</sup> to evaluate DSTC7-AVSD;
- For question answering tasks, we use Exact Match (EM) and Macro-averaged F1 score (F1) for evaluation. We use the provided script for CoQA <sup>6</sup> and SQuAD <sup>7</sup>.
- For question generation tasks, we use BLEU-4, ROUGE-L, and METEOR for evaluation. We use the script provided by Dong et al. (2019) <sup>8</sup>;
- For story generation, we employ nucleus sampling with  $p = 0.9$  and temperature of 0.7 following Guan et al. (2021). We use corpus BLEU-1, BLEU-2, Distinct-1, and Distinct-4 for evaluation. We use NLTK 3.5 to calculate corpus BLEU following Guan et al. (2021);
- For task-oriented dialogue system tasks, we use BLEU(-4), inform (rate), success (rate), and combined score for evaluation. Inform and success are two specially designed accuracy metrics for task-oriented dialogue system, and the combined score is defined as  $(\text{Inform} + \text{Success}) \times 0.5 + \text{BLEU}$  (Budzianowski et al., 2018). We use the script provided by Su et al. (2022) <sup>9</sup>;
- For text summarization tasks, we use ROUGE-1, ROUGE-2, and ROUGE-L for evaluation. We use the toolkit `files2rouge` <sup>10</sup>.

For the experiments in Section 5 (Tables 4 and 5), the fine-tuning and evaluation details are as follows:

- For paraphrase generation tasks, we employ the fine-tuning and evaluation scripts provided by AESOP (Sun et al., 2021) <sup>11</sup>. The evaluation metrics are BLEU-4, ROUGE-1, ROUGE-2, ROUGE-L, and METEOR.
- For text style transfer tasks, we employ the fine-tuning and evaluation scripts provided by SC & BLEU (Lai et al., 2021) <sup>12</sup>. We conduct the informal-to-formal transfer and train the model on the data from both the E&M and F&R domains following Lai et al. (2021). The evaluation metrics are BLEU-4, accuracy, and HM. Accuracy is calculated by a pre-trained TextCNN to evaluate the style strength, and HM denotes the harmonic mean of BLEU-4 and style accuracy (Lai et al., 2021).
- For GLUE tasks, we utilize the fine-tuning code provided by Hugging Face <sup>13</sup>. The hyper-parameters are consistent with original BART (Lewis et al., 2020) <sup>14</sup>. The evaluation is computed by the official website <sup>15</sup>.

<sup>4</sup><https://github.com/wenhuchen/Data-to-text-Evaluation-Metric>

<sup>5</sup><https://github.com/lemuria-wchen/DialogVED/blob/main/src/utils/evaluate.py>

<sup>6</sup><https://github.com/PaddlePaddle/ERNIE/blob/repro/ernie-gen/eval/tasks/coqa/eval.py>

<sup>7</sup><https://github.com/allenai/bi-att-flow/blob/master/squad/evaluate-v1.1.py>

<sup>8</sup><https://github.com/microsoft/unilm/blob/master/unilm-v1/src/qg/eval.py>

<sup>9</sup>[https://github.com/aws-labs/ppod/blob/main/E2E\\_TOD/eval.py](https://github.com/aws-labs/ppod/blob/main/E2E_TOD/eval.py)

<sup>10</sup><https://github.com/pltrdy/files2rouge>

<sup>11</sup><https://github.com/PlusLabNLP/AESOP>

<sup>12</sup><https://github.com/laihuiyuan/pre-trained-formality-transfer>

<sup>13</sup><https://github.com/huggingface/transformers/tree/main/examples/pytorch/text-classification>

<sup>14</sup><https://github.com/facebookresearch/fairseq/blob/main/examples/bart/README.glue.md>

<sup>15</sup><https://gluebenchmark.com/>

Table 9: The results on six seen tasks under full tuning settings. <sup>a</sup> (Nguyen et al., 2021) <sup>b</sup> (Tang et al., 2022) <sup>c</sup> (Gu et al., 2021) <sup>d</sup> (Lewis et al., 2020) <sup>e</sup> (Guan et al., 2021) <sup>f</sup> (Chen et al., 2022) <sup>g</sup> (Chen et al., 2020b) <sup>h</sup> (Raffel et al., 2020) <sup>i</sup> (Xu et al., 2021)

Methods	XSum			SAMSum			CoQA QG		
	R-1	R-2	R-L	R-1	R-2	R-L	B-4	ME	R-L
SOTA	<b>49.57<sup>a</sup></b>	<b>25.08</b>	<b>41.81</b>	53.89 <sup>b</sup>	28.85	49.29	15.78 <sup>c</sup>	40.15	50.98
BART	45.14 <sup>d</sup>	22.27	37.25	51.74 <sup>b</sup>	26.46	48.72	12.34 <sup>c</sup>	35.78	46.88
MVP	<u>45.67</u>	22.53	<u>37.41</u>	<u>53.91</u>	<u>29.28</u>	<u>49.40</u>	<u>22.96</u>	<b>47.14</b>	54.84
MVP+S	45.59	<u>22.54</u>	37.39	<b>53.97</b>	<b>29.44</b>	<b>49.69</b>	<b>23.04</b>	<u>47.01</u>	<b>55.14</b>
Methods	WritingPrompts				DailyDialog			WikiBio	
	B-1	B-2	D-1	D-4	B-1	B-2	D-1	D-2	B-4
SOTA	22.40 <sup>e</sup>	8.40	–	31.30	46.10 <sup>f</sup>	40.70	4.10	22.20	45.10 <sup>g</sup>
BART	22.40 <sup>e</sup>	8.40	–	31.30	44.30 <sup>f</sup>	39.20	3.90	21.10	–
MVP	<b>31.81</b>	<b>12.80</b>	<u>2.58</u>	<u>69.45</u>	<b>52.34</b>	<b>43.93</b>	<u>6.39</u>	<u>35.65</u>	<b>48.42</b>
MVP+S	<u>29.18</u>	<u>11.11</u>	<b>3.71</b>	<b>78.02</b>	<u>51.04</u>	<u>42.87</u>	<b>6.70</b>	<b>36.84</b>	<u>48.19</u>
Methods	DSTC7-AVSD						SQuAD		
	B-1	B-2	B-3	B-4	ME	R-L	CIDEr	F1	EM
SOTA	83.20 <sup>f</sup>	70.50	59.80	50.60	31.40	63.80	1.391	<b>96.22<sup>h</sup></b>	<b>91.26</b>
BART	82.40 <sup>f</sup>	69.10	58.20	48.70	31.30	63.50	1.382	91.56 <sup>i</sup>	84.23
MVP	<b>83.92</b>	<u>70.83</u>	<u>60.06</u>	<u>50.64</u>	<b>31.95</b>	<b>64.87</b>	<b>1.429</b>	93.04	86.44
MVP+S	<u>83.82</u>	<b>70.95</b>	<b>60.23</b>	<b>50.87</b>	<u>31.49</u>	<u>64.29</u>	<u>1.409</u>	<u>93.21</u>	<u>86.78</u>

For the experiments of the GEM benchmark in Appendix D.2 (Table 10), the fine-tuning settings are the same as those described in Section 4. We use BLEU-4, ROUGE-2, and METEOR for evaluation. We use the GEM evaluation scripts <sup>16</sup>.

## D ADDITIONAL RESULTS

In this section, we provide additional results of our MVP model and other baselines.

### D.1 RESULTS OF COMMON DATASETS

We also conduct experiments on eight common datasets under full tuning settings. Due to space limits in Section 4, these results are shown in Table 9. We can see that these results share a similar trend to those in Section 4, and we achieve SOTA performances in 6 of 8 datasets.

### D.2 RESULTS ON THE GEM BENCHMARK

To better compare with ExT5 (Aribandi et al., 2022), we conduct experiments on the GEM benchmark (Gehrmann et al., 2021). For “unseen” commonsense generation and text simplification tasks, we utilize prompts of data-to-text generation and summarization, respectively. The results are presented in Table 10, and our MVP models outperform ExT5 in 26 out of 27 metrics.

### D.3 RESULTS WITHOUT FINE-TUNING

Considering our MVP model has already been pre-trained on several tasks, we conduct experiments on these “seen” tasks without fine-tuning our model. To some degree, this setting can be viewed as

<sup>16</sup><https://github.com/GEM-benchmark/GEM-metrics>



Table 10: The results on the GEM benchmark under full tuning settings. We utilize the large version of T5.1.1 and ExT5, and all the results of them are from Aribandi et al. (2022).

Methods	DART			E2E			ToTTo		
	B-4	R-2	ME	B-4	R-2	ME	B-4	R-2	ME
T5.1.1	34.31	45.22	36.30	<b>42.57</b>	46.60	38.20	39.79	49.90	36.80
ExT5	36.62	48.14	37.60	<u>42.25</u>	46.70	38.10	40.14	50.33	36.90
MVP	<b>39.13</b>	<b>48.92</b>	<b>38.53</b>	37.38	<b>47.96</b>	<b>39.39</b>	<u>50.58</u>	<u>55.24</u>	<u>41.27</u>
MVP+S	<u>38.83</u>	<u>48.49</u>	<u>38.41</u>	37.32	<u>47.40</u>	<u>38.90</u>	<b>50.69</b>	<b>55.52</b>	<b>41.29</b>
Methods	WebNLG			CommonGen			SGD		
	B-4	R-2	ME	B-4	R-2	ME	B-4	R-2	ME
T5.1.1	31.67	43.31	34.40	8.38	17.01	20.20	33.15	36.17	32.40
ExT5	35.03	48.17	36.50	9.68	19.04	21.40	34.74	37.77	33.00
MVP	<b>47.03</b>	<u>59.00</u>	<b>42.34</b>	<u>32.59</u>	<u>37.71</u>	<u>33.00</u>	<b>45.63</b>	<b>48.29</b>	<b>38.48</b>
MVP+S	<b>47.03</b>	<b>59.03</b>	<u>42.28</u>	<b>34.10</b>	<b>37.87</b>	<b>33.11</b>	<u>45.24</u>	<u>48.25</u>	<u>38.47</u>
Methods	WiA-A			WiA-T			WLE		
	B-4	R-2	ME	B-4	R-2	ME	B-4	R-2	ME
T5.1.1	29.30	38.37	30.10	42.12	50.52	36.2	15.55	20.47	19.60
ExT5	29.23	37.98	30.00	41.39	50.38	35.8	16.64	21.16	20.40
MVP	<b>71.55</b>	<b>70.88</b>	<b>48.19</b>	<b>91.73</b>	83.46	<b>57.34</b>	<b>18.80</b>	<b>22.84</b>	<u>21.95</u>
MVP+S	<u>70.37</u>	<u>70.65</u>	<u>47.70</u>	<u>91.12</u>	<b>83.59</b>	<u>56.95</u>	<u>18.52</u>	<u>22.57</u>	<b>22.02</b>

Table 11: The results on seven seen tasks without fine-tuning. Given that T0 has been pre-trained on the CNN/DailyMail dataset, we exclude their results to provide a fair comparison (denoted as “-”).

Methods	CNN/DailyMail			WebNLG			SQuAD (QG)			CoQA	
	R-1	R-2	R-L	B-4	ME	R-L	B-4	ME	R-L	F1	EM
FT BART	44.16	21.28	40.90	64.55	46.51	75.13	22.00	26.40	52.55	68.60	-
FT MVP	44.44	21.61	40.99	67.76	47.74	77.04	26.21	27.20	53.46	86.51	77.62
T0	-	-	-	1.40	10.20	18.43	3.06	12.43	14.91	13.30	6.60
MVP	<b>29.50</b>	<b>11.29</b>	<b>25.92</b>	34.42	31.33	52.33	2.90	13.94	15.48	29.40	18.20
MVP+S	25.60	9.51	22.67	<b>39.43</b>	<b>34.32</b>	<b>55.34</b>	<b>2.96</b>	<b>15.23</b>	<b>18.23</b>	<b>52.40</b>	<b>37.30</b>
Methods	ROCStories				PersonaChat				MultiWOZ		
	B-1	B-2	D-1	D-4	B-1	B-2	D-1	D-2	B-4	Success	Inform
FT BART	30.70	13.30	-	69.90	49.90	40.00	1.30	8.00	17.89	74.91	84.88
FT MVP	33.42	15.54	2.92	75.06	50.07	40.54	1.54	9.86	20.04	79.20	87.20
T0	8.69	3.02	4.37	35.49	23.20	23.57	2.56	12.06	0.02	<b>2.50</b>	22.10
MVP	1.01	0.31	<b>7.18</b>	<b>86.26</b>	35.54	32.71	<b>2.87</b>	<b>16.38</b>	<b>3.08</b>	<b>2.50</b>	<b>22.20</b>
MVP+S	<b>10.52</b>	<b>3.54</b>	2.13	69.55	<b>37.04</b>	<b>33.38</b>	<u>2.66</u>	<u>14.84</u>	0.38	<b>2.50</b>	22.10

zero-shot learning. Nonetheless, it does not conform to the definition of *true zero-shot* settings (Perez et al., 2021). To avoid controversy, we refer to this as *without fine-tuning*.

We include T0-3B (Sanh et al., 2022) as our baseline. The results are listed in Table 11. Our MVP model outperforms T0 in all metrics with a large margin. However, all tasks demonstrate that methods without fine-tuning perform significantly worse than those with full tuning settings. This suggests that zero-shot strategies that are effective for NLU tasks may not produce satisfactory results for NLG tasks. Even though our model has acquired task knowledge, it struggles to perform well in a new domain without being fine-tuned. Thus, we focus mainly on full tuning settings in this paper.

## E QUALITATIVE EXAMPLES

In this section, we showcase the linearized inputs, human-written task prompts, and corresponding outputs of a single dataset for tasks in Section 4. We provide the results of BART, MVP, and MVP+S under full tuning settings. To minimize human intervention, we select the first and second instances of the test set.

Table 12: The first instance from the CNN/Daily Mail dataset. Human-written task prompts are labeled in *italic*. The setting is the same below.

---

**Input**

*Summarize:* Marseille, France (CNN)The French prosecutor leading an investigation into the crash of Germanwings Flight 9525 insisted Wednesday that he was not aware of any video footage from on board the plane. Marseille prosecutor Brice Robin told CNN that "so far no videos were used in the crash investigation." He added, "A person who has such a video needs to immediately give it to the investigators." Robin's comments follow claims by two magazines, German daily Bild and French Paris Match, of a cell phone video showing the harrowing final seconds from on board Germanwings Flight 9525 as it crashed into the French Alps. All 150 on board were killed. Paris Match and Bild reported that the video was recovered from a phone at the wreckage site. The two publications described the supposed video, but did not post it on their websites. The publications said that they watched the video, which was found by a source close to the investigation. "One can hear cries of 'My God' in several languages," Paris Match reported. "Metallic banging can also be heard more than three times, perhaps of the pilot trying to open the cockpit door with a heavy object. Towards the end, after a heavy shake, stronger than the others, the screaming intensifies. Then nothing." "It is a very disturbing scene," said Julian Reichelt, editor-in-chief of Bild online. An official with France's accident investigation agency, the BEA, said the agency is not aware of any such video. Lt. Col. Jean-Marc Menichini, a French Gendarmerie spokesman in charge of communications on rescue efforts around the Germanwings crash site, told CNN that the reports were "completely wrong" and "unwarranted." Cell phones have been collected at the site, he said, but that they "hadn't been exploited yet." Menichini said he believed the cell phones would need to be sent to the Criminal Research Institute in Rosny sous-Bois, near Paris, in order to be analyzed by specialized technicians working hand-in-hand with investigators. But none of the cell phones found so far have been sent to the institute, Menichini said. Asked whether staff involved in the search could have leaked a memory card to the media, Menichini answered with a categorical "no." Reichelt told "Erin Burnett: Outfront" that he had watched the video and stood by the report, saying Bild and Paris Match are "very confident" that the clip is real. He noted that investigators only revealed they'd recovered cell phones from the crash site after Bild and Paris Match published their reports. "That is something we did not know before. ... Overall we can say many things of the investigation weren't revealed by the investigation at the beginning," he said. What was mental state of Germanwings co-pilot? German airline Lufthansa confirmed Tuesday that co-pilot Andreas Lubitz had battled depression years before he took the controls of Germanwings Flight 9525, which he's accused of deliberately crashing last week in the French Alps. Lubitz told his Lufthansa flight training school in 2009 that he had a "previous episode of severe depression," the airline said Tuesday. Email correspondence between Lubitz and the school discovered in an internal investigation, Lufthansa said, included medical documents he submitted in connection with resuming his flight training. The announcement indicates that Lufthansa, the parent company of Germanwings, knew of Lubitz's battle with depression, allowed him to continue training and ultimately put him in the cockpit. Lufthansa, whose CEO Carsten Spohr previously said Lubitz was 100% fit to fly, described its statement Tuesday as a "swift and seamless clarification" and said it was sharing the information and documents – including training and medical records – with public prosecutors. Spohr traveled to the crash site Wednesday, where recovery teams have been working for the past week to recover human remains and plane debris scattered across a steep mountainside. He saw the crisis center set up in Seyne-les-Alpes, laid a wreath in the village of Le Vernet, closer to the crash site, where grieving families have left flowers at a simple stone memorial. Menichini told CNN late Tuesday that no visible human remains were left at the site but recovery teams would keep searching. French President Francois Hollande, speaking Tuesday, said that it should be possible to identify all the victims using DNA analysis by the end of the week, sooner than authorities had previously suggested. In the meantime, the recovery of the victims' personal belongings will start Wednesday, Menichini said. Among those personal belongings could be more cell phones belonging to the 144 passengers and six crew on board. Check out the latest from our correspondents. The details about Lubitz's correspondence with the flight school during his training were among several developments as investigators continued to delve into what caused the crash and Lubitz's possible motive for downing the jet. A Lufthansa spokesperson told CNN on Tuesday that Lubitz had a valid medical certificate, had passed all his examinations and "held all the licenses required." Earlier, a spokesman for the prosecutor's office in Dusseldorf, Christoph Kumpa, said medical records reveal Lubitz suffered from suicidal tendencies at some point before his aviation career and underwent psychotherapy before he got his pilot's license. Kumpa emphasized there's no evidence suggesting Lubitz was suicidal or acting aggressively before the crash. Investigators are looking into whether Lubitz feared his medical condition would cause him to lose his pilot's license, a European government official briefed on the investigation told CNN on Tuesday. While flying was "a big part of his life," the source said, it's only one theory being considered. Another source, a law enforcement official briefed on the investigation, also told CNN that authorities believe the primary motive for Lubitz to bring down the plane was that he feared he would not be allowed to fly because of his medical problems. Lubitz's girlfriend told investigators he had seen an eye doctor and a neuropsychologist, both of whom deemed him unfit to work recently and concluded he had psychological issues, the European government official said. But no matter what details emerge about his previous mental health struggles, there's more to the story, said Brian Russell, a forensic psychologist. "Psychology can explain why somebody would turn rage inward on themselves about the fact that maybe they weren't going to keep doing their job and they're upset about that and so they're suicidal," he said. "But there is no mental illness that explains why somebody then feels entitled to also take that rage and turn it outward on 149 other people who had nothing to do with the person's problems." Germanwings crash compensation: What we know. Who was the captain of Germanwings Flight 9525? CNN's Margot Haddad reported from Marseille and Pamela Brown from Dusseldorf, while Laura Smith-Spark wrote from London. CNN's Frederik Pleitgen, Pamela Boykoff, Antonia Mortensen, Sandrine Amiel and Anna-Maja Rappard contributed to this report.

**Gold**

Match are "very confident" the video clip is real, an editor says. Andreas Lubitz had informed his Lufthansa training school of an episode of severe depression, airline says.

**BART**

Prosecutor: "A person who has such a video needs to immediately give it to the investigators" French Gendarmerie spokesman: Cell phones have been collected at the site, but haven't been exploited.

**MVP**

Prosecutor: "So far no videos were used in the crash investigation" An official with France's accident investigation agency says the agency is not aware of any video.

**MVP+S**

Prosecutor: "So far no videos were used in the crash investigation" Two publications claim to have found a cell phone video of the crash. "It is a very disturbing scene," says editor-in-chief of Bild online.

---

Table 13: The second instance from the CNN/Daily Mail dataset.

**Input**

*Summarize:* The Palestinian Authority officially became the 123rd member of the International Criminal Court on Wednesday, a step that gives the court jurisdiction over alleged crimes in Palestinian territories. The formal accession was marked with a ceremony at The Hague, in the Netherlands, where the court is based. The Palestinians signed the ICC’s founding Rome Statute in January, when they also accepted its jurisdiction over alleged crimes committed “in the occupied Palestinian territory, including East Jerusalem, since June 13, 2014.” Later that month, the ICC opened a preliminary examination into the situation in Palestinian territories, paving the way for possible war crimes investigations against Israelis. As members of the court, Palestinians may be subject to counter-charges as well. Israel and the United States, neither of which is an ICC member, opposed the Palestinians’ efforts to join the body. But Palestinian Foreign Minister Riad al-Malki, speaking at Wednesday’s ceremony, said it was a move toward greater justice. “As Palestine formally becomes a State Party to the Rome Statute today, the world is also a step closer to ending a long era of impunity and injustice,” he said, according to an ICC news release. “Indeed, today brings us closer to our shared goals of justice and peace.” Judge Kuniko Ozaki, a vice president of the ICC, said acceding to the treaty was just the first step for the Palestinians. “As the Rome Statute today enters into force for the State of Palestine, Palestine acquires all the rights as well as responsibilities that come with being a State Party to the Statute. These are substantive commitments, which cannot be taken lightly,” she said. Rights group Human Rights Watch welcomed the development. “Governments seeking to penalize Palestine for joining the ICC should immediately end their pressure, and countries that support universal acceptance of the court’s treaty should speak out to welcome its membership,” said Balkees Jarrah, international justice counsel for the group. “What’s objectionable is the attempts to undermine international justice, not Palestine’s decision to join a treaty to which over 100 countries around the world are members.” In January, when the preliminary ICC examination was opened, Israeli Prime Minister Benjamin Netanyahu described it as an outrage, saying the court was overstepping its boundaries. The United States also said it “strongly” disagreed with the court’s decision. “As we have said repeatedly, we do not believe that Palestine is a state and therefore we do not believe that it is eligible to join the ICC,” the State Department said in a statement. It urged the warring sides to resolve their differences through direct negotiations. “We will continue to oppose actions against Israel at the ICC as counterproductive to the cause of peace,” it said. But the ICC begs to differ with the definition of a state for its purposes and refers to the territories as “Palestine.” While a preliminary examination is not a formal investigation, it allows the court to review evidence and determine whether to investigate suspects on both sides. Prosecutor Fatou Bensouda said her office would “conduct its analysis in full independence and impartiality.” The war between Israel and Hamas militants in Gaza last summer left more than 2,000 people dead. The inquiry will include alleged war crimes committed since June. The International Criminal Court was set up in 2002 to prosecute genocide, crimes against humanity and war crimes. CNN’s Vasco Cotovio, Kareem Khadder and Faith Karimi contributed to this report.

**Gold**

Membership gives the ICC jurisdiction over alleged crimes committed in Palestinian territories since last June. Israel and the United States opposed the move, which could open the door to war crimes investigations against Israelis.

**BART**

Palestinian Authority becomes 123rd member of the International Criminal Court. The move gives the court jurisdiction over alleged crimes in Palestinian territories. Israel and the United States opposed the Palestinians’ efforts to join the body.

**MVP**

“Today brings us closer to our shared goals of justice and peace,” foreign minister says. The Palestinians signed the ICC’s founding Rome Statute in January. The move gives the court jurisdiction over alleged crimes in Palestinian territories.

**MVP+S**

“Today brings us closer to our shared goals of justice and peace,” foreign minister says. The United States says it “strongly” disagrees with the decision. The Palestinian Authority is the 123rd member of the International Criminal Court.

Table 14: The first instance from the WebNLG dataset, which has two gold target sentences.

---

<b>Input</b> <i>Describe the following data:</i> Abilene, Texas — cityServed — Abilene_Regional_Airport
<b>Gold</b> Abilene, Texas is served by the Abilene regional airport. Abilene Regional Airport serves the city of Abilene in Texas.
<b>BART</b> Abilene Regional Airport serves the city of Abilene in Texas.
<b>MVP</b> Abilene Regional Airport serves the city of Abilene, Texas.
<b>MVP+S</b> Abilene Regional Airport serves the city of Abilene, Texas.

---

Table 15: The second instance from the WebNLG dataset, which has three gold target sentences.

---

<b>Input</b> <i>Describe the following data:</i> "Madrid, Paracuellos de Jarama, San Sebastián de los Reyes and Alcobendas" — location — Adolfo_Suárez_Madrid-Barajas_Airport
<b>Gold</b> Adolfo Suárez Madrid-Barajas Airport can be found in Madrid, Paracuellos de Jarama, San Sebastián de los Reyes and Alcobendas. Adolfo Suarez Madrid-Barajas airport is located at Madrid, Paracuellos de Jarama, San Sebastián de los Reyes and Alcobendas. Adolfo Suarez Madrid-Barajas Airport is located in Madrid, Paracuellos de Jarama, San Sebastian de los Reyes and Alcobendas.
<b>BART</b> Adolfo Suárez Madrid-Barajas Airport can be found in Madrid, Paracuellos de Jarama, San Sebastián de los Reyes and Alcobendas.
<b>MVP</b> Adolfo Suárez Madrid-Barajas Airport can be found in Madrid, Paracuellos de Jarama, San Sebastián de los Reyes and Alcobendas.
<b>MVP+S</b> Adolfo Suárez Madrid-Barajas Airport is located in Madrid, Paracuellos de Jarama, San Sebastián de los Reyes and Alcobendas.

---

Table 16: The first instance from the SQuAD dataset.

---

<b>Input</b> <i>Generate the question based on the answer:</i> Saint Bernadette Soubirous [SEP] Architecturally , the school has a Catholic character . Atop the Main Building ’ s gold dome is a golden statue of the Virgin Mary . Immediately in front of the Main Building and facing it , is a copper statue of Christ with arms upraised with the legend ” Venite Ad Me Omnes ” . Next to the Main Building is the Basilica of the Sacred Heart . Immediately behind the basilica is the Grotto , a Marian place of prayer and reflection . It is a replica of the grotto at Lourdes , France where the Virgin Mary reputedly appeared to Saint Bernadette Soubirous in 1858 . At the end of the main drive ( and in a direct line that connects through 3 statues and the Gold Dome ) , is a simple , modern stone statue of Mary .
<b>Gold</b> To whom did the Virgin Mary allegedly appear in 1858 in Lourdes France ?
<b>BART</b> Who is believed to have appeared to the Virgin Mary at Lourdes ?
<b>MVP</b> Who did the Virgin Mary appear to in Lourdes ?
<b>MVP+S</b> The Grotto is a replica of the grotto at Lourdes , France where the Virgin Mary appeared to whom ?

---

Table 17: The second instance from the SQuAD dataset.

---

<b>Input</b> <i>Generate the question based on the answer:</i> a copper statue of Christ [SEP] Architecturally , the school has a Catholic character . Atop the Main Building ’ s gold dome is a golden statue of the Virgin Mary . Immediately in front of the Main Building and facing it , is a copper statue of Christ with arms upraised with the legend ” Venite Ad Me Omnes ” . Next to the Main Building is the Basilica of the Sacred Heart . Immediately behind the basilica is the Grotto , a Marian place of prayer and reflection . It is a replica of the grotto at Lourdes , France where the Virgin Mary reputedly appeared to Saint Bernadette Soubirous in 1858 . At the end of the main drive ( and in a direct line that connects through 3 statues and the Gold Dome ) , is a simple , modern stone statue of Mary .
<b>Gold</b> What is in front of the Notre Dame Main Building ?
<b>BART</b> What is in front of the Main Building and facing it ?
<b>MVP</b> What is immediately in front of the Main Building ?
<b>MVP+S</b> What is immediately in front of the Main Building ?

---

Table 18: The first instance from the CoQA dataset.

<b>Input</b> <i>Answer the following question:</i> what color was cotton ? [X_SEP] once upon a time , in a barn near a farm house , there lived a little white kitten named cotton . cotton lived high up in a nice warm place above the barn where all of the farmer ’ s horses slept . but cotton wasn ’ t alone in her little home above the barn , oh no . she shared her hay bed with her mommy and 5 other sisters . all of her sisters were cute and fluffy , like cotton . but she was the only white one in the bunch . the rest of her sisters were all orange with beautiful white tiger stripes like cotton ’ s mommy . being different made cotton quite sad . she often wished she looked like the rest of her family . so one day , when cotton found a can of the old farmer ’ s orange paint , she used it to paint herself like them . when her mommy and sisters found her they started laughing . ” what are you doing , cotton ? ! ” ” i only wanted to be more like you ” . cotton ’ s mommy rubbed her face on cotton ’ s and said ” oh cotton , but your fur is so pretty and special , like you . we would never want you to be any other way ” . and with that , cotton ’ s mommy picked her up and dropped her into a big bucket of water . when cotton came out she was herself again . her sisters licked her face until cotton ’ s fur was all all dry . ” don ’ t ever do that again , cotton ! ” they all cried . ” next time you might mess up that pretty white fur of yours and we wouldn ’ t want that ! ” then cotton thought , ” i change my mind . i like being special ” .
<b>Gold</b> white
<b>BART</b> white
<b>MVP</b> white
<b>MVP+S</b> white

Table 19: The second instance from the CoQA dataset.

<b>Input</b> <i>Answer the following question:</i> what color was cotton ? [SEP] white [X_SEP] where did she live ? [X_SEP] once upon a time , in a barn near a farm house , there lived a little white kitten named cotton . cotton lived high up in a nice warm place above the barn where all of the farmer ’ s horses slept . but cotton wasn ’ t alone in her little home above the barn , oh no . she shared her hay bed with her mommy and 5 other sisters . all of her sisters were cute and fluffy , like cotton . but she was the only white one in the bunch . the rest of her sisters were all orange with beautiful white tiger stripes like cotton ’ s mommy . being different made cotton quite sad . she often wished she looked like the rest of her family . so one day , when cotton found a can of the old farmer ’ s orange paint , she used it to paint herself like them . when her mommy and sisters found her they started laughing . ” what are you doing , cotton ? ! ” ” i only wanted to be more like you ” . cotton ’ s mommy rubbed her face on cotton ’ s and said ” oh cotton , but your fur is so pretty and special , like you . we would never want you to be any other way ” . and with that , cotton ’ s mommy picked her up and dropped her into a big bucket of water . when cotton came out she was herself again . her sisters licked her face until cotton ’ s fur was all all dry . ” don ’ t ever do that again , cotton ! ” they all cried . ” next time you might mess up that pretty white fur of yours and we wouldn ’ t want that ! ” then cotton thought , ” i change my mind . i like being special ” .
<b>Gold</b> in a barn
<b>BART</b> in a barn
<b>MVP</b> in a barn
<b>MVP+S</b> in a barn

Table 20: The first instance from the ROCStories dataset.

---

**Input**

*Given the story title:* male was out jogging one morning .

---

**Gold**

the weather was crisp and cool . male felt good and energetic . he decided to keep jogging longer than normal . male went several more miles out of his way .

---

**BART**

he saw a man struggling to walk across a field . male stopped to help him . he saw a guy struggling to walking across a fields . the man thanked male profusely !

---

**MVP**

he ran into a man who was holding a knife . male was scared , so he ran away . the man threatened male . male ran back into his house .

---

**MVP+S**

he noticed a sound coming from a tree . male went to investigate . he found a small baby bird stuck in the trunk . male took the bird to the vet and they released it .

---

Table 21: The second instance from the ROCStories dataset.

---

**Input**

*Given the story title:* ken was driving around in the snow .

---

**Gold**

he needed to get home from work . he was driving slowly to avoid accidents . unfortunately the roads were too slick and ken lost control . his tires lost traction and he hit a tree .

---

**BART**

he hit a patch of black ice . ken slid and hurt his knee . he went to the doctor for it . krampus was sent home with a concussion .

---

**MVP**

he was in his car and it was very cold . ken started to back up . the car started to make a loud noise . kema had to back out and call for help .

---

**MVP+S**

he was having trouble starting his car . he had to push it a couple of blocks . his car would n't start . ken had to call a tow truck for help .

---



Table 22: The first instance from the PersonaChat dataset.

---

<b>Input</b> <i>Given the dialog:</i> i love to meet new people . [SEP] i have a turtle named timothy . [SEP] my favorite sport is ultimate frisbee . [SEP] my parents are living in bora bora . [SEP] autumn is my favorite season . [X_SEP] hello , how are you doing tonight ?
<b>Gold</b> i am well an loving this interaction how are you ?
<b>BART</b> i am doing well , how are you ?
<b>MVP</b> i am doing well . how are you ?
<b>MVP+S</b> i am doing well . how are you ?

---

Table 23: The second instance from the PersonaChat dataset.

---

<b>Input</b> <i>Given the dialog:</i> i just bought a brand new house . [SEP] i like to dance at the club . [SEP] i run a dog obedience school . [SEP] i have a big sweet tooth . [SEP] i like taking and posting selkies . [X_SEP] hello , how are you doing tonight ? [SEP] i am well an loving this interaction how are you ?
<b>Gold</b> i am great . i just got back from the club .
<b>BART</b> i am doing well , thanks for asking . what do you do for fun ?
<b>MVP</b> i am doing well . i just got home from dancing at the club .
<b>MVP+S</b> i am doing well . i just bought a new house .

---

Table 24: The first instance from the MultiWOZ dataset. The task-oriented dialogue is composed of dialogue state tracking, dialogue action learning and system response generation.

<p><b>Input</b>  <i>Given the task dialog:</i> Belief state [X_SEP] i would like a taxi from saint john 's college to pizza hut fen ditton .  <i>Given the task dialog:</i> Dialogue action [X_SEP] [db_nores] [X_SEP] i would like a taxi from saint john 's college to pizza hut fen ditton .  <i>Given the task dialog:</i> System response [X_SEP] [db_nores] [X_SEP] i would like a taxi from saint john 's college to pizza hut fen ditton .</p>
<p><b>Gold</b>  [taxi] destination pizza hut fen ditton departure saint john 's college  [taxi] [request] leave arrive  what time do you want to leave and what time do you want to arrive by ?</p>
<p><b>BART</b>  [taxi] destination pizza express fen ditton departure saint johns college  [taxi] [request] leave  what time would you like to leave or arrive ?</p>
<p><b>MVP</b>  [taxi] destination pizza hut fen ditton departure saint john 's college  [taxi] [request] leave  sure ! when would you like to arrive ?</p>
<p><b>MVP+S</b>  [taxi] destination pizza hut fen ditton departure saint john 's college  [taxi] [request] leave arrive  what time would you like to leave ?</p>

Table 25: The second instance from the MultiWOZ dataset.

<p><b>Input</b>  <i>Given the task dialog:</i> Belief state [X_SEP] i would like a taxi from saint john 's college to pizza hut fen ditton . [SEP] what time do you want to leave and what time do you want to arrive by ? [SEP] i want to leave after 17:15 .  <i>Given the task dialog:</i> Dialogue action [X_SEP] [db_nores] [X_SEP] i would like a taxi from saint john 's college to pizza hut fen ditton . [SEP] what time do you want to leave and what time do you want to arrive by ? [SEP] i want to leave after 17:15 .  <i>Given the task dialog:</i> System response [X_SEP] [db_nores] [X_SEP] i would like a taxi from saint john 's college to pizza hut fen ditton . [SEP] what time do you want to leave and what time do you want to arrive by ? [SEP] i want to leave after 17:15 .</p>
<p><b>Gold</b>  [taxi] destination pizza hut fen ditton departure saint john 's college leave 17:15  [taxi] [inform] car phone  booking completed ! your taxi will be [value_car] contact number is [value_phone]</p>
<p><b>BART</b>  [taxi] destination pizza hut fen ditton departure saint john 's college leave 17:15  [taxi] [inform] car phone  booking completed ! booked car type : [value_car] contact number : [value_phone]</p>
<p><b>MVP</b>  [taxi] destination pizza hut fen ditton departure saint john 's college leave 17:15  [taxi] [inform] car phone  booking completed ! booked car type : [value_car] contact number : [value_phone]</p>
<p><b>MVP+S</b>  [taxi] destination pizza hut fen ditton departure saint john 's college leave 17:15  [taxi] [inform] car phone  booking completed ! booked car type : [value_car] contact number : [value_phone]</p>