

Crosslingual Generalization through Multitask Finetuning

Niklas Muennighoff¹ Thomas Wang¹ Lintang Sutawika^{2,3} Adam Roberts⁴
Stella Biderman^{3,5} Teven Le Scao¹ M Saiful Bari⁶ Sheng Shen⁷
Zheng-Xin Yong⁸ Hailey Schoelkopf^{3,9} Xiangru Tang⁹ Dragomir Radev⁹
Alham Fikri Aji¹⁰ Khalid Almubarak¹¹ Samuel Albanie¹² Zaid Alyafeai¹³

Albert Webson⁸ Edward Raff⁵ Colin Raffel¹

¹Hugging Face ²Datasaur.ai ³EleutherAI ⁴Google Research, Brain Team

⁵Booz Allen Hamilton ⁶NTU ⁷UC Berkeley ⁸Brown University

⁹Yale University ¹⁰MBZUAI ¹¹PSAU ¹²University of Cambridge ¹³KFUPM

niklas@hf.co

Abstract

Multitask prompted finetuning (MTF) has been shown to help large language models generalize to new tasks in a zero-shot setting, but so far explorations of MTF have focused on English data and models. We apply MTF to the pretrained multilingual BLOOM and mT5 model families to produce finetuned variants called BLOOMZ and mT0. We find finetuning large multilingual language models on English tasks with English prompts allows for task generalization to non-English languages that appear only in the pretraining corpus. Finetuning on multilingual tasks with English prompts further improves performance on English and non-English tasks leading to various state-of-the-art zero-shot results. We also investigate finetuning on multilingual tasks with prompts that have been machine-translated from English to match the language of each dataset. We find training on these machine-translated prompts leads to better performance on human-written prompts in the respective languages. Surprisingly, we find models are capable of zero-shot generalization to tasks in languages they have never intentionally seen. We conjecture that the models are learning higher-level capabilities that are both task- and language-agnostic. In addition, we introduce xP3, a composite of supervised datasets in 46 languages with English and machine-translated prompts. Our code, datasets and models are freely available at <https://github.com/bigscience-workshop/xmtf>.

1 Introduction

Large language models pretrained on vast amounts of text show some capability of solving tasks expressed in natural language, even without explicit training on these tasks (Brown et al., 2020). Finetuning on groups of language tasks has been shown to significantly boost this zero-shot task generalization of language models (Wei et al., 2021; Sanh

et al., 2022; Min et al., 2021). For example, Sanh et al. (2022) finetune on tasks like summarization and question answering leading to better performance on unseen tasks like natural language inference. Previous work has focused on multitask finetuning in the context of large English language models and tasks.

Multilingual large language models show the same zero-shot learning capabilities for both monolingual and crosslingual tasks (Goyal et al., 2021a; Lin et al., 2021; Patel et al., 2022; Soltan et al., 2022). However, zero-shot performance tends to be significantly lower than finetuned performance. Thus, task-specific or language-specific transfer learning via finetuning remains the predominant practice (Devlin et al., 2018; Conneau et al., 2019; Aribandi et al., 2021). This is particularly challenging for low-resource languages or tasks with limited data available, such as writing a fable that teaches a specified moral. In the spirit of multitask finetuning, it would be desirable to improve the zero-shot task generalization of multilingual models to make them usable on tasks from low-resource languages without requiring further finetuning.

To address this goal, we focus on crosslingual multitask finetuning. Due to the difficulty of collecting supervised task data in low-resource languages, previous work typically aims to transfer capabilities learned from finetuning on English data, which can improve performance on non-English language tasks (Wu and Dredze, 2019; Phang et al., 2020; Chalkidis et al., 2021; Vu et al., 2022). We investigate whether English-only multitask finetuning also improves performance on non-English held-out tasks using the multilingual BLOOM (Scao et al., 2022a) and mT5 (Xue et al., 2020) models. We find that after finetuning on the English-only multitask mixture used for T0 (Sanh et al., 2022) (P3), performance on a diverse set of

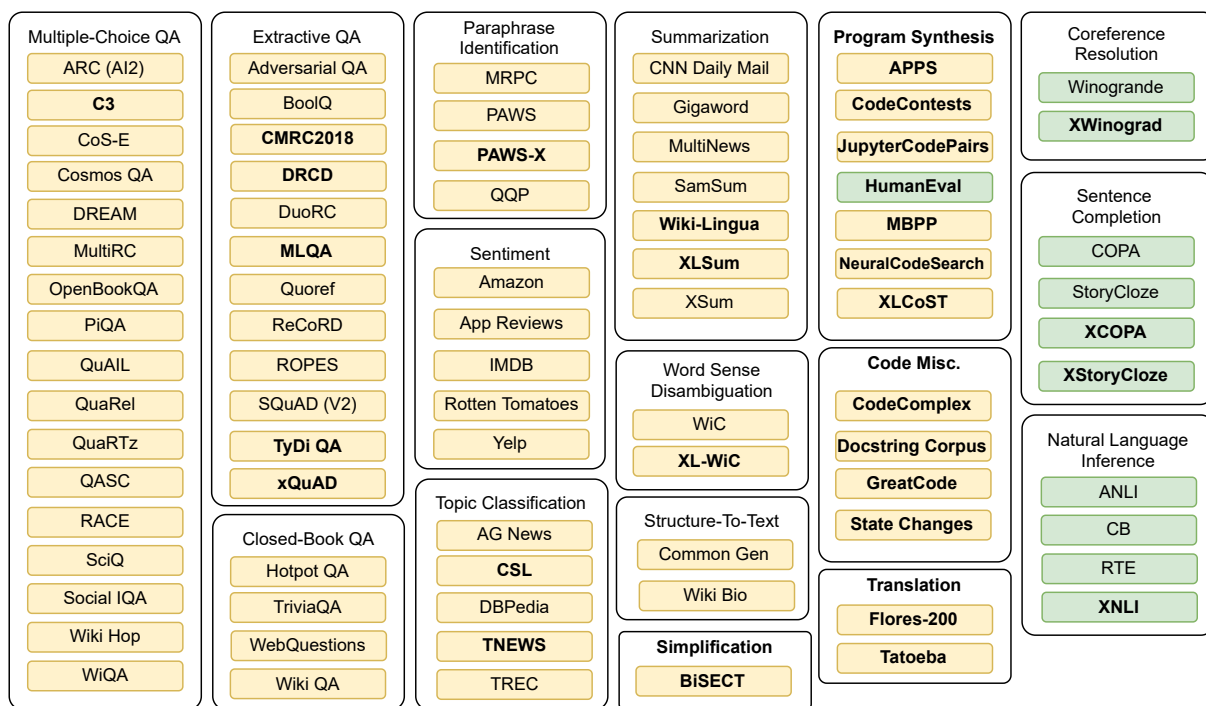


Figure 1: An overview of datasets in xP3. Datasets added to P3 in this work are marked **bold**. Yellow datasets are trained on. Green datasets are held out for evaluation.

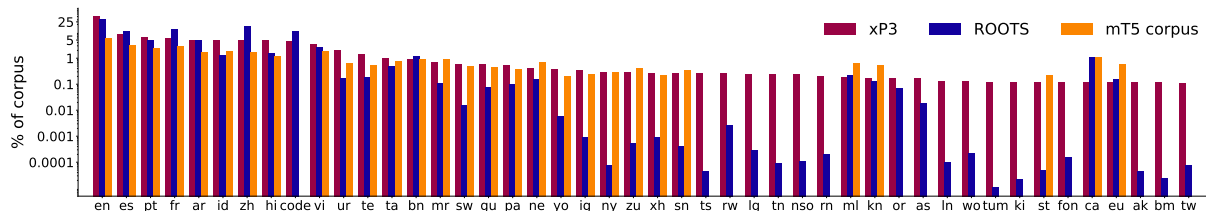


Figure 2: Language composition of xP3, ROOTS, and the corpus of mT5. All ROOTS and xP3 languages are depicted. The mT5 corpus covers additional languages that are not included in the graph.

non-English held-out tasks increases.

To investigate whether multilingual task data can further improve performance, we extend P3 to xP3 by adding datasets from 46 different languages that cover tasks previously not present in P3 (such as translation and program synthesis). Finetuning on xP3 leads to even better zero-shot task generalization in both English and non-English compared to the P3-trained baseline. Models finetuned on xP3 perform best on English prompts, even for non-English samples. Hypothesizing that better performance could be attained by training on non-English prompts, we construct a variant of xP3 with machine-translated prompts called xP3mt. We find that finetuning on machine-translated prompts is enough to significantly increase performance on held-out tasks with non-English human-written prompts. However, reducing the number of En-

glish prompts in the finetuning also worsens English prompt performance on multilingual tasks.

Notably, we also find that models finetuned on xP3 generalize to held-out tasks in languages never intentionally seen during pretraining nor finetuning. We conduct a contamination analysis and find that only small amounts of these languages were included in the pretraining corpus. Thus, we hypothesize the models learn some language- and task-agnostic capabilities.

We publicly release all our datasets and models (URLs in Appendix §C).

2 Related work

2.1 Multitask learning

Multitask finetuning (Sanh et al., 2022) (or instruction tuning (Wei et al., 2021)) has emerged as a

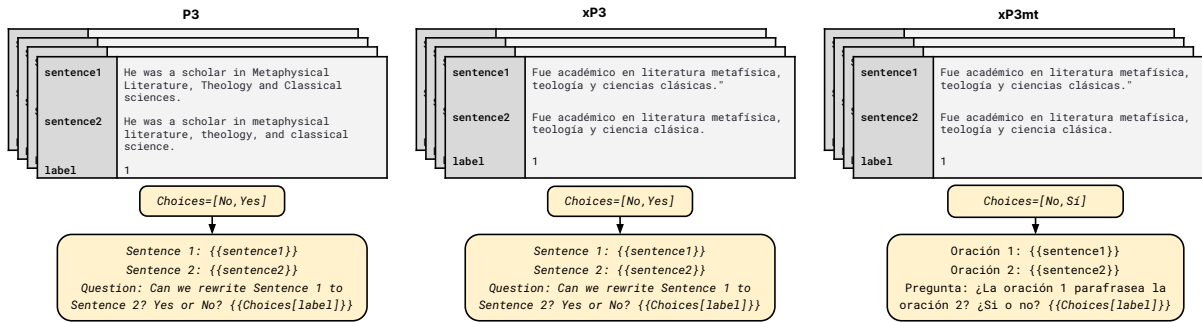


Figure 3: Comparison of dataset variants P3, xP3, and xP3mt on a sample from PAWS for P3 (Zhang et al., 2019) and PAWS-X (Yang et al., 2019) for xP3 and xP3mt. P3 pairs English datasets with English prompts, xP3 pairs multilingual datasets with English prompts and xP3mt pairs multilingual datasets with prompts machine-translated from English to match the dataset language. Expressions in curly brackets are replaced, e.g. for xP3mt the target shown as `{{Choices[label]}}` becomes `Sí`.

recipe for improving the zero-shot task generalization of large language models. Typically, these works define a task as a collection of datasets that require a certain set of skills. To inform large language models which task to perform given an input, a prompt is used to add natural language instructions to dataset instances (Schick and Schütze, 2020; Scao and Rush, 2021). In this line of work, zero-shot task generalization refers to the ability to perform a held-out task based on prompted instructions alone. Our work builds on T0 (Sanh et al., 2022), a variant of T5 (Raffel et al., 2020) that underwent MTF and was subsequently shown to have strong zero-shot task generalization capabilities.

Increasing the number and diversity of finetuning tasks and datasets has been shown to increase model performance (Min et al., 2021; Fries et al., 2022; Wang et al., 2022d; Scialom et al., 2022; Chung et al., 2022; Mishra et al., 2021b). PromptSource (Bach et al., 2022) is a software application that provides a framework for developing and applying prompts. PromptSource was used to construct P3, the training dataset of T0. While most prior work has focused on using English prompts on English datasets, Wang et al. (2022c) trained both English and multilingual models on prompted datasets. Their multilingual model, called mTk-Instruct, attains strong crosslingual performance. In contrast with Wang et al. (2022c), our sole focus is crosslingual zero-shot generalization. Therefore, we consider a wider variety of prompting settings and perform a more detailed evaluation of multilingual capabilities. Separately, Radford et al. (2019) find that accidental inclusion of non-English text gave the GPT-2 model a limited ability to process and generate non-English text. We similarly dis-

cover that our finetuned models can process text in languages not intentionally trained on.

2.2 Multilingual models

Many language models are pretrained on English data only. Multilingual pretrained language models (Lample and Conneau, 2019; Conneau et al., 2019; Fan et al., 2021) aim to enable processing a wide variety of non-English languages. Unlike monolingual models, multilingual models can also be used for crosslingual tasks, such as translation. For language generation, recent efforts have focused on two different model architectures based on the Transformer (Vaswani et al., 2017). On the one hand, encoder-decoder transformers trained with a denoising objective such as mBART (Liu et al., 2020) and mT5 (Xue et al., 2020) learn to predict tokens masked out in the input sequence. Predicting masked tokens is only a pretraining task and these models are generally finetuned on downstream datasets before being used. On the other hand, decoder-only models pretrained on next token prediction such as mGPT (Shliazhko et al., 2022), XGLM (Lin et al., 2021) and BLOOM (Scao et al., 2022a) can be used to solve tasks expressed in natural language directly in a zero-shot or few-shot setting (Brown et al., 2020). XGLM demonstrated competitive few-shot performance even when the model was prompted in a language different than the sample being processed. In particular, using English prompts for multilingual datasets provides better performance with XGLM than human-translating the English prompt to the dataset language.

In this work, we use the BLOOM models (Scao et al., 2022a,b), which were pretrained on the

ROOTS corpus (Laurençon et al., 2022) in 46 natural languages and 13 programming languages. We also finetune mT5 (Xue et al., 2020) to compare encoder-decoder and decoder-only performance. mT5 is pretrained on a corpus sampled from mC4 covering 101 languages.

3 Finetuning data and models

To study crosslingual multitask prompted finetuning, we create xP3 by extending the P3 dataset collection with additional non-English tasks. We finetune both BLOOM and mT5 models on xP3. We refer to Appendix §C for public links to released models and datasets.

3.1 Finetuning data

We build on the P3 (Sanh et al., 2022) task taxonomy and add 30 new multilingual datasets illustrated in Figure 1. We define four task clusters previously not present in P3: translation, simplification, program synthesis, and miscellaneous code datasets. As 11% of BLOOM’s pretraining data is code, we add code datasets classified as program synthesis (text-to-code) or miscellaneous. The latter includes tasks such as estimating the computational complexity of a provided code snippet and generating a name for a given function. We extend the XWinograd dataset (Tikhonov and Ryabinin, 2021) with winograd schemas from CLUE (Xu et al., 2020) to increase its Chinese samples from 16 to 504. Similar to P3, a fraction of our prompts invert the task at hand. For example, a prompt may invert a closed-book QA sample by asking the model to generate a question given an answer.

With xP3 we aim to replicate the language distribution of the ROOTS corpus (Laurençon et al., 2022) used to pretrain BLOOM. Thus, xP3 consists of the same 46 natural languages and code as ROOTS. ROOTS, xP3 and the mT5 corpus (Xue et al., 2020) language distributions are visualized in Figure 2. 39% of xP3 data is English, slightly more than the 30% of English data in ROOTS. Various African languages such as Twi (tw) and Bambara (bm) form the tail of xP3’s language distribution. Many of them are not included in the mT5 pretraining corpus. In xP3, Twi and others are represented solely as a translation task using data from Flores-200 (NLLB Team et al., 2022).

To study the importance of non-English prompts, we construct a machine-translated variant of xP3, xP3mt. We translate prompts of monolingual

datasets into the respective dataset language. For example, for the Chinese dataset C3 (Sun et al., 2020) prompts in xP3mt are in Chinese instead of English in xP3. For crosslingual datasets prompts remain in English in xP3mt (such as Wiki-Lingua, which involves producing a summary in one language based on text in another language). We use the Google Cloud API for machine translation¹. Figure 3 compares the dataset variants we train on.

3.2 Models

We use publicly available pretrained BLOOM models ranging from 560 million to 176 billion parameters. BLOOM models are large decoder-only language models pretrained for around 350 billion tokens with an architecture similar to GPT-3 (Brown et al., 2020). We finetune the models for an additional 13 billion tokens with loss only being computed on target tokens. For example, given the input “Translate to English: Je t’aime.” and a space-separated target “I love you.”, the model is trained to predict only the targets. As targets vary in length from just one to hundreds of tokens, we downscale the loss of each token by the length of the target it belongs to. This ensures short targets (e.g. for multiple-choice QA) get the same weight as long targets (e.g. for translation). We skip samples longer than 2048 tokens and use packing to train efficiently on multiple samples at a time (Kosec et al., 2021). We select the final checkpoint based on validation performance.

For mT5 models, we finetune using the T5X (Roberts et al., 2022) framework on TPUs. mT5 uses the same encoder-decoder architecture, pretraining objective (masked language modeling), and pretraining length (1 trillion tokens) as T5 (Raffel et al., 2020). For finetuning mT5, we follow the same procedure as described above for BLOOM, except that inputs are fed into the encoder and thus are not space-separated from targets.

We produce three core model variants available in different sizes:

- **BLOOMZ-P3 / mT0-P3:** Models finetuned on the English-only P3.
- **BLOOMZ / mT0:** Models finetuned on xP3, which consists of multilingual datasets with English prompts.
- **BLOOMZ-MT / mT0-MT:** Models finetuned on xP3mt, which consists of multi-

¹<https://cloud.google.com/translate>

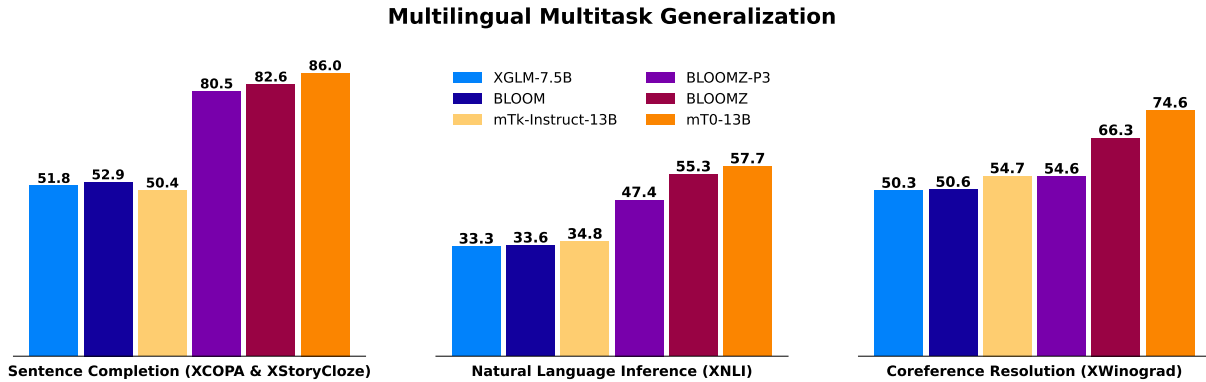


Figure 4: Zero-shot multilingual task generalization with English prompts. BLOOM models have 176 billion parameters. Scores are the language average for each task. Appendix §B breaks down performance by language.

lingual datasets with English and machine-translated prompts.

We evaluate on three held-out tasks: coreference resolution, sentence completion and natural language inference (NLI) as depicted in Figure 1. We also evaluate on HumanEval due to its popularity for code evaluations using the pass@k metric (Chen et al., 2021). For datasets that involve choosing the correct completion from several options, we follow prior work (Sanh et al., 2022; Brown et al., 2020) and use rank classification: We compute the log-likelihood of each possible completion and select the highest scoring option. For each evaluation dataset, we select 5 prompts at random from PromptSource and use them for all language splits of the dataset. We report the median of the 5 prompts for results per language split. Thus, in contrast to XGLM (Lin et al., 2021), we do not tune prompts based on performance on validation data. A selection of prompts can be found in Appendix §M. For evaluation on generative tasks, such as translation, we use lm-evaluation-harness (Gao et al., 2021) and report BLEU scores (Papineni et al., 2002).

4 Results

We first examine generalization to new tasks in languages included in finetuning in §4.1. Then, in §4.2, we look at language generalization: Can models generalize to tasks in languages that (a) they have only seen during pretraining and (b) they have never seen intentionally? In §4.3, we investigate performance on multilingual prompts and finetuning on xP3mt. Scaling laws are analyzed in §4.4. Finally, §4.5 looks at performance on generative tasks and §4.6 at the effect of language

proportions on performance.

4.1 Task generalization

Previous work has shown that large language models finetuned on prompted multitask mixtures generalize to unseen tasks (Zhong et al., 2021; Wei et al., 2021; Mishra et al., 2021b,a; Wang et al., 2022c). In Figure 4, we show that the same applies to multilingual models: Finetuned BLOOMZ and BLOOMZ-P3 models significantly improve over BLOOM and XGLM on held-out tasks. Despite an order of magnitude fewer parameters, mT0 (13 billion parameters) is ahead of BLOOMZ (176 billion parameters). We attribute this to the encoder-decoder architecture paired with a masked language modeling pretraining objective (Wang et al., 2022a; Tay et al., 2022a) as well as the longer pretraining of mT5 (Hoffmann et al., 2022; Su et al., 2022) (1 trillion tokens for mT5 vs. 366 billion for BLOOM). Despite also having gone through crosslingual multitask finetuning, mTk-Instruct performs significantly worse than the same-sized mT0. We attribute this to our prompting style, which aims to replicate natural human communication. mTk-Instruct is finetuned on more structured prompts with specific “Definition”, “Input” and “Output” fields. Similarly, Wang et al. (2022c) find that T0 performs worse than Tk-Instruct on their prompts. We also find models finetuned on the 39% English xP3 (BLOOMZ, mT0-13B) outperform models finetuned on the 100% English P3 (BLOOMZ-P3, mT0-13B-P3) on *English tasks* (Appendix §B). Even the fully English T0-11B model (Sanh et al., 2022) is outperformed by our mT0-13B model on entirely *English tasks*. Ignoring embedding parameters T0-11B and mT0-13B have about the same size. This is likely due to xP3 adding additional

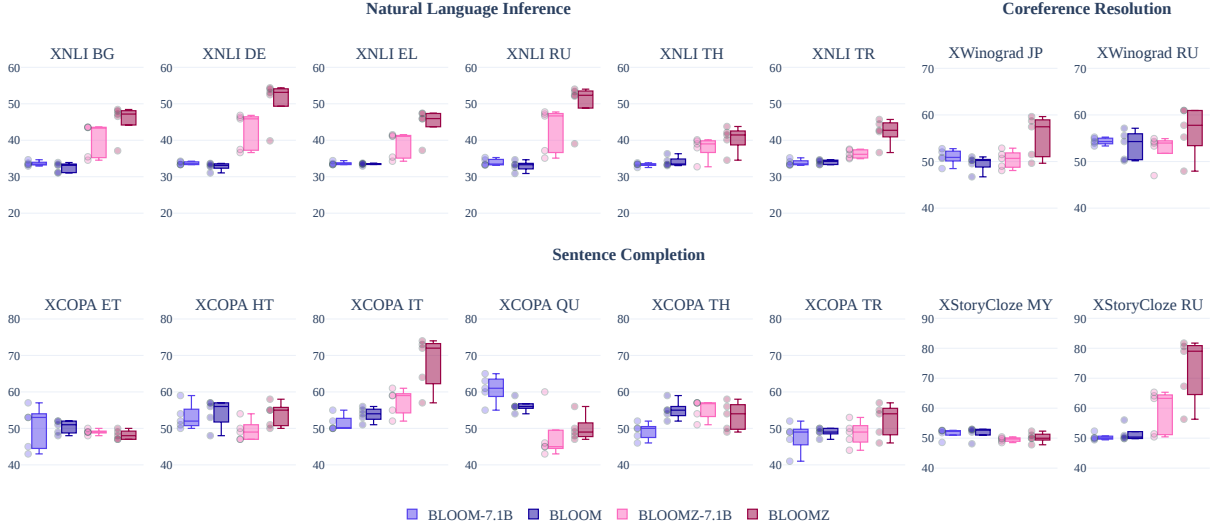


Figure 5: Zero-shot task and language generalization using English prompts on tasks and languages not intentionally seen during pretraining nor finetuning. Language codes are ISO 639-1, except for JP (Japanese).

tasks and prompts, which has been shown to help generalization (Chung et al., 2022; Iyer et al., 2022). mT0-13B beating T0-11B indicates that the benefit of scaling tasks is larger than the benefit of pre-training and finetuning on relatively more English tokens.

4.2 Language generalization

Here we add another layer of generalization: languages. Figure 4 already shows that finetuning on English data only (P3) leads to better performance on non-English data: For example, BLOOMZ-P3 improves by over 50% on multilingual sentence completion compared to BLOOM. Thus, zero-shot task performance in languages only seen during pretraining improves after finetuning on English. This has major practical benefits as it can be more difficult to collect data for low-resource languages.

Next, we investigate performance on languages the model has *never intentionally seen*. Due to the scale of large language model pretraining, it is difficult to label tasks or languages as strictly unseen. It is likely that the training data unintentionally includes small fractions of these languages (just as many tasks might appear “implicitly” in the pretraining corpus (Sanh et al., 2022)). In Figure 5 we show that after multitask finetuning on xP3, the models can perform unseen tasks in languages that were not intentionally trained on. After probing the pretraining corpus of BLOOM, we do find small amounts of these languages that were unintentionally included (Appendix §D). However, for

XNLI, performance increases across all languages, many of which only show up in tiny fractions in our language contamination analysis, such as Thai with 0.006%. If we extrapolate this proportion to the entire ROOTS corpus, the BLOOM models would have seen a mere 20 million tokens of Thai during pretraining. One possibility is that better-than-random XNLI performance can be attained with little or no language understanding. In Appendix §H, we investigate edit distances of XNLI samples and find that there are differences across labels, however, likely not significant enough to enable this kind of generalization.

4.3 Multilingual prompting

Task	Prompt	Average accuracy			
		BLOOMZ	BLOOMZ-MT	mT0-13B	mT0-13B-MT
XNLI	EN	52.99	49.01	48.24	51.29
	MT	37.56	41.16	39.31	41.66
	HT	40.4	43.88	44.95	46.87
XCOPA	EN	72.52	73.24	81.4	80.36
	MT	70.04	71.84	81.16	79.64
XStoryCloze	EN	81.73	81.39	81.99	82.3
	MT	80.89	81.76	83.37	82.86
XWinograd	EN	60.07	59.15	70.49	73.24
	MT	58.48	60.14	66.89	72.33

Table 1: Comparison between EN (English), MT (machine-translated) and HT (human-translated) prompts for 176B BLOOMZ and 13B mT0 models finetuned on either only English or English and machine-translated multilingual prompts (-MT).

Since all prompts in xP3 are in English (even for

multilingual datasets), we created xP3mt, an extension with machine-translated prompts. To investigate performance on non-English prompts, we additionally human- and machine-translated the English evaluation prompts from Figure 4. In Table 1, we report performance on these. Results on machine-translated prompts in languages that are not part of the finetuning corpus, such as those in Figure 5, are in Appendix §I. Table 1 shows that BLOOMZ performs much better on English than on non-English prompts. BLOOMZ-MT, which is finetuned on xP3mt, significantly improves on multilingual prompts. On XNLI, BLOOMZ-MT raises the average performance on human-translated prompts from 41.13 to 45.55. This comes at the cost of a reduction in its performance on English prompts, from 53.58 to 49.74. For mT0, the MT version provides similar performance gains on XNLI and XWinograd non-English prompts, while results on XCOPA and XStoryCloze are mixed. Similar to Lin et al. (2021), we also find that models perform better on human-translated prompts than machine-translated ones for XNLI.

4.4 Scaling

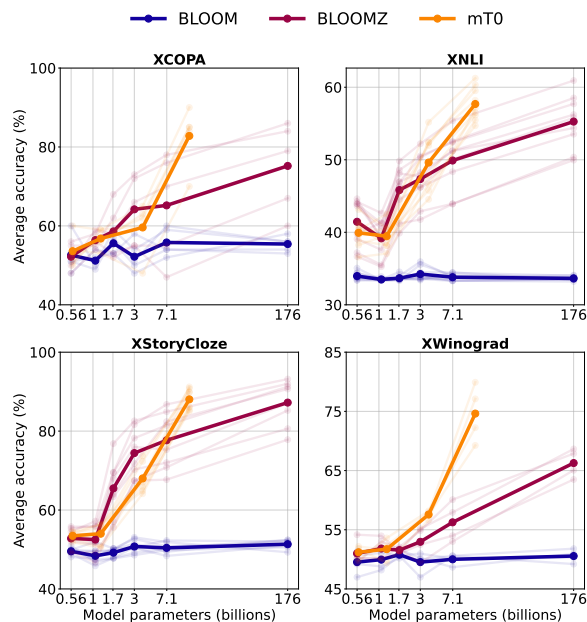


Figure 6: Aggregate performance vs. size. Transparent lines correspond to individual languages, while thick lines are average accuracy scores.

In Figure 4, the average performance of BLOOM is near the random baselines of 0.50 for Sentence Completion and Coreference Resolution and

0.33 for NLI. We think this is due to all of our experiments being zero-shot and using untuned prompts (Perez et al., 2021a). We find in Figure 6 that even at 560M parameters, multitask finetuning improves zero-shot generalization. The gap between pretrained and multitask finetuned models grows significantly as parameters increase. Scaling up parameters benefits all languages evaluated.

4.5 Generation tasks

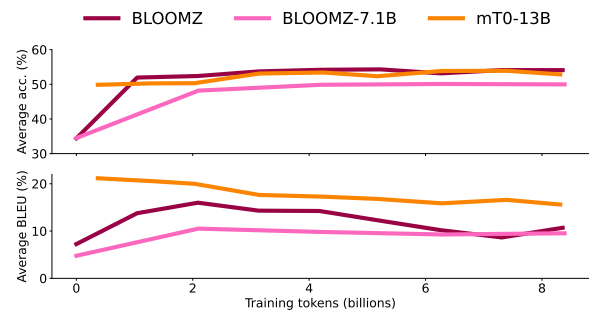


Figure 7: Validation performance during training on natural language understanding (NLU) and natural language generation (NLG) tasks. The former are scored using accuracy and the latter using BLEU (Papineni et al., 2002). The NLG tasks measured are translation and summarization. For BLOOMZ(-7.1B) the performance at 0 training tokens corresponds to the performance of BLOOM(-7.1B). For mT0 there is no data point at 0 tokens, as its base model, mT5, is not suitable for evaluation without finetuning. Performance on individual tasks is in Appendix §K.

In this section, we investigate the impact of multitask finetuning on generative tasks. In Figure 7, we plot validation performance throughout the training process. We find that while performance on natural language understanding tasks continues to increase, generative performance jumps initially and then decreases. Relatedly, in Table 2, we find that multitask finetuning does not improve performance on HumanEval (Chen et al., 2021). Only for small models, such as BLOOM-560M vs. BLOOMZ-560M, there are meaningful performance gains. When no code data is included in finetuning (BLOOMZ-P3) performance decreases significantly. mT0 models, which have not been pretrained on code, fail to solve any HumanEval problems (see full results in Appendix §K). Given a Python docstring, HumanEval requires models to complete a function. Inspecting generations reveals that the multitask finetuned models are biased towards short generations. In Appendix §E, we show example solutions from HumanEval and compute

average length statistics. BLOOMZ tries to solve problems with 70% fewer characters than BLOOM.

	Pass@ <i>k</i>		
	<i>k</i> = 1	<i>k</i> = 10	<i>k</i> = 100
GPT-Neo 1.3B	4.79%	7.47%	16.30%
GPT-Neo 2.7B	6.41%	11.27%	21.37%
GPT-J 6B	11.62%	15.74%	27.74%
GPT-NeoX 20B	15.4%	25.6%	41.2%
Codex-300M	13.17%	20.37%	36.27%
Codex-679M	16.22%	25.7%	40.95%
Codex-2.5B	21.36%	35.42%	59.5%
Codex-12B	28.81%	46.81%	72.31%
BLOOM-560M	0.82%	3.02%	5.91%
BLOOM-1.1B	2.48%	5.93%	9.62%
BLOOM-1.7B	4.03%	7.45%	12.75%
BLOOM-3B	6.48%	11.35%	20.43%
BLOOM-7.1B	7.73%	17.38%	29.47%
BLOOM	15.52%	32.20%	55.45%
BLOOMZ-560M	2.18 %	4.11%	9.00%
BLOOMZ-1.1B	2.63%	6.22%	11.68%
BLOOMZ-1.7B	4.38%	8.73%	16.09%
BLOOMZ-3B	6.29%	11.94%	19.06%
BLOOMZ-7.1B	8.06%	15.03%	27.49%
BLOOMZ	12.06%	26.53%	48.44%
BLOOMZ-P3	6.13%	11.79%	18.73%

Table 2: Code continuation on HumanEval. Non-BLOOM results come from prior work (Chen et al., 2021; Fried et al., 2022). Codex is a language model finetuned on code, while the GPT models (Black et al., 2021; Wang and Komatsuzaki, 2021; Black et al., 2022) are trained on a mix of code and text like BLOOM. Following Chen et al. (2021) we generate 200 samples for each problem with top $p = 0.95$ and compute pass rates. We perform this evaluation three times for temperatures 0.2, 0.6 and 0.8 and pick the best pass rate.

This bias towards short answers and the performance drop on generative tasks come from finetuning on short texts. Most tasks in our finetuning dataset, xP3, are single sentences. We show in Appendix §G that finetuning on fewer short tasks via early stopping, adding long tasks or upweighting long tasks leads to longer generations and slightly better performance. We find it most effective, however, to force a minimum generation length at inference. This is done by ignoring any probability mass the model assigns to its end-of-sequence token for a desired number of tokens. Only after the generation has reached the desired length, can the model generate the end-of-sequence token, thus finishing the generation. Forcing a minimum generation length improves the BLEU score on a translation task by 9 points, see Appendix §G for quantitative and Figure 15 for qualitative results.

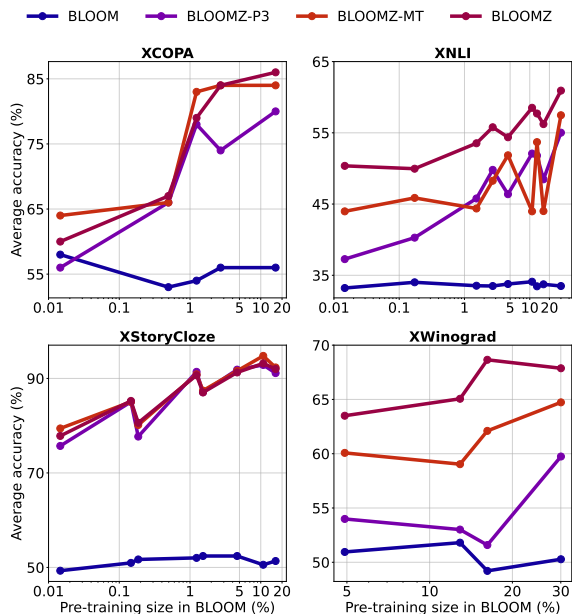


Figure 8: Performance across languages by size in the BLOOM pretraining corpus, ROOTS.

4.6 Effect of language proportions

In Figure 8, we find that finetuned BLOOM models perform better on languages seen extensively during pretraining. As the language distribution in the finetuning dataset, xP3, closely follows that of pretraining, these languages are also seen most frequently during finetuning. Specifically, XCOPA and XNLI show significantly better performance on these high-resource languages, such as English, Spanish or French, which all make up more than 10% of pretraining individually. The trend is less consistent for XWinograd. This may be caused by the fact that XWinograd language subsets are not translations of each other and have a significantly different number of samples. Thus, some language subsets of XWinograd may be inherently more difficult than others.

5 Conclusion

In this work we investigated crosslingual multitask finetuning. We developed xP3, a corpus consisting of tasks in 46 languages. Further, we have extended xP3 to xP3mt with machine-translated prompts. We have finetuned pretrained BLOOM and mT5 models on the newly created corpora as well as the English-only P3 corpus to produce BLOOMZ and mT0 models.

We found that English-only finetuning suffices for a multilingual pretrained large language model to generalize to tasks in other pretrained languages.

However, finetuning on multiple languages using xP3 provided even better performance. We have further observed finetuned models to be capable of generalization to new tasks in languages they have never intentionally seen. We investigated multilingual prompting and found performance after finetuning on English prompts only to be poor. However, finetuning on a corpus with machine-translated prompts (xP3mt) lead to significantly better performance on human-written non-English prompts. Comparing models from 560 million up to 176 billion parameters revealed that the performance gap between only pretraining and finetuning widens as parameters increase. Lastly, we found multitask finetuning on billions of short targets biases models to produce short answers, which can hurt performance on generative tasks. We proposed a simple workaround by forcing a minimum generation length at inference.

To contribute to future progress on improving zero-shot generalization, we release all datasets and models introduced in this work.

6 Limitations

We highlight several limitations of our work:

Unnatural prompting format The choice to separate inputs and targets using a space character has proven effective to multitask finetune our decoder-only models. Nonetheless, poorly formatted prompts may result in undesirable behavior. For example, given the following prompt: “Translate to English: Je t’aime”, the model may continue the input with additional French content before starting to solve the task, i.e. translating the input from French to English. This can be mitigated by improving the prompts with a trailing full stop or a newline symbol. Encoder-decoder models, such as our mT0, do not suffer from this problem, as inputs and targets are fed into different parts of the model.

Limited languages in xP3 The pretraining corpus of mT0 contains more than 101 languages (Xue et al., 2020), however, we finetune on only 46 languages. Likely, finetuning on the full 101 languages mT0 has seen during pretraining would lead to better performance. However, we decided to use only the languages of BLOOM in order to study language generalization (§4.2). Similarly, one could likely attain better performance by enhancing xP3 with more datasets, such as via BIG-Bench (Srivastava et al., 2022; Suzgun et al., 2022), or more

prompts, such as via NL-Augmenter (Dhole et al., 2021). We have released an extended version of xP3 dubbed xP3x that covers 277 languages and is around ten times larger than xP3, but are yet to finetune models on it.

Performance While our models show strong capabilities of performing tasks zero-shot, there remain numerous failure modes that are common in large language models (Rae et al., 2021; Bommasani et al., 2021; Zhang et al., 2022; Smith et al., 2022; Ouyang et al., 2022; Taylor et al., 2022; Chowdhery et al., 2022; Biderman et al., 2023; Allal et al., 2023; Li et al., 2023). In Figure 16 of Appendix §F, BLOOMZ fails to understand the moral of a fable resulting in an undesirable generation. Similarly, in Figure 15, mT0-13B is asked to provide an explanation, but answers with a question. We have made several modifications to the multitask finetuning recipe, such as loss weighting, mixing in long tasks, and various multilingual aspects, leading to the strong zero-shot performance of our models. However, there are many other changes to the multitask finetuning procedure that are worth exploring to get better models (Honovich et al., 2022; Wang et al., 2022b; Longpre et al., 2023a; Liu et al., 2023; Dettmers et al., 2023; Yin et al., 2023). Further, the pre-trained models we use, BLOOM and mT5, are suboptimal in many aspects such as compute allocation (Hoffmann et al., 2022; Muennighoff et al., 2023), pre-training datasets (Longpre et al., 2023b; Touvron et al., 2023; Chung et al., 2023), pre-training objective (Tay et al., 2022b) and possibly model architecture (Komatsuzaki et al., 2022; Shen et al., 2023). Future work should investigate multitask finetuning better base models.

Learning new languages during finetuning While we have investigated generalization to languages only seen during pretraining, we did not investigate generalization to languages only seen during finetuning. Our mT0 models are finetuned on several new languages not seen in pretraining (see Figure 2). Out of those, we only evaluated on code (HumanEval), where mT0 performed at the random baseline (0.00 in Table 10). We point to follow-up work that has investigated the question of teaching BLOOMZ new languages (Yong et al., 2022; Cahyawijaya et al., 2023) and work investigating adaptation of BLOOM (Ennen et al., 2023; Yong and Nikoulina, 2022).

Acknowledgments

This work was granted access to the HPC resources of Institut du développement et des ressources en informatique scientifique (IDRIS) du Centre national de la recherche scientifique (CNRS) under the allocation 2021-A0101012475 made by Grand équipement national de calcul intensif (GENCI). In particular, all the evaluations and data processing ran on the Jean Zay cluster of IDRIS, and we want to thank the IDRIS team for responsive support throughout the project, in particular Rémi Lacroix.

We thank the XGLM team for providing access to XStoryCloze. We thank volunteers who human-translated XNLI prompts. We thank Noah Constant and Douwe Kiela for feedback on drafts of this paper. We thank Victor Sanh, Stephen Bach, Sasha Rush and Jordan Clive for support throughout the project.

References

2018. [Neural code search evaluation dataset](#). page arXiv:1908.09804 [cs.SE].
2020. [Wikilingua: A new benchmark dataset for multilingual abstractive summarization](#). *arXiv preprint arXiv:2010.03093*.
- Loubna Ben Allal, Raymond Li, Denis Kocetkov, Chenghao Mou, Christopher Akiki, Carlos Munoz Ferrandis, Niklas Muennighoff, Mayank Mishra, Alex Gu, Manan Dey, et al. 2023. Santacoder: don't reach for the stars! *arXiv preprint arXiv:2301.03988*.
- Vamsi Aribandi, Yi Tay, Tal Schuster, Jinfeng Rao, Huaixiu Steven Zheng, Sanket Vaibhav Mehta, Honglei Zhuang, Vinh Q. Tran, Dara Bahri, Jianmo Ni, Jai Prakash Gupta, Kai Hui, Sebastian Ruder, and Donald Metzler. 2021. [Ext5: Towards extreme multi-task scaling for transfer learning](#). *CoRR*, abs/2111.10952.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2019. [On the cross-lingual transferability of monolingual representations](#). *CoRR*, abs/1910.11856.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Stephen H. Bach, Victor Sanh, Zheng-Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak, Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, Zaid Alyafeai, Manan Dey, Andrea Santilli, Zhiqing Sun, Srulik Ben-David, Canwen Xu, Gunjan Chhablani, Han Wang, Jason Alan Fries, Maged S. Al-shaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangru Tang, Xiangru Tang, Mike Tian-Jian Jiang, and Alexander M. Rush. 2022. [Promptsources: An integrated development environment and repository for natural language prompts](#).
- Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. 2023. Pythia: A suite for analyzing large language models across training and scaling. *arXiv preprint arXiv:2304.01373*.
- Sid Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, et al. 2022. Gpt-neox-20b: An open-source autoregressive language model. *arXiv preprint arXiv:2204.06745*.
- Sid Black, Leo Gao, Phil Wang, Connor Leahy, and Stella Biderman. 2021. Gpt-neo: Large scale autoregressive language modeling with mesh-tensorflow. *If you use this software, please cite it using these metadata*, 58.
- 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. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv: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. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Samuel Cahyawijaya, Holy Lovenia, Tiezheng Yu, Willy Chung, and Pascale Fung. 2023. Instruct-align: Teaching novel languages with llms through alignment-based cross-lingual instruction. *arXiv preprint arXiv:2305.13627*.
- Ilias Chalkidis, Manos Fergadiotis, and Ion Androutsopoulos. 2021. Multieurlex—a multi-lingual and multi-label legal document classification dataset for zero-shot cross-lingual transfer. *arXiv preprint arXiv:2109.00904*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.

- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- Hyung Won Chung, Noah Constant, Xavier Garcia, Adam Roberts, Yi Tay, Sharan Narang, and Orhan Firat. 2023. Unimax: Fairer and more effective language sampling for large-scale multilingual pretraining. *arXiv preprint arXiv:2304.09151*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *arXiv preprint arXiv:2210.11416*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. Xnli: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Yiming Cui, Ting Liu, Li Xiao, Zhipeng Chen, Wentao Ma, Wanxiang Che, Shijin Wang, and Guoping Hu. 2018. A span-extraction dataset for chinese machine reading comprehension. *arXiv preprint arXiv:1810.07366*.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Kaustubh D Dhole, Varun Gangal, Sebastian Gehrmann, Aadish Gupta, Zhenhao Li, Saad Mahamood, Abinaya Mahendiran, Simon Mille, Ashish Srivastava, Samson Tan, et al. 2021. NI-augmenter: A framework for task-sensitive natural language augmentation. *arXiv preprint arXiv:2112.02721*.
- Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, et al. 2022. Delta tuning: A comprehensive study of parameter efficient methods for pre-trained language models. *arXiv preprint arXiv:2203.06904*.
- Philipp Ennen, Po-Chun Hsu, Chan-Jan Hsu, Chang-Le Liu, Yen-Chen Wu, Yin-Hsiang Liao, Chin-Tung Lin, Da-Shan Shiu, and Wei-Yun Ma. 2023. Extending the pre-training of bloom for improved support of traditional chinese: Models, methods and results. *arXiv preprint arXiv:2303.04715*.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, et al. 2021. Beyond english-centric multilingual machine translation. *J. Mach. Learn. Res.*, 22(107):1–48.
- Daniel Fried, Armen Aghajanyan, Jessy Lin, Sida Wang, Eric Wallace, Freda Shi, Ruiqi Zhong, Wen-tau Yih, Luke Zettlemoyer, and Mike Lewis. 2022. Incoder: A generative model for code infilling and synthesis. *arXiv preprint arXiv:2204.05999*.
- Jason Alan Fries, Leon Weber, Natasha Seelam, Gabriel Altay, Debajyoti Datta, Samuele Garda, Myungsun Kang, Ruisi Su, Wojciech Kusa, Samuel Cahyawijaya, et al. 2022. Bigbio: A framework for data-centric biomedical natural language processing. *arXiv preprint arXiv:2206.15076*.
- Leo Gao, Jonathan Tow, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Kyle McDonell, Niklas Muennighoff, Jason Phang, Laria Reynolds, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2021. [A framework for few-shot language model evaluation](#).
- Naman Goyal, Jingfei Du, Myle Ott, Giri Anantharaman, and Alexis Conneau. 2021a. Larger-scale transformers for multilingual masked language modeling. *arXiv preprint arXiv:2105.00572*.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzm’an, and Angela Fan. 2021b. The flores-101 evaluation benchmark for low-resource and multilingual machine translation.
- Francisco Guzm’an, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. Two new evaluation datasets for low-resource machine translation: Nepali-english and sinhala-english.
- Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, and Rifat Shahriyar. 2021. [XL-sum: Large-scale multilingual abstractive summarization for 44 languages](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4693–4703, Online. Association for Computational Linguistics.

- Vincent J. Hellendoorn, Charles Sutton, Rishabh Singh, Petros Maniatis, and David Bieber. 2020. [Global relational models of source code](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, Dawn Song, and Jacob Steinhardt. 2021. Measuring coding challenge competence with apps. *NeurIPS*.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Or Honovich, Thomas Scialom, Omer Levy, and Timo Schick. 2022. Unnatural instructions: Tuning language models with (almost) no human labor. *arXiv preprint arXiv:2212.09689*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Srinivasan Iyer, Xi Victoria Lin, Ramakanth Pasunuru, Todor Mihaylov, Dániel Simig, Ping Yu, Kurt Shuster, Tianlu Wang, Qing Liu, Punit Singh Koura, et al. 2022. Opt-impl: Scaling language model instruction meta learning through the lens of generalization. *arXiv preprint arXiv:2212.12017*.
- Joongwon Kim, Mounica Maddela, Reno Kriz, Wei Xu, and Chris Callison-Burch. 2021. [BiSECT: Learning to split and rephrase sentences with bitexts](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6193–6209, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Aran Komatsuzaki, Joan Puigcerver, James Lee-Thorp, Carlos Riquelme Ruiz, Basil Mustafa, Joshua Ainslie, Yi Tay, Mostafa Dehghani, and Neil Houlsby. 2022. Sparse upcycling: Training mixture-of-experts from dense checkpoints. *arXiv preprint arXiv:2212.05055*.
- Matej Kosec, Sheng Fu, and Mario Michael Krell. 2021. Packing: Towards 2x nlp bert acceleration. *arXiv preprint arXiv:2107.02027*.
- Guillaume Lample and Alexis Conneau. 2019. [Cross-lingual language model pretraining](#).
- Hugo Launçon, Lucile Saulnier, Thomas Wang, Christopher Akiki, Albert Villanova del Moral, Teven Le Scao, Leandro Von Werra, Chenghao Mou, Eduardo González Ponferrada, Huu Nguyen, et al. 2022. The bigscience roots corpus: A 1.6 tb composite multilingual dataset. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Vladimir I Levenshtein et al. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710. Soviet Union.
- Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2019. Mlqa: Evaluating cross-lingual extractive question answering. *arXiv preprint arXiv:1910.07475*.
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, et al. 2023. Starcoder: may the source be with you! *arXiv preprint arXiv:2305.06161*.
- Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shruti Bhosale, Jingfei Du, et al. 2021. Few-shot learning with multilingual language models. *arXiv preprint arXiv:2112.10668*.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohata, Tenghao Huang, Mohit Bansal, and Colin Raffel. 2022. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *arXiv preprint arXiv:2205.05638*.
- Qian Liu, Fan Zhou, Zhengbao Jiang, Longxu Dou, and Min Lin. 2023. From zero to hero: Examining the power of symbolic tasks in instruction tuning. *arXiv preprint arXiv:2304.07995*.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Robert L Logan, Ivana Balažević, Eric Wallace, Fabio Petroni, Sameer Singh, and Sebastian Riedel. 2021. Cutting down on prompts and parameters: Simple few-shot learning with language models. *arXiv preprint arXiv:2106.13353*.
- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, et al. 2023a. The flan collection: Designing data and methods for effective instruction tuning. *arXiv preprint arXiv:2301.13688*.
- Shayne Longpre, Gregory Yauney, Emily Reif, Katherine Lee, Adam Roberts, Barret Zoph, Denny Zhou, Jason Wei, Kevin Robinson, David Mimno, et al. 2023b. A pretrainer’s guide to training data: Measuring the effects of data age, domain coverage, quality, & toxicity. *arXiv preprint arXiv:2305.13169*.
- Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hananeh Hajishirzi. 2021. Metaicl: Learning to learn in context. *arXiv preprint arXiv:2110.15943*.

- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2021a. Cross-task generalization via natural language crowdsourcing instructions. *arXiv preprint arXiv:2104.08773*.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2021b. [Natural instructions: Benchmarking generalization to new tasks from natural language instructions](#). *CoRR*, abs/2104.08773.
- Niklas Muennighoff. 2022. Sgpt: Gpt sentence embeddings for semantic search. *arXiv preprint arXiv:2202.08904*.
- Niklas Muennighoff, Alexander M. Rush, Boaz Barak, Teven Le Scao, Aleksandra Piktus, Nouamane Tazi, Sampo Pyysalo, Thomas Wolf, and Colin Raffel. 2023. [Scaling data-constrained language models](#).
- Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. 2022. [Mteb: Massive text embedding benchmark](#). *arXiv preprint arXiv:2210.07316*.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#). *arXiv preprint 2207.04672*.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Ajay Patel, Bryan Li, Mohammad Sadegh Rasooli, Noah Constant, Colin Raffel, and Chris Callison-Burch. 2022. Bidirectional language models are also few-shot learners. *arXiv preprint arXiv:2209.14500*.
- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. 2021a. True few-shot learning with language models. *Advances in Neural Information Processing Systems*, 34:11054–11070.
- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. 2021b. [True few-shot learning with language models](#). *CoRR*, abs/2105.11447.
- Jason Phang, Iacer Calixto, Phu Mon Htut, Yada Pruksachatkun, Haokun Liu, Clara Vania, Katharina Kann, and Samuel R Bowman. 2020. English intermediate-task training improves zero-shot cross-lingual transfer too. *arXiv preprint arXiv:2005.13013*.
- Edoardo M. Ponti, Goran Glavas, Olga Majewska, Qianchu Liu, Ivan Vulic, and Anna Korhonen. 2020. [XCOPA: A multilingual dataset for causal commonsense reasoning](#). *arXiv preprint*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. 2021. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Alessandro Raganato, Tommaso Pasini, Jose Camacho-Collados, and Mohammad Taher Pilehvar. 2020. Xliw: A multilingual benchmark for evaluating semantic contextualization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7193–7206.
- Adam Roberts, Hyung Won Chung, Anselm Levskaya, Gaurav Mishra, James Bradbury, Daniel Andor, Sharan Narang, Brian Lester, Colin Gaffney, Afroz Mohiuddin, Curtis Hawthorne, Aitor Lewkowycz, Alex Salcianu, Marc van Zee, Jacob Austin, Sebastian Goodman, Livio Baldini Soares, Haitang Hu, Sasha Tsvyashchenko, Aakanksha Chowdhery, Jasmijn Bastings, Jannis Bulian, Xavier Garcia, Jianmo Ni, Andrew Chen, Kathleen Kenealy, Jonathan H. Clark, Stephan Lee, Dan Garrette, James Lee-Thorp, Colin Raffel, Noam Shazeer, Marvin Ritter, Maarten Bosma, Alexandre Passos, Jeremy Maitin-Shepard, Noah Fiedel, Mark Omernick, Brennan Saeta, Ryan Sepassi, Alexander Spiridonov, Joshua Newlan, and Andrea Gesmundo. 2022. [Scaling up models and data with t5x and seqio](#). *arXiv preprint arXiv:2203.17189*.
- Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S Gordon. 2011. [Choice of plausible alternatives: An evaluation of commonsense causal reasoning](#). In *2011 AAAI Spring Symposium Series*.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. 2022. Multitask prompted training enables zero-shot task generalization. In *The Tenth International Conference on Learning Representations*.

- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2022a. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.
- Teven Le Scao and Alexander M Rush. 2021. How many data points is a prompt worth? *arXiv preprint arXiv:2103.08493*.
- Teven Le Scao, Thomas Wang, Daniel Hesslow, Lucile Saulnier, Stas Bekman, M Saiful Bari, Stella Bideman, Hady Elsahar, Niklas Muennighoff, Jason Phang, et al. 2022b. What language model to train if you have one million gpu hours? *arXiv preprint arXiv:2210.15424*.
- Timo Schick and Hinrich Schütze. 2020. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*.
- Timo Schick and Hinrich Schütze. 2020. [Exploiting cloze questions for few-shot text classification and natural language inference](#). *CoRR*, abs/2001.07676.
- Thomas Scialom, Tuhin Chakrabarty, and Smaranda Muresan. 2022. Continual-t0: Progressively instructing 50+ tasks to language models without forgetting. *arXiv preprint arXiv:2205.12393*.
- Sheng Shen, Le Hou, Yanqi Zhou, Nan Du, Shayne Longpre, Jason Wei, Hyung Won Chung, Barret Zoph, William Fedus, Xinyun Chen, et al. 2023. Flan-moe: Scaling instruction-finetuned language models with sparse mixture of experts. *arXiv preprint arXiv:2305.14705*.
- Oleh Shliakhko, Alena Fenogenova, Maria Tikhonova, Vladislav Mikhailov, Anastasia Kozlova, and Tatiana Shavrina. 2022. mgpt: Few-shot learners go multilingual. *arXiv preprint arXiv:2204.07580*.
- Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhunoye, George Zerveas, Vijay Korthikanti, et al. 2022. Using deep-speed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990*.
- Saleh Soltan, Shankar Ananthkrishnan, Jack FitzGerald, Rahul Gupta, Wael Hamza, Haidar Khan, Charith Peris, Stephen Rawls, Andy Rosenbaum, Anna Rumshisky, Chandana Satya Prakash, Mukund Sridhar, Fabian Trifunovic, Apurv Verma, Gokhan Tur, and Prem Natarajan. 2022. [Alexatm 20b: Few-shot learning using a large-scale multilingual seq2seq model](#).
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*.
- Hui Su, Xiao Zhou, Houjing Yu, Yuwen Chen, Zilin Zhu, Yang Yu, and Jie Zhou. 2022. Welm: A well-read pre-trained language model for chinese. *arXiv preprint arXiv:2209.10372*.
- Kai Sun, Dian Yu, Dong Yu, and Claire Cardie. 2020. [Investigating prior knowledge for challenging chinese machine reading comprehension](#). *Trans. Assoc. Comput. Linguistics*, 8:141–155.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, et al. 2022. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*.
- Yi Tay, Mostafa Dehghani, Vinh Q Tran, Xavier Garcia, Dara Bahri, Tal Schuster, Huaixiu Steven Zheng, Neil Houlsby, and Donald Metzler. 2022a. [Unifying language learning paradigms](#). *arXiv preprint arXiv:2205.05131*.
- Yi Tay, Mostafa Dehghani, Vinh Q Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Dara Bahri, Tal Schuster, Steven Zheng, et al. 2022b. [U12: Unifying language learning paradigms](#). In *The Eleventh International Conference on Learning Representations*.
- Yi Tay, Jason Wei, Hyung Won Chung, Vinh Q Tran, David R So, Siamak Shakeri, Xavier Garcia, Huaixiu Steven Zheng, Jinfeng Rao, Aakanksha Chowdhery, et al. 2022c. [Transcending scaling laws with 0.1% extra compute](#). *arXiv preprint arXiv:2210.11399*.
- Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, Anthony Hartshorn, Elvis Saravia, Andrew Poulton, Viktor Kerkez, and Robert Stojnic. 2022. [Galactica: A large language model for science](#). *arXiv preprint arXiv:2211.09085*.
- Jörg Tiedemann. 2020. [The Tatoeba Translation Challenge – Realistic data sets for low resource and multilingual MT](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1174–1182. Association for Computational Linguistics.
- Alexey Tikhonov and Max Ryabinin. 2021. [It’s all in the heads: Using attention heads as a baseline for cross-lingual transfer in commonsense reasoning](#).
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. [Llama: Open and efficient foundation language models](#). *arXiv preprint arXiv:2302.13971*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all

- you need. *Advances in neural information processing systems*, 30.
- Tu Vu, Aditya Barua, Brian Lester, Daniel Cer, Mohit Iyyer, and Noah Constant. 2022. Overcoming catastrophic forgetting in zero-shot cross-lingual generation. *arXiv preprint arXiv:2205.12647*.
- Ben Wang and Aran Komatsuzaki. 2021. Gpt-j-6b: A 6 billion parameter autoregressive language model.
- Thomas Wang, Adam Roberts, Daniel Hesslow, Teven Le Scao, Hyung Won Chung, Iz Beltagy, Julien Launay, and Colin Raffel. 2022a. What language model architecture and pretraining objective work best for zero-shot generalization? *arXiv preprint arXiv:2204.05832*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2022b. Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560*.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoor-molabashi, Yeganeh Kordi, Amirreza Mirzaei, Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Gary Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krima Doshi, Maitreya Patel, Kuntal Kumar Pal, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Shailaja Keyur Sampat, Savan Doshi, Siddhartha Mishra, Sujan Reddy, Sumanta Patro, Tanay Dixit, Xudong Shen, Chitta Baral, Yejin Choi, Noah A. Smith, Hannaneh Hajishirzi, and Daniel Khashabi. 2022c. Super-naturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks. *arXiv preprint arXiv:2204.07705*.
- Zhenhailong Wang, Xiaoman Pan, Dian Yu, Dong Yu, Jianshu Chen, and Heng Ji. 2022d. Zemi: Learning zero-shot semi-parametric language models from multiple tasks. *arXiv preprint arXiv:2210.00185*.
- Albert Webson and Ellie Pavlick. 2021. [Do prompt-based models really understand the meaning of their prompts?](#)
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [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)*, pages 1112–1122. Association for Computational Linguistics.
- Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT](#). 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)*, pages 833–844, Hong Kong, China. Association for Computational Linguistics.
- Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, et al. 2020. Clue: A chinese language understanding evaluation benchmark. *arXiv preprint arXiv:2004.05986*.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2020. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*.
- Yinfei Yang, Yuan Zhang, Chris Tar, and Jason Baldridge. 2019. PAWS-X: A Cross-lingual Adversarial Dataset for Paraphrase Identification. In *Proc. of EMNLP*.
- Da Yin, Xiao Liu, Fan Yin, Ming Zhong, Hritik Bansal, Jiawei Han, and Kai-Wei Chang. 2023. Dynosaur: A dynamic growth paradigm for instruction-tuning data curation. *arXiv preprint arXiv:2305.14327*.
- Zheng-Xin Yong and Vassilina Nikoulina. 2022. Adapting bigscience multilingual model to unseen languages. *arXiv preprint arXiv:2204.04873*.
- Zheng-Xin Yong, Hailey Schoelkopf, Niklas Muenighoff, Alham Fikri Aji, David Ifeoluwa Adelani, Khalid Almubarak, M Saiful Bari, Lintang Sutawika, Jungo Kasai, Ahmed Baruwu, et al. 2022. Bloom+1: Adding language support to bloom for zero-shot prompting. *arXiv preprint arXiv:2212.09535*.
- Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. 2021. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. Glm-130b: An open bilingual pre-trained model. *arXiv preprint arXiv:2210.02414*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. Paws: Paraphrase adversaries from word scrambling. *arXiv preprint arXiv:1904.01130*.
- Ruiqi Zhong, Kristy Lee, Zheng Zhang, and Dan Klein. 2021. [Meta-tuning language models to answer prompts better](#). *CoRR*, abs/2104.04670.

Ming Zhu, Aneesh Jain, Karthik Suresh, Roshan Ravindran, Sindhu Tipirneni, and Chandan K. Reddy. 2022. [Xlcost: A benchmark dataset for cross-lingual code intelligence.](#)

Contents

1	Introduction	1
2	Related work	2
2.1	Multitask learning	2
2.2	Multilingual models	3
3	Finetuning data and models	4
3.1	Finetuning data	4
3.2	Models	4
4	Results	5
4.1	Task generalization	5
4.2	Language generalization	6
4.3	Multilingual prompting	6
4.4	Scaling	7
4.5	Generation tasks	7
4.6	Effect of language proportions	8
5	Conclusion	8
6	Limitations	9
A	Contributions	18
B	Task generalization breakdown	18
C	Artifacts	20
D	ROOTS language contamination	20
E	Code generations	21
F	Qualitative examples	21
G	Increasing generation length	24
H	XNLI edit distances	24
I	Multilingual prompting in unseen languages	25
J	Ideas that did not work	26
K	Full results	26
L	Version control	28
M	Prompts used	28

A Contributions

This research was conducted under the BigScience project for open research, a year-long initiative targeting the study of large models and datasets. The goal of the project is to research language models in a public environment. The project has hundreds of researchers from more than 50 countries and over 250 institutions. The BigScience project was initiated by Thomas Wolf at Hugging Face, and this collaboration would not have been possible without his effort. In the following, we list contributions made to this work.

Niklas Muennighoff evaluated all models, created xP3 and wrote most of the paper.

Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts and Hailey Schoelkopf wrote the training and evaluation code.

Niklas Muennighoff and Adam Roberts trained the models.

Niklas Muennighoff, Teven Le Scao, Hailey Schoelkopf, Zheng-Xin Yong, Thomas Wang, Khalid Almubarak, Alham Fikri Aji, M Saiful Bari and Zaid Alyafeai contributed prompts or datasets.

Lintang Sutawika, Stella Biderman, Zheng-Xin Yong, Khalid Almubarak, M Saiful Bari and Albert Webson initiated the project.

Sheng Shen conducted the contamination analysis.

Samuel Albanie wrote the prompt appendix.

Thomas Wang and Zheng-Xin Yong converted checkpoints.

Colin Raffel, Thomas Wang, Teven Le Scao, M Saiful Bari, Edward Raff and Dragomir Radev advised the project.

Niklas Muennighoff, Lintang Sutawika, Teven Le Scao, Colin Raffel, Stella Biderman, Alham Fikri Aji, Adam Roberts, Samuel Albanie, Sheng Shen, M Saiful Bari, Albert Webson, Xiangru Tang, Dragomir Radev and Edward Raff contributed to the paper.

B Task generalization breakdown

In Figure 9, we compare performance on English held-out tasks. We find that (a) finetuning on xP3 outperforms P3 (b) multilingual mT0 is better than monolingual T0 on *English tasks*. We think both improvements come from xP3 having more prompts and datasets than P3 (Chung et al., 2022).

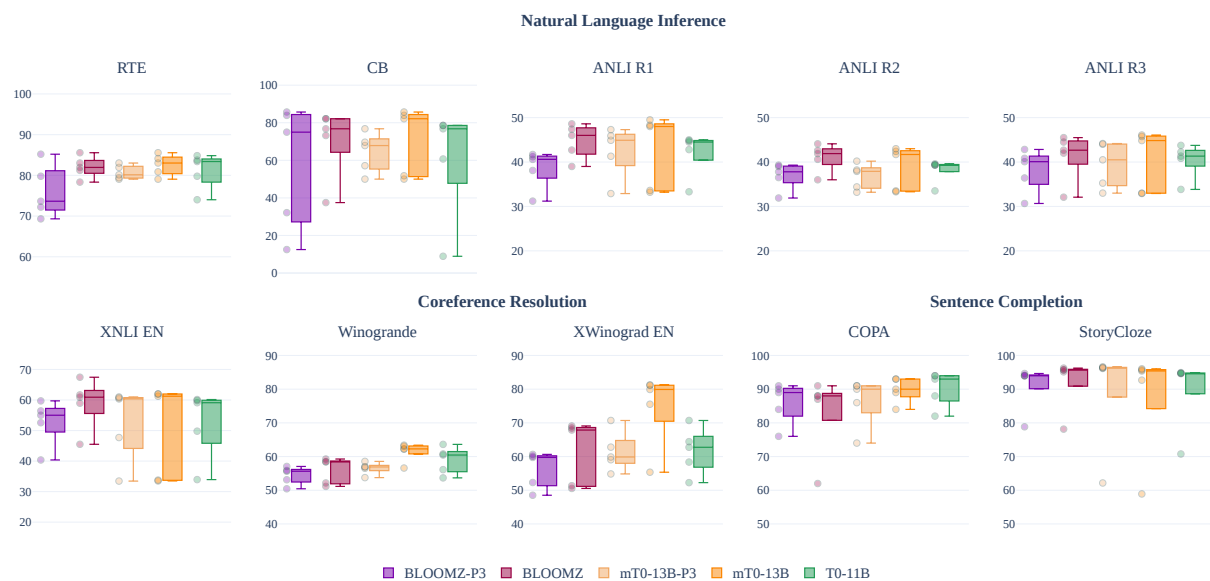


Figure 9: Zero-shot English task generalization. Each dot represents performance on one English evaluation prompt.

In Figure 10, we visualize task generalization to multilingual datasets. The same data is aggregated in Figure 4. Performance by prompt varies substantially highlighting that prompt engineering may still be necessary after MTF. We also find that mT0 consistently outperforms BLOOMZ on Swahili (SW), possibly due to it being a larger part of its pretraining corpus (see Figure 2 and §4.6).



Figure 10: Zero-shot multilingual task generalization on languages seen during pretraining and finetuning. Each dot represents performance on one English evaluation prompt.

C Artifacts

Table 3 lists all artifacts used or released in this work. We make all our work accessible under the most permissive licenses available to us.

Artifact	Explanation	Public link
ROOTS mC4 P3 xP3 xP3all xP3mt xP3megds xP3x	Multilingual pretraining corpus of BLOOM Multilingual pretraining corpus used for mT5 Multitask finetuning dataset with English data & English prompts Multitask finetuning dataset with multilingual data & English prompts Same as xP3 with held-out evaluation sets Same as xP3 with English & multilingual machine-translated prompts Processed version of xP3 for easy usage with Megatron-DeepSpeed Extension of xP3 to 277 languages	https://huggingface.co/bigscience-data https://huggingface.co/datasets/mc4 https://huggingface.co/datasets/bigscience/P3 https://huggingface.co/datasets/bigscience/xP3 https://huggingface.co/datasets/bigscience/xP3a11 https://huggingface.co/datasets/bigscience/xP3mt https://huggingface.co/datasets/bigscience/xP3megds https://huggingface.co/datasets/Muennighoff/xP3x
XGLM-7.5B T0-11B mTk-Instruct-3.7B mTk-Instruct-13B	7.5B parameter pretrained multilingual transformer 11B parameter model finetuned on P3 3.7B parameter multitask finetuned multilingual transformer 13B parameter multitask finetuned multilingual transformer	https://huggingface.co/facebook/xglm-7.5B https://huggingface.co/bigscience/t0 https://huggingface.co/allenai/mtk-instruct-3b-def-pos https://huggingface.co/allenai/mtk-instruct-11b-def-pos
BLOOM-560M BLOOM-1.1B BLOOM-1.7B BLOOM-3B BLOOM-7.1B BLOOM	560M parameter model pretrained on ROOTS 1.1B parameter model pretrained on ROOTS 1.7B parameter model pretrained on ROOTS 3B parameter model pretrained on ROOTS 7.1B parameter model pretrained on ROOTS 176B parameter model pretrained on ROOTS	https://huggingface.co/bigscience/bloom-560m https://huggingface.co/bigscience/bloom-1b1 https://huggingface.co/bigscience/bloom-1b7 https://huggingface.co/bigscience/bloom-3b https://huggingface.co/bigscience/bloom-7b1 https://huggingface.co/bigscience/bloom
BLOOMZ-560M BLOOMZ-1.1B BLOOMZ-1.7B BLOOMZ-3B BLOOMZ-7.1B BLOOMZ-7.1B-MT BLOOMZ-7.1B-P3 BLOOMZ BLOOMZ-MT BLOOMZ-P3	560M parameter model finetuned on xP3 1.1B parameter model finetuned on xP3 1.7B parameter model finetuned on xP3 3B parameter model finetuned on xP3 7.1B parameter model finetuned on xP3 7.1B parameter model finetuned on xP3mt 7.1B parameter model finetuned on P3 176B parameter model finetuned on xP3 176B parameter model finetuned on xP3mt 176B parameter model finetuned on P3	https://huggingface.co/bigscience/bloomz-560m https://huggingface.co/bigscience/bloomz-1b1 https://huggingface.co/bigscience/bloomz-1b7 https://huggingface.co/bigscience/bloomz-3b https://huggingface.co/bigscience/bloomz-7b1 https://huggingface.co/bigscience/bloomz-7b1-mt https://huggingface.co/bigscience/bloomz-7b1-p3 https://huggingface.co/bigscience/bloomz https://huggingface.co/bigscience/bloomz-mt https://huggingface.co/bigscience/bloomz-p3
mT5-300M mT5-580M mT5-1.2B mT5-3.7B mT5-13B	300M parameter model pretrained on a sampled version of mC4 580M parameter model pretrained on a sampled version of mC4 1.2B parameter model pretrained on a sampled version of mC4 3.7B parameter model pretrained on a sampled version of mC4 13B parameter model pretrained on a sampled version of mC4	https://huggingface.co/google/mt5-small https://huggingface.co/google/mt5-base https://huggingface.co/google/mt5-large https://huggingface.co/google/mt5-xl https://huggingface.co/google/mt5-xxl
mT0-300M mT0-580M mT0-1.2B mT0-3.7B mT0-13B mT0-13B-MT mT0-13B-P3	300M parameter model finetuned on xP3 580M parameter model finetuned on xP3 1.2B parameter model finetuned on xP3 3.7B parameter model finetuned on xP3 13B parameter model finetuned on xP3 13B parameter model finetuned on xP3mt 13B parameter model finetuned on P3	https://huggingface.co/bigscience/mt0-small https://huggingface.co/bigscience/mt0-base https://huggingface.co/bigscience/mt0-large https://huggingface.co/bigscience/mt0-xl https://huggingface.co/bigscience/mt0-xxl https://huggingface.co/bigscience/mt0-xxl-mt https://huggingface.co/bigscience/mt0-xxl-p3

Table 3: Links to all models & datasets used as part of this work. BLOOMZ models have an additional repository containing the final optimizer states for training with Megatron-DeepSpeed that can be found by appending “-optimizer-states” to the respective URL. BLOOM(Z) models are released under the RAIL license, while mT5 / mT0 models are licensed under Apache 2.0

D ROOTS language contamination

While the BLOOM ROOTS corpus (Laurençon et al., 2022) was collected from 46 natural languages and 13 programming languages, we find that sentences from the same document do not always belong to the collected (meta) language. Some sentences use languages like Russian or Japanese that were not the intentionally collected parts. This “language contamination” may stem from “code-mixing” or different languages being used in code comments. To investigate the extent of contamination, we randomly sample 1% of the documents from ROOTS for a total of 51M documents. For each document, we use cld3² (Xue et al., 2020) to identify the languages used in each sentence and compare them with the meta language of the document. We summarize our results in Figure 11. It shows that ROOTS contains unintentionally collected languages, such as Burmese (my: 0.00003%), Thai (th: 0.006%), Turkish (tr: 0.03%), Greek (el: 0.03%), Russian (ru: 0.03%), Bulgarian (bg: 0.05%), Estonian (et: 0.06%), Haitian (ht: 0.12%), German (de: 0.21%), Italian (it: 0.28%) and Japanese (ja: 0.54%). These “unseen” languages only have

²<https://github.com/google/cld3>

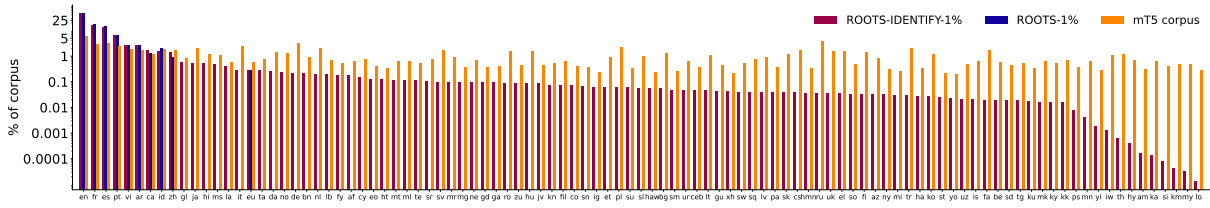


Figure 11: Language composition of ROOTS-IDENTIFY-1%, ROOTS-1% and the mT5 corpus. All mT5 languages are depicted. ROOTS-1% is a random 1% sample of ROOTS with its assigned meta-languages. ROOTS-IDENTIFY-1% are the actual languages in ROOTS-1% re-identified using c1d3.

small sentence proportions in our subsample compared to English (en: 46.23%), French (fr: 15.73%) and Spanish (es: 13.38%). Yet, they may help the language generalization of BLOOMZ models described in §4.2. Japanese is mostly mixed in the meta English documents (47%), meta Code documents (8%) and meta Chinese documents (5%). Meanwhile, Russian is mostly mixed in the meta English documents (52%), meta Code documents (19%) and meta French documents (11%).

E Code generations

Table 4 provides statistics on code generations and code data. We find that BLOOM generates on average 70% more characters and 17x more comments than BLOOMZ for a given problem from HumanEval. Figure 12 compares an example solution from BLOOM and BLOOMZ. While both solutions are correct, BLOOMZ is biased towards short and concise answers.

```

from typing import List

def has_close_elements(numbers: List[float], threshold: float) -> bool:
    """ Check if in given list of numbers, are any two numbers closer to each other than
    given threshold.
    """
    >>> has_close_elements([1.0, 2.0, 3.0], 0.5)
    False
    >>> has_close_elements([1.0, 2.8, 3.0, 4.0, 5.0, 2.0], 0.3)
    True
    """
    if len(numbers) <= 1:
        return False
    numbers = sorted(numbers)
    i = 0
    while i < len(numbers) - 1:
        # Get distance between current and next elements
        distance = numbers[i] - numbers[i + 1]

        # Check if distance is closer than threshold
        if distance < threshold:
            return True
        i += 1
    return False

```

(a) BLOOM

```

from typing import List

def has_close_elements(numbers: List[float], threshold: float) -> bool:
    """ Check if in given list of numbers, are any two numbers closer to each other than
    given threshold.
    """
    >>> has_close_elements([1.0, 2.0, 3.0], 0.5)
    False
    >>> has_close_elements([1.0, 2.0, 3.0, 4.0, 5.0, 2.0], 0.3)
    True
    """
    numbers = sorted(numbers)
    return any(numbers[i + 1] - numbers[i] < threshold for i in range(len(numbers) - 1))

```

(b) BLOOMZ

Figure 12: Code generations of BLOOM and BLOOMZ on HumanEval. The model is prompted to generate after the final """ . The generation is stopped after an end-of-sequence token or a return statement followed by a newline.

Data (→)	HumanEval generations		Fine-tuning data in xP3 (code data)
	BLOOM	BLOOMZ	
Average characters	247	144	531
Average Python comments (#)	0.69	0.04	0.85

Table 4: Number of characters and comments for generations and fine-tuning data. For finetuning data, the statistics are computed for the targets that the model is tasked to generate, not the input.

F Qualitative examples

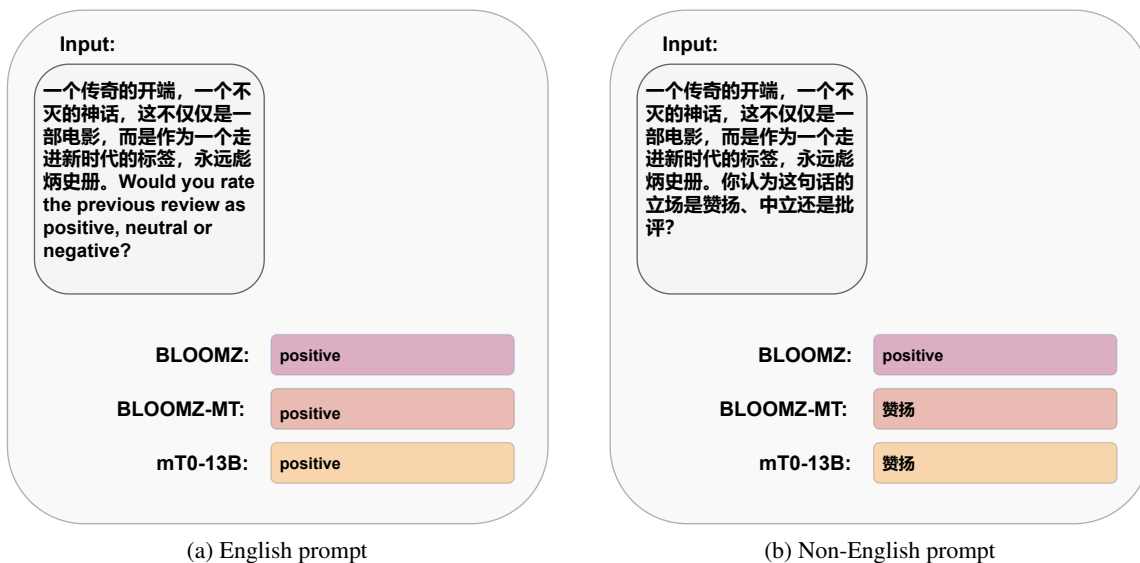


Figure 13: Greedy generations for sentiment analysis, a task trained on. BLOOMZ and mT0-13B have not been trained on non-English prompts, but are still able to handle them. BLOOMZ, however, answers in English. The review is a five star review of Star Wars Episode IV.

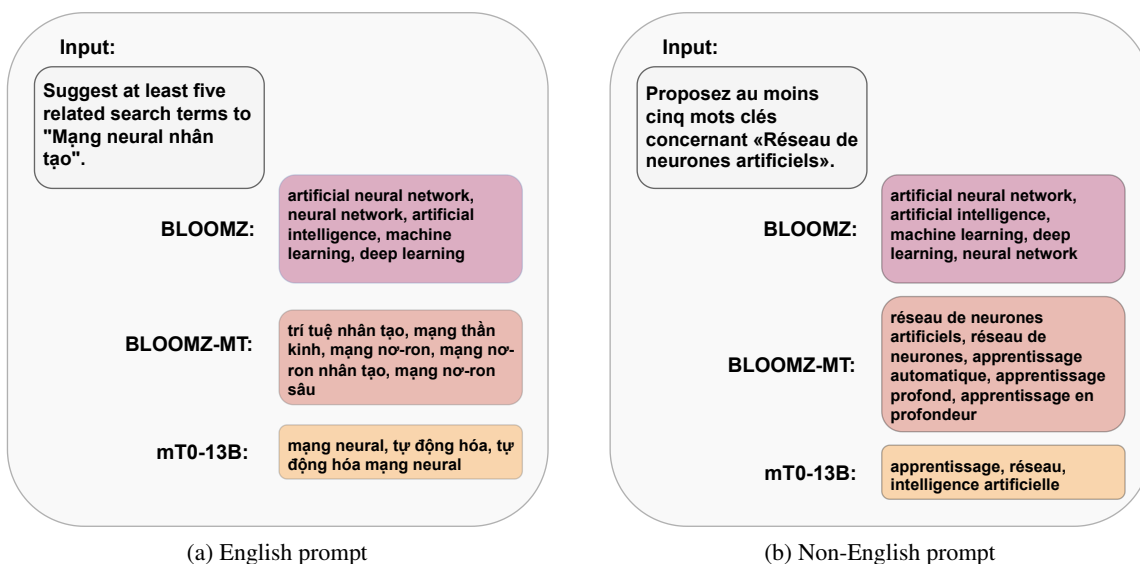


Figure 14: Greedy generations for zero-shot query expansion, a task not trained on. The models sometimes fail to output at least five terms as requested in the prompt.

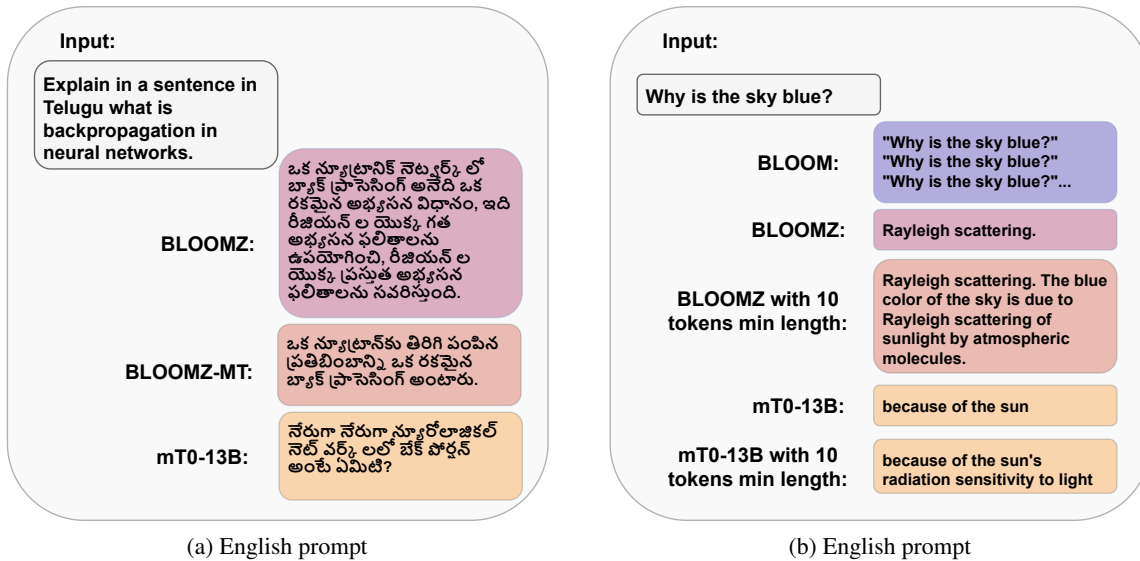


Figure 15: Greedy generations on question answering, a task trained on. **Left:** Specifying the language in the prompt is an effective way to force the output language. **Right:** Setting a minimum token length as a generation hyperparameter is an effective way to force long generations. The output of BLOOM is shortened (marked with ...).

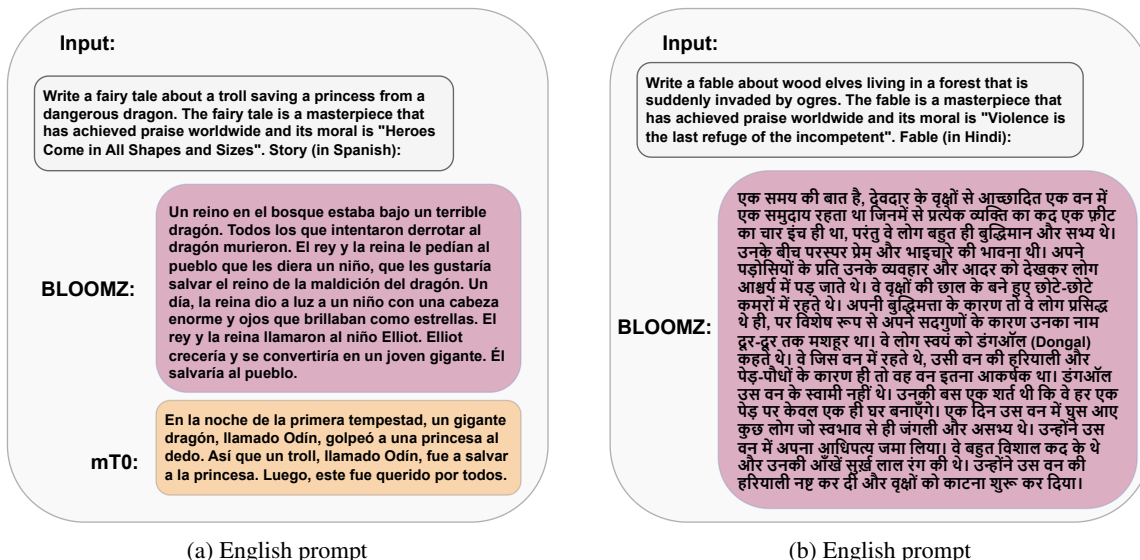


Figure 16: Non-greedy fable generations given a moral, a task not trained on. The generations are cherry-picked from 16 outputs with no minimum length, a temperature of 0.9 and top k of 40. **Left:** BLOOMZ generates an interesting fable with the desired moral. mT0 is significantly worse at writing stories likely due to its different pretraining objective. **Right:** BLOOMZ does not seem to understand the moral correctly.

G Increasing generation length

In §4.5, we found performance on generative tasks to worsen in later stages of training. To investigate this problem further, we study a 7.1 billion parameter BLOOM model that is finetuned for 13 billion tokens, which results in a low BLEU score of 0 and very short generations as shown in Table 5 (Default). We can solve this problem with two high-level strategies: **(a)** Reducing short tasks during finetuning and **(b)** Forcing a minimum generation length.

For **(a)**, we do so by either early stopping, upweighting long tasks or adding new long tasks. As the majority of our finetuning data are single sentences, early stopping has the effect of finetuning on fewer short sentences. Upweighting long tasks is done by removing the loss normalization explained in §3.2. This has the effect of each token getting equal weight regardless of the task, which upweights long tasks, as they have more tokens. Finally, for adding long tasks, we add tasks that require multi-sentence generations, such as generating an entire news article given a title. These long tasks collectively make up 10% of finetuning data for this ablation. All three solutions result in longer average generations as shown in Table 5 and slightly better BLEU scores, albeit effects are still small.

For **(b)**, we force the model to generate a minimum number of tokens at inference. Our benchmarking task, MultiEURLEX (Chalkidis et al., 2021), requires multi-sentence generations with an average target length of 1965 characters (about 491 tokens). By forcing the model to generate at least 768 tokens, we ensure that the generation is at least as long as the target. This boosts the BLEU score significantly to 9.05. This approach is thus an effective strategy to maintain long generations of good quality.

For our final models, we employ early stopping, adding of long tasks and recommend forcing a minimum generation length at inference for long generations. We do not upweight longer tasks, as it worsens accuracy on our NLU validation tasks by 10%. The number of tokens our final models are fine-tuned for are displayed in Table 6.

Model	Finetuning tokens	BLEU Score	Average generation length (characters)
Default	13 billion	0.00	122
Early stopping	6 billion	0.00	155
Upweight longer tasks	13 billion	0.06	364
Add more long tasks	13 billion	0.06	136
Forcing 768 tokens at inference	13 billion	9.05	3072

Table 5: 7.1 billion parameter BLOOMZ models with various modifications benchmarked on MultiEURLEX English-French translation (Chalkidis et al., 2021). We benchmark three prompts on both English to French and French to English translation. We then take the median performance across the three prompts for each translation direction and average the two scores to arrive at the BLEU score reported.

Model	mT0-300M	mT0-560M	mT0-1.2B	mT0-3.7B	mT0-13B	
Tokens	4.62	4.62	4.62	1.85	1.29	
Model	BLOOMZ-560M	BLOOMZ-1.1B	BLOOMZ-1.7B	BLOOMZ-3B	BLOOMZ-7.1B	BLOOMZ
Tokens	3.67	0.502	8.39	8.39	4.19	2.09

Table 6: Tokens in billions that final models are finetuned for. We early-stop models based on validation performance. For -MT and -P3 variants we take the checkpoint after the same number of steps as for their default versions.

H XNLI edit distances

As models are surprisingly capable of solving XNLI in languages they were never intentionally trained on (§4.2), we investigate whether XNLI can be solved without any language understanding. To do so, we compute edit distances using the Levenshtein methodology (Levenshtein et al., 1966) between premise

Premise	Hypothesis	Lev. distance	Label
probably so probably so um-hum	probably yes so uh-huh	13	Entailment
equivalent to increasing national saving to 19 .	National savings are 18 now .	34	Neutral
The Inglethorps did not appear .	The Inglethorps were the first ones to turn up .	26	Contradiction

Table 7: Three samples from the English XNLI split. To solve XNLI models need to classify whether the premise entails, is neutral to or contradicts the hypothesis. Samples are cherry-picked.

and hypothesis. Table 7 shows three samples from the English XNLI and their edit distances. Our hypothesis is that entailment pairs generally need to cover similar content, and thus have similar distance. Contradiction pairs still need to cover similar content but differ in at least one major way. Meanwhile for neutral pairs, hypothesis and premise may be about completely different topics, hence they should have the highest distance. In Table 8 we compute distances across all Thai, Turkish and Greek samples, three languages where we found language generalization to occur for BLOOMZ. Results confirm our hypothesis that distances are generally largest for neutral samples and smallest for entailment samples. However, the aggregate differences are very small with only a few edits difference. For example, Thai contradiction samples only have 2.5 edits more on average than entailment samples. Thus, comparing characters based on edit distance alone is likely not sufficient to fully explain the language generalization of models in §4.2.

Label (→)	Entailment	Neutral	Contradiction
Language (↓)			
Thai (th)	79.08	82.64	81.52
Turkish (tr)	76.93	80.59	80.24
Greek (el)	90.90	95.10	93.93

Table 8: Levenshtein distances between hypothesis and premise averaged across samples from different XNLI labels. Each label has 830 samples per language subset.

I Multilingual prompting in unseen languages

Table 9 shows aggregate performances on languages not intentionally seen during pretraining nor fine-tuning for BLOOMZ and only seen during pretraining for mT0. For BLOOMZ, performance drops significantly when translating the prompts to the respective unseen languages. Unlike on translated prompts for seen languages (§4.3), BLOOMZ-MT performs worse than BLOOMZ for machine-translated prompts in unseen languages. This is likely because BLOOMZ-MT has not been finetuned on prompts in these languages. For mT0 differences are less significant.

Task	Prompt	Average accuracy			
		BLOOMZ	BLOOMZ-MT	mT0-13B	mT0-13B-MT
XNLI	EN	45.65	43.2	48.52	51.33
	MT	36.48	35.67	41.86	39.78
XCOPA	EN	54.27	53.67	72.67	71.6
	MT	53.2	53.0	71.57	70.87
XStoryCloze	EN	61.59	61.36	79.31	80.13
	MT	60.5	59.91	80.21	80.28
XWinograd	EN	55.98	54.54	70.81	72.0
	MT	53.11	52.46	67.86	70.45

Table 9: Comparison between EN (English) and MT (machine-translated) prompts for 176B BLOOMZ and 13B mT0 models finetuned on either only English or English and machine-translated multilingual prompts (-MT). For BLOOMZ the evaluation languages averaged are never intentionally seen, such as Japanese and Russian for XWinograd (see Figure 5). For mT0 the evaluation languages are only seen during pretraining.

J Ideas that did not work

We list several experiments that did not improve over baseline results:

Non-causal In a non-causal or prefix language model, the model attends bidirectionally over input tokens and only causally over target tokens. Given a pretrained causal decoder, other work found that multitask finetuning in a non-causal setup performed better than causal finetuning (Wang et al., 2022a; Tay et al., 2022c). However, in our experiments, non-causal finetuning did not improve over causal finetuning.

Special tokens Instead of separating inputs and targets with a space, we experimented with special tokens. Using the end-of-sequence token as a separator or a completely new token that the model would learn during finetuning significantly worsened results. The models may need to train on more tokens, possibly even during pretraining, to learn these new special tokens (Zeng et al., 2022).

Fixing prompts PromptSource has been written with encoder-decoder models in mind, where inputs and targets are fed into different models. As a consequence, human-written prompts in PromptSource often lack separators between input and target. For our decoder models, we decided to separate them with a space. We additionally experimented with leaving them as is or rewriting a significant amount of prompts, but neither improved significantly over space separation.

BitFit Previous work has shown bias-only finetuning (Zaken et al., 2021) of large language models to be sufficient for strong downstream performance (Logan et al., 2021; Hu et al., 2021; Muennighoff, 2022; Liu et al., 2022; Ding et al., 2022; Muennighoff et al., 2022). We found multitask finetuning of only biases to perform 15 absolute percentage points worse on the average of held-out tasks for BLOOMZ-7.1B.

K Full results

Table 10 shows all evaluation results on test datasets. Table 11 displays evaluation results on validation datasets which we use for checkpoint selection.

Task	Dataset	Config	Split	Prompt	Metric	mT0-300M		mT0-560M		mT0-1.2B		mT0-3.7B		mT0-13B		BLOOMZ-560M		BLOOMZ-1.1B		BLOOMZ-1.7B		BLOOMZ-3B		BLOOMZ-7.1B		BLOOMZ					
						Median acc.	Max acc.	Median acc.	Max acc.	Median acc.	Max acc.	Median acc.	Max acc.	Median acc.	Max acc.	Median acc.	Max acc.	Median acc.	Max acc.	Median acc.	Max acc.	Median acc.	Max acc.	Median acc.	Max acc.	Median acc.	Max acc.	Median acc.	Max acc.	Median acc.	Max acc.
Extractive QA	craftslist_bargains	bargains	validation	EN	Median acc.	30.49	23.95	22.61	39.61	25.96	28.14	47.99	28.14	22.86	46.48	26.47															
Extractive QA	craftslist_bargains	bargains	validation	EN	Max acc.	49.41	28.14	31.32	50.92	40.54	72.53																				
Grammar Correction	blimp_adjunct	island	validation	EN	Median acc.	50.40	51.60	51.80	53.80	55.10	51.60	52.30	50.60	49.20	49.90	49.80															
Grammar Correction	blimp_adjunct	island	validation	EN	Max acc.	50.90	57.00	58.00	59.10	56.80	77.10	60.90	62.30	59.90	57.60	51.60															
Grammar Correction	glue	cola	validation	EN	Median acc.	30.97	38.26	36.57	35.19	45.83	31.26	57.81	31.16	31.75	44.58	53.15															
Grammar Correction	glue	cola	validation	EN	Max acc.	64.33	51.01	62.80	47.17	58.29	41.71	67.98	46.40	65.39	56.86	63.37															
Multiple-Choice QA	aqua_rat	raw	validation	EN	Median acc.	27.95	25.20	24.80	20.47	16.14	19.29	22.83	22.05	22.44	24.41	27.56															
Multiple-Choice QA	aqua_rat	raw	validation	EN	Max acc.	29.53	26.38	25.50	21.65	18.90	20.08	22.83	22.83	22.44	25.20	28.35															
Multiple-Choice QA	codah	codah	train	EN	Median acc.	25.25	25.43	26.48	55.04	75.58	24.93	24.35	57.17	64.12	73.60	80.66															
Multiple-Choice QA	codah	train	EN	Max acc.	25.32	26.15	27.13	55.44	76.22	25.04	24.60	57.31	64.41	73.67	80.91																
Multiple-Choice QA	commonsense_qa	qa	validation	EN	Median acc.	31.20	37.43	36.61	56.35	69.53	43.98	38.90	69.86	84.44	83.05	80.26															
Multiple-Choice QA	commonsense_qa	qa	validation	EN	Max acc.	31.53	37.51	39.72	60.03	69.94	44.47	42.42	72.40	84.60	84.36	83.05															
Multiple-Choice QA	head_qa	en	validation	EN	Median acc.	24.89	24.38	23.43	27.53	36.02	26.72	27.16	27.53	30.01	38.58	53.15															
Multiple-Choice QA	head_qa	en	validation	EN	Max acc.	25.55	25.62	26.87	31.55	36.16	27.75	27.67	33.31	35.21	40.92	53.95															
Multiple-Choice QA	head_qa	es	validation	EN	Median acc.	24.60	24.45	23.94	27.89	34.92	26.94	25.04	24.45	26.21	34.41	50.81															
Multiple-Choice QA	head_qa	es	validation	EN	Max acc.	26.21	26.21	24.74	29.50	37.04	28.26	26.28	29.87	33.02	39.75	51.76															
Multiple-Choice QA	math_qa	qa	test	EN	Median acc.	21.11	20.00	22.18	23.25	23.69	19.66	21.21	20.97	21.81	21.14	21.84															
Multiple-Choice QA	math_qa	qa	test	EN	Max acc.	22.21	26.03	35.64	24.89	26.60	45.56	27.94	35.24	43.28	38.12	47.37															
Multiple-Choice QA	mnsc	mnsc	validation	EN	Median acc.	50.00	52.44	54.88	60.98	74.39	53.66	52.44	56.10	58.54	62.20	71.95															
Multiple-Choice QA	mnsc	mnsc	validation	EN	Max acc.	52.44	53.66	57.32	63.85	79.27	58.54	57.32	58.54	63.41	69.51	80.49															
Multiple-Choice QA	pubmed_qa	labeled	train	EN	Median acc.	45.55	54.50	55.75	58.35	65.35	55.75	58.90	66.75	66.80	67.15	71.80															
Multiple-Choice QA	pubmed_qa	labeled	train	EN	Max acc.	48.60	57.60	58.30	58.60	66.20	57.50	63.50	72.10	69.80	69.50	74.40															
Multiple-Choice QA	riddle_sense	riddle_sense	validation	EN	Median acc.	24.39	22.04	23.41	29.63	43.14	22.87	24.53	30.02	35.11	39.47	50.64															
Multiple-Choice QA	riddle_sense	sense	validation	EN	Max acc.	34.48	33.30	33.01	39.18	47.50	37.41	39.86	43.58	47.60	48.09	59.26															
Sentiment	amazon_reviews_multi	en	validation	EN	Median acc.	40.60	50.80	51.12	49.00	53.24	46.52	42.46	50.48	49.88	51.00	50.90															
Sentiment	amazon_reviews_multi	en	validation	EN	Max acc.	41.34	53.88	54.18	55.92	57.04	50.44	47.74	55.94	53.74	55.08	54.16															
Sentiment	amazon_reviews_multi	es	validation	EN	Median acc.	39.56	48.70	49.02	47.56	52.30	37.60	38.92	45.08	43.28	44.44	45.26															
Sentiment	amazon_reviews_multi	es	validation	EN	Max acc.	42.66	51.00	50.42	50.68	53.58	39.10	40.24	47.98	46.28	47.76	44.48															
Sentiment	amazon_reviews_multi	fr	validation	EN	Median acc.	38.74	48.44	48.32	46.12	51.12	38.78	38.38	44.36	45.84	44.92	43.92															
Sentiment	amazon_reviews_multi	fr	validation	EN	Max acc.	40.66	49.64	49.30	49.30	52.40	41.16	40.04	46.66	46.80	47.42	44.90															
Sentiment	amazon_reviews_multi	zh	validation	EN	Median acc.	38.58	42.58	42.58	39.63	45.30	34.54	34.54	41.56	44.78	40.48	48.48															
Sentiment	amazon_reviews_multi	zh	validation	EN	Max acc.	37.88	44.36	44.74	43.66	47.14	39.48	35.24	43.52	39.64	47.12	42.10															
Sentiment	financial_phrasebank	allagree	train	EN	Median acc.	18.33	28.98	28.09	25.44	35.25	31.10	29.28	34.76	35.91	34.89	24.82															
Sentiment	financial_phrasebank	allagree	train	EN	Max acc.	22.22	57.51	52.25	68.15	37.77	44.79	34.81	54.37	59.23	37.15	37.23															
Sentiment	financial_phrasebank	allagree	train	EN	Median acc.	79.70	83.49	83.37	83.80	93.58	87.96	83.72	92.09	94.50	94.04	93.92															
Sentiment	glue	sst2	validation	EN	Max acc.	81.88	87.96	86.81	91.51	94.84	92.89	89.79	94.15	95.87	94.61	95.07															
Sentiment	lince	spang	validation	EN	Median acc.	43.63	43.09	49.11	41.69	54.81	58.04	53.85	52.82	50.19	58.15	59.60															
Sentiment	lince	spang	validation	EN	Max acc.	56.91	56.05	56.37	57.78	56.80	58.53	55.35	56.37	54.60	60.09	58.47															
Sentiment	movie_rationales	rationales	validation	EN	Median acc.	63.50	78.00	81.00	69.50	90.00	93.50	97.50	98.50	97.50	98.50	98.50															
Sentiment	movie_rationales	rationales	validation	EN	Max acc.	94.50	95.50	98.50	99.50	100.00	98.50	97.50	100.00	99.50	99.00	99.50															
Sentiment	poem_sentiment	sentiment	validation	EN	Median acc.	17.14	18.10	16.19	20.67	25.67	20.95	29.52	24.76	24.76	23.81	23.81															
Sentiment	poem_sentiment	sentiment	validation	EN	Max acc.	18.10	23.81	20.00	27.62	27.62	22.86	33.33	29.52	31.43	29.52	24.76															
Summarization	nlsum	es	validation	EN	Median BLEU	0.18	0.18	0.18	0.19	0.19	0.20	0.18	0.19	0.19	0.20	0.19															
Summarization	nlsum	es	validation	EN	Max BLEU	2.91	3.51	3.46	3.72	4.21	3.62	2.87	3.23	3.84	4.82	4.16															
Text Classification	art	art	validation	EN	Median acc.	50.85	50.85	50.46	53.33	68.99	51.50	50.07	52.68	54.57	58.42	66.58															
Text Classification	art	art	validation	EN	Max acc.	51.04	51.83	51.76	56.07	69.71	52.68	50.65	54.24	57.31	61.10	67.43															
Text Classification	climate_fever	fever	test	EN	Median acc.	10.62	25.28	10.94	26.78	29.97	45.34	10.36	51.92	10.81	43.97	18.63															
Text Classification	climate_fever	fever	test	EN	Max acc.	42.41	43.78	20.98	43.32	51.01	63.97	30.94	65.54	32.12	47.69	36.61															
Text Classification	conv_ai_3	3	validation	EN	Median acc.	35.15	38.5																								

In the following, dataset naming conventions follow those used in the Hugging Face datasets library. Since xP3 expands upon the P3 dataset employed by Sanh et al. (2022), we refer the reader to that work for example prompts from datasets that fall within P3. Here, we provide prompts curated for datasets that belong to xP3 but not to P3. The prompts provided are not exhaustive. Code will be released to provide a canonical reference. For each dataset considered, a dataset example is provided for additional context. Next, it is noted if the prompt does not match the original task formulation of the dataset. This is followed by a reference for the data, an input template and a target template. For prompts with predefined answer choices, these are also included. To provide examples of both human-translated and machine-translated prompts, samples of each kind are included for the XNLI ES dataset.

CONTENTS

1 Prompts

1.1	Simplification
1.1.1	GEM/BiSECT en
1.1.2	GEM/BiSECT es
1.1.3	GEM/BiSECT fr
1.2	Summarization
1.2.1	GEM/wiki_lingua en
1.2.2	GEM/wiki_lingua es
1.2.3	GEM/xlsum bengali
1.2.4	GEM/xlsum english
1.3	Translation
1.3.1	Helsinki-NLP/tatoeba_mt ben-eng
1.3.2	Helsinki-NLP/tatoeba_mt eng-fra
1.3.3	facebook/flores ben_Beng-eng_Latn
1.3.4	facebook/flores ben_Beng-fra_Latn
1.4	Program Synthesis
1.4.1	Muennighoff/mbpp sanitized
1.4.2	codeparrot/apps all
1.4.3	codeparrot/github-jupyter-text-code-pairs
1.4.4	codeparrot/xlcost-text-to-code C++-program-level
1.4.5	codeparrot/xlcost-text-to-code C-program-level
1.4.6	codeparrot/xlcost-text-to-code Csharp-program-level
1.4.7	codeparrot/xlcost-text-to-code Java-program-level
1.4.8	codeparrot/xlcost-text-to-code Javascript-program-level
1.4.9	codeparrot/xlcost-text-to-code PHP-program-level
1.4.10	codeparrot/xlcost-text-to-code Python-program-level
1.4.11	neural_code_search evaluation_dataset
1.4.12	teven/code_contests
1.5	Coreference Resolution

1.5.1	Muennighoff/xwinograd en
1.5.2	Muennighoff/xwinograd fr
1.6	Question Answering Multiple Choice
1.6.1	clue c3
1.7	Question Answering Extractive
1.7.1	clue cmrc2018
1.7.2	clue drcd
1.7.3	mlqa mlqa.vi.vi
1.7.4	mlqa mlqa.zh.zh
1.7.5	xquad xquad.vi
1.7.6	xquad xquad.zh
1.8	Topic Classification
1.8.1	clue csl
1.8.2	clue tnews
1.9	Code Misc.
1.9.1	codeparrot/codecomplex codeparrot-codecomplex
1.9.2	great_code
1.9.3	teven/code_docstring_corpus top_level
1.10	Word Sense Disambiguation
1.10.1	pasinit/xlwic xlwic_en_zh
1.10.2	pasinit/xlwic xlwic_fr_fr
1.11	Paraphrase Identification
1.11.1	paws-x en
1.11.2	paws-x es
1.12	Sentence Completion
1.12.1	xcopa vi
1.12.2	xcopa zh
1.13	Natural Language Inference
1.13.1	xnli en
1.13.2	xnli es
	1.13.2.1 Human-translated prompts
	1.13.2.2 Machine-translated prompts

1 PROMPTS

1.1 SIMPLIFICATION

1.1.1 GEM/BISECT EN

Dataset from Kim et al. (2021). Used in training.

Data Example

Key	Value
gem_id	BiSECT_en-train-1
source	To view any of the video clips belo...
target	If you want to watch one of the vid...
references	If you want to watch one of the vid...

Prompts

Input Template:

```
Split and simplify the following sentence while retaining its full meaning:  
{{source}}  
Simplified version:
```

Target Template:

```
{{target}}
```

Input Template:

```
{{source}}  
The above sentence is very complicated. Please provide me a simplified synonymous  
version consisting of multiple sentences:
```

Target Template:

```
{{target}}
```

Input Template:

```
{{source}}. This sentence is hard to understand. A simpler version with equivalent  
meaning is the following:
```

Target Template:

```
{{target}}
```

1.1.2 GEM/BiSECT ES

Data Example

Key	Value
gem_id	BiSECT_es-train-1
source	Al final de la Santa Misa , mientra...
target	Al finalizar la santa misa , mientr...
references	Al finalizar la santa misa , mientr...

Prompts

Input Template:

```
{{source}}. Esta frase es difícil de entender. Una versión más simple con significado equivalente es la siguiente:
```

Target Template:

```
{{target}}
```

Input Template:

```
Divida y simplifique la siguiente oración conservando su significado completo:  
{{source}}  
Versión simplificada:
```

Target Template:

```
{{target}}
```

Input Template:

```
{{source}}  
La frase anterior es muy complicada. Por favor, proporcione una versión sinónima simplificada que consta de varias oraciones:
```

Target Template:

```
{{target}}
```

1.1.3 GEM/BiSECT FR

Data Example

Prompts

Input Template:

Key	Value
gem_id	BiSECT_fr-train-1
source	N'ayez pas peur de poser des questi...
target	Il ne faut pas avoir peur de poser ...
references	Il ne faut pas avoir peur de poser ...

Divisez et simplifiez la phrase suivante tout en conservant son sens complet :
{{source}}
Version simplifiée :

Target Template:

{{target}}

Input Template:

{{source}}
La phrase ci-dessus est très compliquée. Veuillez me fournir une version synonyme simplifiée composée de plusieurs phrases :

Target Template:

{{target}}

Input Template:

{{source}}. Cette phrase est difficile à comprendre. Une version plus simple avec une signification équivalente est la suivante :

Target Template:

{{target}}

1.2 SUMMARIZATION

1.2.1 GEM/WIKI_LINGUA EN

Dataset from lad (2020). Used in training.

Data Example

Prompts

Notes: xsum DOC_write_summary_of_above template

Key	Value
gem_id	wikilingua_multilingual-train-42437...
gem_parent_id	wikilingua_multilingual-train-42437...
source_language	en
target_language	en
source	Go online and simply search "Decor ...
target	Take a quiz online to find your sty...
references	Take a quiz online to find your sty...

Input Template:

```
{{source}}
```

```
===
```

Write a summary of the text above in English :

Target Template:

```
{{target}}
```

Notes: xsum 'article_DOC_summary' template

Input Template:

```
Article in English: {{source}}
```

```
Summary in English:
```

Target Template:

```
{{target}}
```

Notes: xsum 'DOC_how_would_you_rephrase_few_words' template

Input Template:

```
{{source}}
```

```
How would you rephrase that briefly in English?
```

Target Template:

```
{{target}}
```

Notes: xsum 'DOC_tldr' template

Input Template:

```
{{source}}
```

TL;DR in English:

Target Template:

```
{{target}}
```

Notes: xsum 'read_below_DOC_write_abstract' template

Input Template:

```
First, read the English article below.
```

```
{{source}}
```

```
Now, please write a short abstract for it in English.
```

Target Template:

```
{{target}}
```

Input Template:

```
{{target}}
```

```
Given the above abstract, write an English article for it.
```

Target Template:

```
{{source}}
```

Input Template:

```
{{target}}
```

```
I'm interested in that, but I only have a few mins. Can you give me the first 500 characters of an article about that?
```

Target Template:

```
{{source[:500]}}
```

1.2.2 GEM/WIKI_LINGUA ES

Data Example

Key	Value
gem_id	wikilingua_multilingual-train-34808...
gem_parent_id	wikilingua_multilingual-train-34808...
source_language	es
target_language	es
source	Navega en la web y simplemente busc...
target	Haz un cuestionario en línea para e...
references	Haz un cuestionario en línea para e...

Prompts

Notes: xsum templates

Input Template:

```
{{source}}
```

```
===
```

Write a summary of the text above in Spanish:

Target Template:

```
{{target}}
```

Notes: xsum templates

Input Template:

First, read the Spanish article below.

```
{{source}}
```

Now, please write a short abstract for it in Spanish.

Target Template:

```
{{target}}
```

Notes: xsum templates

Input Template:

```
{{source}}
```

TL;DR in Spanish:

Target Template:

```
{{target}}
```

Notes: xsum templates

Input Template:

```
Article in Spanish: {{source}}
```

```
Summary in Spanish:
```

Target Template:

```
{{target}}
```

Notes: xsum templates

Input Template:

```
{{source}}
```

```
How would you rephrase that briefly in Spanish?
```

Target Template:

```
{{target}}
```

1.2.3 GEM/XLSUM BENGALI

Dataset from Hasan et al. (2021). Used in training.

Data Example

Key	Value
gem_id	xlsum_bengali-train-2
url	https://www.bbc.com/bengali/news-50...
title	রাশিয়ায় ক্ষমতার ২০ বছর যেভাবে কেট...
target	ভ্লাদিমির পুতিন তাঁর ক্ষমতায় থাকার...
references	ভ্লাদিমির পুতিন তাঁর ক্ষমতায় থাকার...
text	গত ২০ বছরে তিনি রাশিয়ার প্রেসিডেন্...

Prompts

Input Template:

একটি নিবন্ধের নীচের শিরোনাম এবং সারাংশ দেওয়া, একটি ছোট নিবন্ধ তৈরি করুন বা তাদের সাথে যেতে একটি দীর্ঘ নিবন্ধের শুরু করুন। শিরোনাম: `{{title}}` সারাংশ: `{{target}}`

Target Template:

`{{text[:500]}}`

Input Template:

বিষয়বস্তু: `{{text[:7000]}}`

Target Template:

`{{target}}`

Input Template:

ডক সংক্ষিপ্ত করার জন্য: `{{text[:8500]}}`

Target Template:

`{{target}}`

Input Template:

...`{{text[3000:3500]}}`

Target Template:

`{{text[5000:]}}`

Input Template:

শিরোনাম: `{{title}}`

Target Template:

`{{text[:7000]}}`

Input Template:

`{{text}}`

Target Template:

`{{title}}`

Input Template:

শিরোনাম: `{{title}}`

Target Template:

`{{text[:700]}}`

Input Template:

`{{title}}{{text[:500]}}`

Target Template:

`{{target}}`

Input Template:

`{{text[:1000]}}`

Target Template:

`{{text[1000:5000]}}`

1.2.4 GEM/XLSUM ENGLISH

Data Example

Key	Value
gem_id	xlsum_english-train-2
url	https://www.bbc.com/news/uk-scotland
title	Huge tidal turbine installed at Ork...
target	The massive tidal turbine AK1000 ha...
references	The massive tidal turbine AK1000 ha...
text	Atlantis Resources unveiled the mar...

Prompts

Input Template:

Doc to summarize: `{{text[:8500]}}`\nSummary in the same language as the doc:

Target Template:

`{{target}}`

Input Template:

Content: `{{text[:7000]}}`\n\nThe previous content can be summarized as follows:

Target Template:

`{{target}}`

Input Template:

`{{title}}`\n\n`{{text[:5000]}}`\n\nnt1;dr:

Target Template:

`{{target}}`

Input Template:

`{{text}}` \n\nGive me a good title for the article above.

Target Template:

`{{title}}`

Input Template:

Given the below title and summary of an article, generate a short article or the beginning of a long article to go along with them. Title: `{{title}}`\n\nSummary: `{{target}}`\n\nArticle (Max 500 characters):

Target Template:

`{{text[:500]}}`

Input Template:

Title: `{{title}}`\n\nGiven the above title of an imaginary article, imagine the article.\n

Target Template:

`{{text[:7000]}}`

Input Template:

Title: `{{title}}`\nGiven the above title of an imaginary article, imagine the article.\n

Target Template:

`{{text[:7000]}}`

Input Template:

`{{text[:1000]}}`... Continue the article for another 4000 characters max:

Target Template:

`{{text[1000:5000]}}`

Input Template:

...`{{text[3000:3500]}}`... Write the rest of the article:

Target Template:

`{{text[5000:]}}`

1.3 TRANSLATION

1.3.1 HELSINKI-NLP/TATOEBA__MT BEN-ENG

Dataset from Tiedemann (2020). Used in training.

Data Example

Key	Value
sourceLang	ben
targetlang	eng
sourceString	Tatoebaর অর্থ কী?
targetString	What does "Tatoeba" mean?

Prompts

Input Template:

Translate the following text from English to Bengali `{{ targetString }}`

Target Template:

```
{{ sourceString }}
```

Input Template:

```
Translate the following text from Bengali to English {{ sourceString }}
```

Target Template:

```
{{ targetString }}
```

1.3.2 HELSINKI-NLP/TATOEBA_MT ENG-FRA

Data Example

Key	Value
sourceLang	eng
targetlang	fra
sourceString	Aah. Now I understand.
targetString	Ah! Maintenant, je comprends.

Prompts

Input Template:

```
Translate the following text from French to English {{ targetString }}
```

Target Template:

```
{{ sourceString }}
```

Input Template:

```
Translate the following text from English to French {{ sourceString }}
```

Target Template:

```
{{ targetString }}
```

1.3.3 FACEBOOK/FLORES BEN_BENG-ENG_LATN

Dataset from NLLB (2022). Used in training.

Data Example

Key	Value
id	2
URL	https://en.wikinews.org/wiki/Scient...
domain	wikinews
topic	health
has_image	0
has_hyperlink	0
sentence_ben_Beng	শীর্ষ গবেষকরা বলছেন, এটি নিম্ন-আয়ে...
sentence_eng_Latn	Lead researchers say this may bring...

Prompts

Input Template:

```
{{sentence_ben_Beng}}
```

Target Template:

```
{{sentence_eng_Latn}}
```

Input Template:

```
A text in Bengali: {{sentence_ben_Beng}}
```

Target Template:

```
{{sentence_eng_Latn}}
```

Input Template:

```
{{sentence_ben_Beng}}
```

Target Template:

```
{{sentence_eng_Latn}}
```

1.3.4 FACEBOOK/FLORES BEN_BENG-FRA_LATN

Data Example

Key	Value
id	2
URL	https://en.wikinews.org/wiki/Scient...
domain	wikinews
topic	health
has_image	0
has_hyperlink	0
sentence_ben_Beng	শীর্ষ গবেষকরা বলছেন, এটি নিম্ন-আয়ে...
sentence_fra_Latn	Selon les chercheurs principaux, ce...

Prompts

Input Template:

```
{{sentence_ben_Beng}}
```

Target Template:

```
{{sentence_fra_Latn}}
```

Input Template:

```
{{sentence_ben_Beng}}
```

Target Template:

```
{{sentence_fra_Latn}}
```

Input Template:

```
A text in Bengali: {{sentence_ben_Beng}}
```

Target Template:

```
{{sentence_fra_Latn}}
```

1.4 PROGRAM SYNTHESIS

1.4.1 MUENNIGHOFF/MBPP SANITIZED

Dataset from Austin et al. (2021). Used in training.

Data Example

Prompts

Input Template:

Key	Value
source_file	Benchmark Questions Verification V2...
task_id	3
prompt	Write a python function to identify...
code	import math def is_not_prime(n): ...
test_imports	
test_list	assert is_not_prime(2) == False;ass...

`{{ prompt }}` Here is a solution in Python:

Target Template:

`{{ code }}`

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

`{{ prompt }}` This can be solved in Python with the following code:

Target Template:

`{{ code }}`

1.4.2 CODEPARROT/APPS ALL

Dataset from Hendrycks et al. (2021). Used in training.

Data Example

Key	Value
problem_id	1
question	Mikhail walks on a Cartesian plane...
solutions	["q=int(input())\n\nfor e in range(...
input_output	{ "inputs": ["3\n2 2 3\n4 3 ...
difficulty	interview
url	https://codeforces.com/problemset/p...
starter_code	

Prompts

Input Template:

Solve in Python:
`{{ question }}`

Target Template:

`{{ solution }}`

Input Template:

`{{ question }}`

Can you solve the above problem using Python?

Target Template:

`{{ solution }}`

Input Template:

I found an interesting problem on `{{url}}`:
`{{ question }}`

I tried it in Python, but could not do it. Can you solve it?

Target Template:

`{{ solution }}`

1.4.3 CODEPARROT/GITHUB-JUPYTER-TEXT-CODE-PAIRS

Data Example

Key	Value
markdown	Extract the dataset from the compre...
code	num_classes = 10 np.random.seed(133...
path	machine-learning/deep-learning/udac...
repo_name	pk-ai/training
license	mit

Prompts

Input Template:

```
"{{ markdown }}"
Please write code following the instructions in jupyter notebook style.
```

Target Template:

```
{{ code }}
```

Input Template:

```
I am working on the file "{{ path }}".
The first task is:
{{ markdown }}
Can you write Python code for it?
```

Target Template:

```
{{ code }}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
{{ markdown }}
```

Target Template:

```
{{ code }}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
{{ code }}
Given the above code, generate some markdown instructions for it.
```

Target Template:

```
{{ markdown }}
```

1.4.4 CODEPARROT/XLCOST-TEXT-TO-CODE C++-PROGRAM-LEVEL

Dataset from Zhu et al. (2022). Used in training.

Data Example

Key	Value
text	Check if a number can be represente...
code	#include <bits/stdc++.h> NEW_LINE u...

Prompts

Input Template:

```
"{{ text }}"  
Solution in C++:
```

Target Template:

```
{{ code_clean }}
```

Input Template:

```
"{{ text }}"  
How can the above be solved in C++?
```

Target Template:

```
{{ code_clean }}
```

1.4.5 CODEPARROT/XLCOST-TEXT-TO-CODE C-PROGRAM-LEVEL

Data Example

Key	Value
text	Logarithm tricks for Competitive Pr...
code	#include <stdio.h> NEW_LINE #includ...

Prompts

Input Template:

```
"{{ text }}"  
Solution in C:
```

Target Template:

```
{{ code_clean }}
```

Input Template:

```
{{ text }}  
How can the above be solved in C?
```

Target Template:

```
{{ code_clean }}
```

1.4.6 CODEPARROT/XLCOST-TEXT-TO-CODE CSHARP-PROGRAM-LEVEL

Data Example

Key	Value
text	Check if a number can be represente...
code	using System ; class GFG { static b...

Prompts

Input Template:

```
"{{ text }}"  
Solution in C#:
```

Target Template:

```
{{ code_clean }}
```

Input Template:

```
"{{ text }}"  
How can the above be solved in C-Sharp?
```

Target Template:

```
{{ code_clean }}
```

1.4.7 CODEPARROT/XLCOST-TEXT-TO-CODE JAVA-PROGRAM-LEVEL

Data Example

Key	Value
text	Check if a number can be represente...
code	import java . io . * ; class GFG { ...

Prompts

Input Template:

```
"{{ text }}"  
Solution in Java:
```

Target Template:

```
{{ code_clean }}
```

Input Template:

```
"{{ text }}"  
How can the above be solved in Java?
```

Target Template:

```
{{ code_clean }}
```

1.4.8 CODEPARROT/XLCOST-TEXT-TO-CODE JAVASCRIPT-PROGRAM-LEVEL

Data Example

Key	Value
text	Check if a number can be represente...
code	function sumOfTwoCubes (n) { var ...

Prompts

Input Template:

```
"{{ text }}"  
Solution in Javascript:
```

Target Template:

```
{{ code_clean }}
```

Input Template:

```
"{{ text }}"  
How can the above be solved in JS?
```

Target Template:

```
{{ code_clean }}
```

1.4.9 CODEPARROT/XLCOST-TEXT-TO-CODE PHP-PROGRAM-LEVEL

Data Example

Key	Value
text	Rearrange the array to maximize the...
code	< ? php function solve (\$ a , \$ n ...

Prompts

Input Template:

```
"{{ text }}"  
Solution in php:
```

Target Template:

```
{{ code_clean }}
```

Input Template:

```
"{{ text }}"  
How can the above be solved in PHP?
```

Target Template:

```
{{ code_clean }}
```

1.4.10 CODEPARROT/XLCOST-TEXT-TO-CODE PYTHON-PROGRAM-LEVEL

Data Example

Key	Value
text	Check if a number can be represente...
code	import math NEW_LINE def sumOfTwoCu...

Prompts

Input Template:

```
"{{ text }}"
Solution in Python:
```

Target Template:

```
{{ code_clean }}
```

Input Template:

```
"{{ text }}"
How can the above be solved in Python?
```

Target Template:

```
{{ code_clean }}
```

1.4.11 NEURAL_CODE_SEARCH EVALUATION_DATASET

Dataset from hug (2018). Used in training.

Data Example

Key	Value
stackoverflow_id	4616095
question	How to get the build/version number...
question_url	https://stackoverflow.com/questions...
question_author	Fahad Ali Shaikh
question_author_url	https://stackoverflow.com/users/565...
answer	try { PackageInfo pInfo = this.ge...
answer_url	https://stackoverflow.com/a/6593822
answer_author	plus-
answer_author_url	https://stackoverflow.com/users/709...
examples	4130029;3398176;2320640
examples_url	https://github.com/altanzio/Concei...

Prompts

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
Description:  
{{ question }}  
  
Implementation:
```

Target Template:

```
{{ answer }}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
Given the following code:  
{{ answer }}  
Describe it:
```

Target Template:

```
{{ question }}
```

1.4.12 TEVEN/CODE_CONTESTS

Data Example

Key	Value
name	1575_A. Another Sorting Problem
description	Andi and Budi were given an assignm...
source	2
difficulty	7
solution	#include <bits/stdc++.h> using name...
language	CPP

Prompts

Input Template:

```
{{description}}
```

Target Template:

```
{{solution}}
```

Input Template:

```
Can you solve the below in {{language}}?  
{{description}}
```

Target Template:

```
{{solution}}
```

Input Template:

```
{{description}}  
The above is tricky. Write me a correct solution in {{language}}.
```

Target Template:

```
{{solution}}
```

Input Template:

```
{{description}}  
Solve the task in {{language}}.
```

Target Template:

```
{{solution}}
```

Input Template:

```
{{description}}  
Using {{language | lower}} can you solve the prior task?
```

Target Template:

```
{{solution}}
```

Input Template:

```
{{description}}  
{{solution[:5]}}
```

Target Template:

```
{{solution[5:]}}
```

Input Template:

```
{{language}} solution for "{{description}}":
```

Target Template:

```
{{solution}}
```

1.5 COREFERENCE RESOLUTION

1.5.1 MUENNIGHOFF/XWINOGRAD EN

Dataset from Tikhonov and Ryabinin (2021). Used in evaluation.

Data Example

Key	Value
sentence	The city councilmen refused the dem...
option1	The city councilmen
option2	the demonstrators
answer	2

Prompts

Input Template:

```
{{sentence}}  
Replace the _ in the above sentence with the correct option:  
- {{option1}}  
- {{option2}}
```

Target Template:

```
{% if answer == '1' %} {{option1}} {% else %} {{ option2 }} {% endif %}
```

Answer Choices Template:

```
{{option1}} ||| {{option2}}
```

Input Template:

Fill in the _ in the below sentence:
{{sentence}}

Choices:
- {{ option1 }}
- {{ option2 }}

Answer:

Target Template:

```
{% if answer == '1' %} {{option1}} {% else %} {{ option2 }} {% endif %}
```

Answer Choices Template:

```
{{option1}} ||| {{option2}}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
The _ in the sentence below refers to {{option1}}. True or False?  
{{sentence}}
```

Target Template:

```
{{answer_choices[answer|int - 1]}}
```

Answer Choices Template:

```
True ||| False
```

Input Template:

```
{{ sentence }} In the previous sentence, does _ refer to {{ option1 }} or {{  
option2 }}?
```

Target Template:

```
{% if answer == '1' %} {{option1}} {% else %} {{ option2 }} {% endif %}
```

Answer Choices Template:

```
{{ option1 }} ||| {{ option2 }}
```

Input Template:

```
{{sentence}}
What does the _ in the above sentence refer to? {{ option1 }} or {{ option2 }}?
```

Target Template:

```
{% if answer == '1' %} {{option1}} {% else %} {{ option2 }} {% endif %}
```

Answer Choices Template:

```
{{option1}} ||| {{option2}}
```

Input Template:

```
In the sentence below, does the _ stand for {{answer_choices[0]}} or
{{answer_choices[1]}}?
{{sentence}}
```

Target Template:

```
{{answer_choices[answer | int - 1]}}
```

Answer Choices Template:

```
{{option1}} ||| {{option2}}
```

1.5.2 MUENNIGHOFF/XWINOGRAD FR

Data Example

Key	Value
sentence	La coupe n'entre pas dans la valise...
option1	La coupe
option2	la valise
answer	2

Prompts

Input Template:

```
{{ sentence }} Dans la phrase précédente, _ fait-il référence à {{ option1 }} ou
{{ option2 }} ?
```

Target Template:

```
{% if answer == '1' %} {{option1}} {% else %} {{ option2 }} {% endif %}
```

Answer Choices Template:

```
{{ option1 }} ||| {{ option2 }}
```

Input Template:

```
Dans la phrase ci-dessous, le _ signifie-t-il {{answer_choices[0]}} ou  
{{answer_choices[1]}} ?  
{{sentence}}
```

Target Template:

```
{{answer_choices[answer | int - 1]}}
```

Answer Choices Template:

```
{{option1}} ||| {{option2}}
```

Input Template:

```
{{sentence}}  
Remplacez le _ dans la phrase ci-dessus par l'option correcte :  
- {{option1}}  
- {{option2}}
```

Target Template:

```
{% if answer == '1' %} {{option1}} {% else %} {{ option2 }} {% endif %}
```

Answer Choices Template:

```
{{option1}} ||| {{option2}}
```

Input Template:

```
{{sentence}}  
À quoi le _ dans la phrase ci-dessus fait-il référence ? {{ option1 }} ou {{  
option2 }} ?
```

Target Template:

```
{% if answer == '1' %} {{option1}} {% else %} {{ option2 }} {% endif %}
```

Answer Choices Template:

```
{{option1}} ||| {{option2}}
```

Input Template:

```
Le _ dans la phrase ci-dessous fait référence à {{option1}}. Vrai ou faux?  
{{sentence}}
```

Target Template:

```
{{answer_choices[answer|int - 1]}}
```

Answer Choices Template:

```
Vrai ||| Faux
```

1.6 QUESTION ANSWERING MULTIPLE CHOICE

1.6.1 CLUE C3

Dataset from Sun et al. (2020). Used in training.

Data Example

Key	Value
id	1
context	男：足球比赛是明天上午八点开始吧?;女：因为天气不好，比赛改到后天下午...
question	根据对话，可以知道什么?
choice	今天天气不好;比赛时间变了;校长忘了时间
answer	比赛时间变了

Prompts

Input Template:

```
{% for statement in context %}  
{{ statement }}  
{% endfor %}  
鉴于上面的对话/段落，问题 “{{question}}” 的答案是什么
```

Target Template:

```
{{ answer }}
```

Input Template:

```
段落: {% for statement in context %}
{{ statement }}
{% endfor %}
什么样的问题会引起 {{ answer }} 的回答响应?
```

Target Template:

```
{{ question }}
```

Input Template:

```
{% for statement in context %}
{{ statement }}
{% endfor %}
Given the dialogue / passage above, use the following options to answer the
question "{{question}}".
Options:
- {{ answer_choices | join('\n- ') }}
```

Target Template:

```
{{ answer }}
```

Answer Choices Template:

```
{{ choice | join(" ||| ") }}
```

Input Template:

```
{% for statement in context %}
{{ statement }}
{% endfor %}
鉴于上面的对话/段落, 使用以下选项回答问题 "{{question}}".
选项:
- {{ answer_choices | join('
- ') }}
```

Target Template:

```
{{ answer }}
```

Answer Choices Template:

```
{{ choice | join(" ||| ") }}
```

Input Template:

```
Passage: {% for statement in context %}
{{ statement }}
{% endfor %}
Question: "{{question}}"
Answer choices: {{ answer_choices[:-1] | join(', ') }}, or {{ answer_choices[-1]
}}?
```

Target Template:

```
{{ answer }}
```

Answer Choices Template:

```
{{ choice | join(" ||| ") }}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
Passage: {% for statement in context %}
{{ statement }}
{% endfor %}
What kind of question would elicit an answer response of {{ answer }}?
```

Target Template:

```
{{ question }}
```

Input Template:

```
段落: {% for statement in context %}
{{ statement }}
{% endfor %}
问题: "{{question}}"
答案选择: {{ answer_choices[:-1] | join(', ') }} 还是 {{ answer_choices[-1] }}?
```

Target Template:

```
{{ answer }}
```

Answer Choices Template:

```
{{ choice | join(' ||| ') }}
```

Input Template:

```
{% for statement in context %}
{{ statement }}
{% endfor %}
Given the dialogue / passage above, what is the answer for the question
"{{question}}"
Answer choices: {{ answer_choices[:-1] | join(', ') }}, or {{ answer_choices[-1]
}}?
```

Target Template:

```
{{ answer }}
```

Answer Choices Template:

```
{{ choice | join(' ||| ') }}
```

Input Template:

```
{% for statement in context %}
{{ statement }}
{% endfor %}
鉴于上面的对话/段落，问题“{{question}}”的答案是什么
答案选择：{{ answer_choices[:-1] | join(', ') }} 还是 {{ answer_choices[-1] }}？
```

Target Template:

```
{{ answer }}
```

Answer Choices Template:

```
{{ choice | join(' ||| ') }}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
{% for statement in context %}
{{ statement }}
{% endfor %}
Given the dialogue / passage above, what is the answer for the question
"{{question}}"
```

Target Template:

```
{{ answer }}
```

1.7 QUESTION ANSWERING EXTRACTIVE

1.7.1 CLUE CMRC2018

Dataset from Cui et al. (2018). Used in training.

Data Example

Key	Value
id	TRAIN_186_QUERY_1
context	范廷颂枢机 (), 圣名保禄·若瑟 (), 是越南罗马天主教枢机。1963年...
question	1990年, 范廷颂担任什么职务?
answers	{'text': ['1990年被擢升为天主教河内总教区宗座署理'], ...}

Prompts

Input Template:

```
问: {{ question }}你能写一些上下文来回答这个问题吗?
```

Target Template:

```
{{ context }}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
Given this context "{{ context }}", generate a question that would return the answer of "{{ answers['text'][0] }}".
```

Target Template:

```
{{ question }}
```

Input Template:

```
{{ context }}  
{{ question }} 的答案在上面的段落中。它是什么?
```

Target Template:

```
{{ answers['text'][0] }}
```

Input Template:

In an exam, you are asked `{{ question }}`, and you are tasked to find the answer from the following passage.
`{{ context }}`
What's the answer?

Target Template:

```
{{ answers['text'][0] }}
```

Input Template:

```
{{ context }}  
The answer to {{ question }} is in the passage above. What is it?
```

Target Template:

```
{{ answers['text'][0] }}
```

Input Template:

```
Answer the question using the given context.  
Question: {{ question }}  
Context: {{ context }}  
Answer:
```

Target Template:

```
{{ answers['text'][0] }}
```

Input Template:

```
Q: {{ question }} Can you write some context to answer the question?
```

Target Template:

```
{{ context }}
```

Input Template:

```
{{ context[:answers["answer_start"][0]-5] }}... How would you continue the prior text to answer "{{ question }}"?
```

Target Template:

```
{{ context[answers["answer_start"][0]-5:] ]}}
```

Input Template:

```
{{ context[:answers["answer_start"][0]-5] ]}}... 你将如何继续前面的文本来回答 “{{ question ]}}” ?
```

Target Template:

```
{{ context[answers["answer_start"][0]-5:] ]}}
```

Input Template:

```
使用给定的上下文回答问题。  
问题: {{ question ]}}  
上下文: {{ context ]}}  
答案:
```

Target Template:

```
{{ answers['text'][0] ]}}
```

Input Template:

```
在考试中, 你被问到 {{ question ]}}, 你的任务是从以下段落中找到答案。  
{{ context ]}}  
答案是什么 ?
```

Target Template:

```
{{ answers['text'][0] ]}}
```

Input Template:

```
给定这个上下文 “{{ context ]}}”, 生成一个返回 “{{ answers['text'][0] ]}}” 答案的问题。
```

Target Template:

```
{{ question ]}}
```

1.7.2 CLUE DRCD

Dataset from Xu et al. (2020). Used in training.

Data Example

Key	Value
id	1001-10-2
context	2010年引進的廣州快速公交運輸系統，屬世界第二大快速公交系統，日常載...
question	從哪一天開始在廣州市區騎摩托車會被回收？
answers	{'text': ['2007年1月16日'], 'answer_st...

Prompts

Input Template:

```
{{ context }}
{{ question }} 的答案在上面的段落中。它是什么？
```

Target Template:

```
{{ answers['text'][0] }}
```

Input Template:

```
Answer the question using the given context.
Question: {{ question }}
Context: {{ context }}
Answer:
```

Target Template:

```
{{ answers['text'][0] }}
```

Input Template:

```
{{context[:answers["answer_start"]-5]}}... 你将如何继续前面的文本来回答 “{{
question}}”？
```

Target Template:

```
{{context[answers["answer_start"]-5:]}}
```

Input Template:

```
在考试中，你被问到 {{ question }}，你的任务是找到回答问题的段落。写这样一段话：
```

Target Template:

```
{{ context }}
```

Input Template:

```
{{ context }}  
The answer to {{ question }} is in the passage above. What is it?
```

Target Template:

```
{{ answers['text'][0] }}
```

Input Template:

```
在考试中，你被问到 {{ question }}，你的任务是从以下段落中找到答案。  
{{ context }}  
答案是什么？
```

Target Template:

```
{{ answers['text'][0] }}
```

Input Template:

```
给定这个上下文 “{{ context }}”，生成一个返回 “{{ answers['text'][0] }}” 答案的问题。
```

Target Template:

```
{{ question }}
```

Input Template:

```
{{context[:answers["answer_start"]-5]}}... How would you continue the prior text  
to answer "{{ question }}"?
```

Target Template:

```
{{context[answers["answer_start"]-5:]}}
```

Input Template:

```
使用给定的上下文回答问题。  
问题： {{ question }}  
上下文： {{ context }}  
答案：
```

Target Template:

```
{{ answers['text'][0] }}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
Given this context "{{ context }}", generate a question that would return the answer of "{{ answers['text'][0] }}".
```

Target Template:

```
{{ question }}
```

Input Template:

```
In an exam, you are asked {{ question }}, and you are tasked to find the answer from the following passage.  
{{ context }}  
What's the answer?
```

Target Template:

```
{{ answers['text'][0] }}
```

Input Template:

```
In an exam, you are asked {{ question }}, and you are tasked to find a passage answering the question. Write such a passage:
```

Target Template:

```
{{ context }}
```

1.7.3 MLQA MLQA.VI.VI

Dataset from Lewis et al. (2019). Used in training.

Data Example

Key	Value
context	Thành phố Miêu Lật tiếng Trung: 苗栗市,...
question	Miaoli có tỷ lệ cao loại người nào?
answers	{'answer_start': [311], 'text': ['K...]
id	2f0d6ff162619164bb113c0cadbcca06a50...

Prompts

Input Template:

```
{{ context[:answers.answer_start[0]-5]}} ... Tiếp tục ở trên, sao cho nó trả lời
"{{question}}":
```

Target Template:

```
{{ context[answers.answer_start[0]-5:]}}
```

Input Template:

```
{{context}}
```

Với sự tham chiếu đến ngữ cảnh trên, {{question}}

Target Template:

```
{{answers.text[0]}}
```

Input Template:

```
{{context}}
```

H: {{question}}

Đề cập đến đoạn văn trên, câu trả lời đúng cho câu hỏi đã cho trong ngôn ngữ của đoạn văn là

Target Template:

```
{{answers["text"][0]}}
```

Input Template:

Câu hỏi: {{question}}

Ngữ cảnh: {{context}}

Câu trả lời từ ngữ cảnh:

Target Template:

```
{{answers.text[0]}}
```

Input Template:

Tham khảo đoạn văn dưới đây và sau đó trả lời câu hỏi sau đó bằng ngôn ngữ tương tự như đoạn văn:

Đoạn: `{{context}}`

Câu hỏi: `{{question}}`

Target Template:

```
{{answers["text"][0]}}
```

Input Template:

Tôi đã tìm thấy một văn bản trả lời "`{{question}}`" bằng `{{answers.text[0]}}`. Nó bắt đầu bằng "`{{ context[:10] }}`". Bạn có thể tiếp tục nó không?

Target Template:

```
{{ context[10:] }}
```

Input Template:

Đọc đoạn văn sau và sau đó trả lời câu hỏi tiếp theo bằng cách trích một phần đúng trong đoạn văn:

`{{context}}`
`{{question}}`

Target Template:

```
{{answers.text[0]}}
```

Input Template:

D: `{{context}}`

H: `{{question}}`

A:

Target Template:

```
{{answers["text"][0]}}
```

1.7.4 MLQA MLQA.ZH.ZH

Data Example

Key	Value
context	楚河州包括有整个楚河河谷及邻近的山脉与峡谷。河谷的黑土非常肥沃，而且被...
question	哪水体有助土地如此多产？
answers	{'answer_start': [36], 'text': ['楚河...']}
id	1aee17dd937cc1043e3ff47c38396541fc3...

Prompts

Input Template:

```
阅读下面的短文，然后从短文中选出正确的部分来回答下面的问题：  
{{context}}  
{{question}}
```

Target Template:

```
{{answers.text[0]}}
```

Input Template:

```
{{ context[:answers.answer_start[0]-5]}}... 继续上述操作，使其回答 “{{question}}”:
```

Target Template:

```
{{ context[answers.answer_start[0]-5:]}}
```

Input Template:

```
D: {{context}}  
问: {{question}}  
答:
```

Target Template:

```
{{answers["text"][0]}}
```

Input Template:

我找到了一个用 `{{answers.text[0]}}` 回答 “`{{answers.text[0]}}`” 的文本。它以 “`{{context[:10]}}`” 开头。可以继续吗？

Target Template:

`{{ context[:10] }}`

Input Template:

问题: `{{question}}`
上下文: `{{context}}`
从上下文中回答:

Target Template:

`{{answers.text[0]}}`

Input Template:

参考下面的段落，然后用与段落相同的语言回答问题：

段落: `{{context}}`

问题: `{{question}}`

Target Template:

`{{answers["text"][0]}}`

Input Template:

`{{context}}`

参考上述上下文，`{{question}}`

Target Template:

`{{answers.text[0]}}`

Input Template:

```
{{context}}
```

```
问: {{question}}
```

```
参考上面的段落, 用该段落的语言对给定问题的正确答案是
```

Target Template:

```
{{answers["text"][0]}}
```

1.7.5 XQUAD XQUAD.VI

Dataset from Artetxe et al. (2019). Used in training.

Data Example

Key	Value
id	56beb4343aeaaa14008c925c
context	Đội thủ của Panthers chỉ thua 308 đ...
question	Jared Allen có bao nhiêu lần vật ng...
answers	{'text': ['136'], 'answer_start': [...

Prompts

Input Template:

```
{{context}}
```

```
Với sự tham chiếu đến ngữ cảnh trên, {{question}}
```

Target Template:

```
{{answers.text[0]}}
```

Input Template:

```
Đưa ra câu trả lời {{answers.text[0]}} cho {{question}}, hãy viết một văn bản giải thích điều này. Câu trả lời phải bắt đầu ở số ký tự {{answers.answer_start[0]}}.  
Văn bản:
```

Target Template:

```
{{context}}
```

Input Template:

`{{question}}` Rõ ràng là `{{answers.text[0]}}`. Bạn có thể cung cấp cho tôi một số bối cảnh?

Target Template:

`{{context}}`

Input Template:

`{{context}}`

H: `{{question}}`

Đề cập đến đoạn văn trên, câu trả lời chính xác cho câu hỏi được đưa ra là

Target Template:

`{{answers["text"][0]}}`

Input Template:

`{{context}}`

H: `{{question}}`

A:

Target Template:

`{{answers["text"][0]}}`

Input Template:

Đọc đoạn văn sau và trả lời câu hỏi sau:

`{{context}}`

`{{question}}`

Target Template:

`{{answers.text[0]}}`

Input Template:

Tham khảo đoạn văn dưới đây và trả lời câu hỏi sau:

Đoạn: `{{context}}`

Câu hỏi: `{{question}}`

Target Template:

```
{{answers["text"][0]}}
```

Input Template:

```
{{context}}
```

Từ đoạn văn trên, một câu hỏi hợp lý với "`{{answers["text"][0]}}`" như câu trả lời sẽ là:

Target Template:

```
{{question}}
```

Input Template:

```
{{context}}
```

Tạo câu hỏi từ đoạn văn trên:

Target Template:

```
{{question}}
```

1.7.6 XQUAD XQUAD.ZH

Data Example

Key	Value
id	56beb4343aeaaa14008c925c
context	黑豹队的防守只丢了 308分，在联赛中排名第六，同时也以 24 次拦截...
question	贾里德在职业生涯中有多少次擒杀？
answers	<code>{'text': ['136 次'], 'answer_start':...</code>

Prompts

Input Template:

阅读下面的短文，回答下面的问题：

`{{context}}`
`{{question}}`

Target Template:

`{{answers.text[0]}}`

Input Template:

`{{context}}`

问： `{{question}}`

参考上面的段落，给定问题的正确答案是

Target Template:

`{{answers["text"][0]}}`

Input Template:

参考下面的短文，回答下列问题：

段落： `{{context}}`

问题： `{{question}}`

Target Template:

`{{answers["text"][0]}}`

Input Template:

`{{context}}`

从上面的段落中，一个以 "`{{answers["text"][0]}}`" 为答案的合理问题将是：

Target Template:

`{{question}}`

Input Template:

`{{context}}`

从上面的段落中产生一个问题:

Target Template:

`{{question}}`

Input Template:

`{{context}}`

参考上述上下文, `{{question}}`

Target Template:

`{{answers.text[0]}}`

Input Template:

`{{context}}`

问: `{{question}}`

答:

Target Template:

`{{answers["text"][0]}}`

1.8 TOPIC CLASSIFICATION

1.8.1 CLUE CSL

Data Example

Key	Value
idx	1
corpus_id	2565
abst	针对核函数参数选择的重要性,提出了粒子群(PSO)模式搜索算法来搜索最...
label	-1
keyword	模式搜索;支持向量机;核参数选取

Prompts

Input Template:

```
After John wrote the abstract "{{abst}}", he wrote these keywords "{{ keyword | join(', ') }}" . Do you think his choice of keywords was correct? Answer {{ answer_choices[1] }} or {{ answer_choices[0] }}.
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
no ||| yes
```

Input Template:

```
Do these keywords "{{ keyword | join(', ') }}" represent key concepts in the abstract "{{ abst }}"?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
no ||| yes
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
Given the abstract {{abst}}, list out {{ keyword | length }} keywords for it.
```

Target Template:

```
{% if label == 1 %}  
{{ keyword | join(', ') }}  
{% endif %}
```

Input Template:

```
一位学者使用 "{{ keyword | join(', ') }}" 作为搜索词。你认为搜索引擎会返回摘要 "{{abst}}" 吗？回答 {{ answer_choices[1] }} 或 {{ answer_choices[0] }}。
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
不 ||| 是的
```

Input Template:

```
给定抽象{{abst}}, 列出{{ keyword | length }} 关键字。
```

Target Template:

```
{% if label == 1 %}  
{{ keyword | join(', ') }}  
{% endif %}
```

Input Template:

```
写一篇关于 “{{ keyword | join(', ') }}” 的摘要:
```

Target Template:

```
{% if label == 1 %} {{abst}} {% endif %}
```

Answer Choices Template:

```
不 ||| 是的
```

Input Template:

```
在约翰写完摘要 “{{abst}}” 之后, 他写了这些关键字 “{{ keyword | join(', ') }}”。你认为他选择的关键词是正确的吗? 回答 {{ answer_choices[1] }} 或 {{ answer_choices[0] }}。
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
不 ||| 是的
```

Input Template:

这些关键字 “`{{ keyword | join(', ') }}`” 是否代表抽象 “`{{ abst }}`” 中的关键概念？

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
不 ||| 是的
```

Input Template:

```
A scholar used "{{ keyword | join(', ') }}" as search terms. Do you think the search engine would return the abstract "{{abst}}"? Answer {{ answer_choices[1] }} or {{ answer_choices[0] }}.
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
no ||| yes
```

Input Template:

```
Write an abstract about "{{ keyword | join(', ') }}":
```

Target Template:

```
{% if label == 1 %} {{abst}} {% endif %}
```

Answer Choices Template:

```
no ||| yes
```

1.8.2 CLUE TNEWS

Data Example

Prompts

Input Template:

Key	Value
sentence	买套房不香吗？为什么会有人愿花600万买部手机？
label	-1
idx	1

将标题 “`{{ sentence }}`” 分为以下主题：
- `{{ answer_choices | join('\n- ') }}`
主题：

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
故事 ||| 文化 ||| 娱乐 ||| 运动的 ||| 金融 ||| 房地产 ||| 车 ||| 教育 ||| 技术 |||  
军队 ||| 旅行 ||| 世界新闻 ||| 股票 ||| 农业 ||| 游戏
```

Input Template:

```
Classify the title "{{ sentence }}" into the following topics:  
- {{ answer_choices | join('\n- ') }}  
Topic:
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
story ||| culture ||| entertainment ||| sports ||| finance ||| real estate ||| car  
||| education ||| tech ||| military ||| travel ||| world news ||| stock |||  
agriculture ||| game
```

Input Template:

```
Given the topics of {{answer_choices[:-1] | join(', ') }}, and {{  
answer_choices[-1] }}, specify which of them best represents the following  
sentence:  
{{ sentence }}  
Best:
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
story ||| culture ||| entertainment ||| sports ||| finance ||| real estate ||| car  
||| education ||| tech ||| military ||| travel ||| world news ||| stock |||  
agriculture ||| game
```

Input Template:

```
以下新闻标题 “{{ sentence }}" 属于什么主题？ {{ answer_choices[0] | capitalize }},  
{{ answer_choices[1:-1] | join(', ') }} 还是 {{ answer_choices[-1] }}?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
故事 ||| 文化 ||| 娱乐 ||| 运动的 ||| 金融 ||| 房地产 ||| 车 ||| 教育 ||| 技术 |||  
军队 ||| 旅行 ||| 世界新闻 ||| 股票 ||| 农业 ||| 游戏
```

Input Template:

```
鉴于 {{answer_choices[:-1] | join(', ') }} 和 {{ answer_choices[-1] }}, 指定它们中  
的哪一个最能代表以下句子:  
{{ sentence }}
```

最佳:

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
故事 ||| 文化 ||| 娱乐 ||| 运动的 ||| 金融 ||| 房地产 ||| 车 ||| 教育 ||| 技术 |||  
军队 ||| 旅行 ||| 世界新闻 ||| 股票 ||| 农业 ||| 游戏
```

Input Template:

```
What topic does the following news title "{{ sentence }}" belong to? {{  
answer_choices[0] | capitalize }}, {{ answer_choices[1:-1] | join(', ') }}, or {{  
answer_choices[-1] }}?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
story ||| culture ||| entertainment ||| sports ||| finance ||| real estate ||| car  
||| education ||| tech ||| military ||| travel ||| world news ||| stock |||  
agriculture ||| game
```

1.9 CODE MISC.

1.9.1 CODEPARROT/CODECOMPLEX CODEPARROT-CODECOMPLEX

Data Example

Key	Value
src	<pre>import java.util.Scanner; public ...</pre>
complexity	linear
problem	1197_B. Pillars
from	CODEFORCES

Prompts

Input Template:

```
{{ code }} What is the time complexity of the previous code?
```

Target Template:

```
{{ complexity }}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
Identify the time complexity of the following code as constant, linear, quadratic,  
cubic, log(n), nlog(n) or NP-hard. {{ code }} Complexity:
```

Target Template:

```
{{ complexity }}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
{{ code }} Which one is the correct time complexity of the code snippet: constant,  
linear, quadratic, cubic, log(n), nlog(n) or NP-hard?
```

Target Template:

```
{{ complexity }}
```

1.9.2 GREAT_CODE

Dataset from Hellendoorn et al. (2020). Used in training.

Data Example

Key	Value
id	1
source_tokens	#NEWLINE#;def test_get_params(;self...
has_bug	True
error_location	76
repair_candidates	2;76;4;11;18;22;30;40;48;58;66;80;8...
bug_kind	1
bug_kind_name	VARIABLE_MISUSE
repair_targets	4;11;18;22;30;40;48;58;66;80;88;103...
edges	[{'before_index': 1, 'after_index':...
provenances	{'datasetProvenance': {'datasetName...

Prompts

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
{% set mask = 'def <FUNC_NAME> (' %}
{% set indent = ' ' %}
{% set ns = namespace(indent_size=0, result=[], masked=false, target='') %}
{% for token in source_tokens %}
    {% if ns.masked is false and token.startswith('def') %}
        {% set ns.target = token.split('def ')[1][:-1] %}
        {% set token = mask %}
        {% set ns.masked = true %}
    {% endif %}
    {% if token == '#INDENT#' %}
        {% set ns.indent_size = ns.indent_size + 1 %}
        {% set ns.result = ns.result + [indent * ns.indent_size] %}
    {% elif token == '#NEWLINE#' %}
        {% set ns.result = ns.result + ["\n"] %}
    {% elif token == '#UNINDENT#' %}
        {% set ns.indent_size = ns.indent_size - 1 %}
    {% else %}
        {% if not loop.first and loop.previtem == '#NEWLINE#' %}
            {% set ns.result = ns.result + [indent * ns.indent_size] %}
        {% endif %}
        {% set ns.result = ns.result + [token | replace('\n', '\n'), " "] %}
    {% endif %}
{% endfor %}
{{ns.result | join("") | replace(" . ", ".") | replace(" , ", ", ") | replace("("
, "(") | replace(" )", ")") | replace("[ ", "[") | replace(" ]", "]"}}

What is the function name?
```

Target Template:

```
{{ ns.target }}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
{% set result = "" %}
{% set indent = ' ' %}
{% set ns = namespace(indent_size=0, line_number=0, buggy_line=0, bug_location=0,
bug_len=0, result=[], result_lines=[]) %}
{% set fixed_token = source_tokens[repair_targets[0]] %}
{% set buggy_line_content = "" %}
{% set fixed_buggy_line_content = "" %}

{% if has_bug and (repair_targets | length > 0) %}
  {% for token in source_tokens %}
    {% if loop.index0 == error_location %}
      {% set ns.buggy_line = ns.line_number %}
      {% set ns.bug_location = (ns.result | join("") | length) %}
      {% set ns.bug_len = (token | length) %}
    {% endif %}
    {% if token == '#INDENT#' %}
      {% set ns.indent_size = ns.indent_size + 1 %}
      {% set ns.result = ns.result + [indent * ns.indent_size] %}
    {% elif token == '#NEWLINE#' %}
      {% set ns.result_lines = ns.result_lines + [ns.result | join("")] %}
      {% set ns.result = [] %}
      {% set ns.line_number = ns.line_number + 1 %}
    {% elif token == '#UNINDENT#' %}
      {% set ns.indent_size = ns.indent_size - 1 %}
    {% else %}
      {% if not loop.first and loop.previtem == '#NEWLINE#' %}
        {% set ns.result = ns.result + [indent * ns.indent_size] %}
      {% endif %}
      {% set ns.result = ns.result + [token | replace('\n', '\n'), " "] %}
    %}
  {% endif %}
  {% endfor %}
  {% set ns.result_lines = ns.result_lines + [ns.result | join("")] %}
  {% set result = ns.result_lines | join("\n") %}
  {{result | replace(" . ", ". ") | replace(" , ", ", ") | replace("( ", "(") |
replace(" )", ")") | replace("[ ", "[") | replace(" ]", "]"}}}

  {% set buggy_line_content = ns.result_lines[ns.buggy_line] | trim | replace("
", ". ") | replace(" , ", ", ") | replace("( ", "(") | replace(" )", ")") |
replace("[ ", "[") | replace(" ]", "]" %}
  {% set fixed_buggy_line_content =
(ns.result_lines[ns.buggy_line][:ns.bug_location] + fixed_token +
ns.result_lines[ns.buggy_line][ns.bug_location + ns.bug_len:]) | trim | replace("
", ". ") | replace(" , ", ", ") | replace("( ", "(") | replace(" )", ")") |
replace("[ ", "[") | replace(" ]", "]" %}

  Fix the buggy line: {{buggy_line_content}}
```

Target Template:

```
{{fixed_buggy_line_content}}
{% endif %}
```

Input Template:

```
{% set mask = '<MASK>' %}
{% set indent = ' ' %}
{% set ns = namespace(indent_size=0, result=[]) %}

{% if has_bug %}
  {% for token in source_tokens %}
    {% if loop.index0 == error_location %}
      {% set token = mask %}
    {% endif %}
    {% if token == '#INDENT#' %}
      {% set ns.indent_size = ns.indent_size + 1 %}
      {% set ns.result = ns.result + [indent * ns.indent_size] %}
    {% elif token == '#NEWLINE#' %}
      {% set ns.result = ns.result + ["\n"] %}
    {% elif token == '#UNINDENT#' %}
      {% set ns.indent_size = ns.indent_size - 1 %}
    {% else %}
      {% if not loop.first and loop.previtem == '#NEWLINE#' %}
        {% set ns.result = ns.result + [indent * ns.indent_size] %}
      {% endif %}
      {% set ns.result = ns.result + [token | replace('\n', '\n'), " "] %}
    {% endif %}
  {% endfor %}
  {{ns.result | join("") | replace(" . ", ".") | replace(" , ", ", ") |
  replace("(" , "(") | replace(" )", ")") | replace("[ ", "[") | replace(" ]",
  "]"}}}

Given the code above, what is a proper replacement for {{mask}}?
```

Target Template:

```
{{source_tokens[repair_targets[0]]}}
{% endif %}
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```
{% set indent = ' ' %}
{% set ns = namespace(indent_size=0, result=[]) %}
{% for token in source_tokens %}
  {% if token == '#INDENT#' %}
    {% set ns.indent_size = ns.indent_size + 1 %}
    {% set ns.result = ns.result + [indent * ns.indent_size] %}
  {% elif token == '#NEWLINE#' %}
    {% set ns.result = ns.result + ["\n"] %}
  {% elif token == '#UNINDENT#' %}
    {% set ns.indent_size = ns.indent_size - 1 %}
  {% else %}
    {% if not loop.first and loop.previtem == '#NEWLINE#' %}
      {% set ns.result = ns.result + [indent * ns.indent_size] %}
    {% endif %}
  {% endif %}
{% endfor %}
```

```

        {% endif %}
        {% set ns.result = ns.result + [token | replace('\n', '\n'), " "] %}
    {% endif %}
{% endfor %}
{{ns.result | join("") | replace(" . ", ".") | replace(" , ", ", ") | replace("("
, "(") | replace(" )", ")") | replace("[ ", "[" | replace(" ]", "]"}}}

```

Is there a bug in the code above?

Target Template:

```

{{ {True: "Yes", False: "No"}[has_bug] }}

```

Answer Choices Template:

```

Yes ||| No

```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Input Template:

```

{% set mask = '<MASK>' %}
{% set indent = ' ' %}
{% set ns = namespace(indent_size=0, result=[]) %}

{% if has_bug %}
    {% for token in source_tokens %}
        {% if loop.index0 == error_location %}
            {% set token = mask %}
        {% endif %}
        {% if token == '#INDENT#' %}
            {% set ns.indent_size = ns.indent_size + 1 %}
            {% set ns.result = ns.result + [indent * ns.indent_size] %}
        {% elif token == '#NEWLINE#' %}
            {% set ns.result = ns.result + ["\n"] %}
        {% elif token == '#UNINDENT#' %}
            {% set ns.indent_size = ns.indent_size - 1 %}
        {% else %}
            {% if not loop.first and loop.previtem == '#NEWLINE#' %}
                {% set ns.result = ns.result + [indent * ns.indent_size] %}
            {% endif %}
            {% set ns.result = ns.result + [token | replace('\n', '\n'), " "] %}
        %}
    {% endif %}
{% endfor %}
{{ns.result | join("") | replace(" . ", ".") | replace(" , ", ", ") |
replace("(" , "(") | replace(" )", ")") | replace("[ ", "[" | replace(" ]",
"]")}}}

Given the code above, what is a proper replacement for {{mask}}? Choose among:
{{answer_choices | join(", ")}}

```

Target Template:

```

{{source_tokens[repair_targets[0]}}
{% endif %}

```

Answer Choices Template:

```
{% if has_bug %}      {% set nss = namespace(choices=[]) %}      {% for i in
repair_candidates %}      {% set nss.choices = nss.choices + [source_tokens[(i
| int)]] %}      {% endfor %}      {{nss.choices | unique | join(" ||| ")}} {% endif
%}
```

1.9.3 TEVEN/CODE_DOCSTRING_CORPUS TOP_LEVEL

Data Example

Key	Value
desc	'XXX22: This has to be present'
decl	def XXX11():
bodies	pass

Prompts

Input Template:

```
Complete the below
{{decl}}
'''{{desc | replace('
', '
')}}'''
```

Target Template:

```
{{bodies}}
```

Input Template:

```
I wrote the below code
{{bodies}}
What's a good function header?
```

Target Template:

```
{{decl}}
```

Input Template:

```
{{decl}}
```

Target Template:

```
"""{{desc | replace('
', '
 ') | replace('""', '')}}"""
{{bodies}}
```

1.10 WORD SENSE DISAMBIGUATION

1.10.1 PASINIT/XLWIC XLWIC_EN_ZH

Dataset from Raganato et al. (2020). Used in training.

Data Example

Key	Value
id	EN_1
context_1	We like to summer in the Mediterran...
context_2	We summered in Kashmir.
target_word	summer
pos	V
target_word_location_1	{'char_start': 11, 'char_end': 17}
target_word_location_2	{'char_start': 3, 'char_end': 11}
language	EN
label	1

Prompts

Input Template:

```
第 1 句: {{context_1}}
句子 2: {{context_2}}

确定单词 “{{target_word}}” 在两个句子中的使用是否相同。是还是不是？
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

```
不 ||| 是的
```

Input Template:

家庭作业

判断 “`{{target_word}}`” 这个词在以下两个句子中的含义是否相同。回答是或否。

`{{context_1}}`
`{{context_2}}`

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

不 ||| 是的

Input Template:

```
{{context_1}}
{{context_2}}
{{target_word}} 的类似意义？
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

不 ||| 是的

Input Template:

```
“{{target_word}}” 这个词有多种含义。第 1 句和第 2 句的意思相同吗？是还是不是？
第 1 句: {{context_1}}
句子 2: {{context_2}}
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

不 ||| 是的

Input Template:

确定以下两个句子中是否以相同的方式使用了单词 “`{{target_word}}`”。

```
{{context_1}}
{{context_2}}
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

不 ||| 是的

Input Template:

```
{{context_1}}
{{context_2}}
问题: “{{target_word}}” 这个词在上面两个句子中的含义是否相同?
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

不 ||| 是的

Input Template:

```
“{{target_word}}” 这个词在这两个句子中是否具有相同的含义? 是的, 不是吗?
{{context_1}}
{{context_2}}
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

不 ||| 是的

Input Template:

```
“{{target_word}}” 这个词在这两个句子中是否具有相同的含义？  
{{context_1}}  
{{context_2}}
```

Target Template:

```
{% if label != -1%}  
{{answer_choices[label]}}  
{% endif %}
```

Answer Choices Template:

不 ||| 是的

Input Template:

```
句子 A: {{context_1}}  
句子 B: {{context_2}}  
“{{target_word}}” 在句子 A 和 B 中具有相似的含义。对还是错？
```

Target Template:

```
{% if label != -1%}  
{{answer_choices[label]}}  
{% endif %}
```

Answer Choices Template:

错误的 ||| 真的

Input Template:

```
{{context_1}}  
{{context_2}}  
问题: “{{target_word}}” 这个词在上面两个句子中的含义是否相同？是的，不是吗？
```

Target Template:

```
{% if label != -1%}  
{{answer_choices[label]}}  
{% endif %}
```

Answer Choices Template:

```
不 ||| 是的
```

1.10.2 PASINIT/XLWIC XLWIC_FR_FR

Data Example

Key	Value
id	FR_1
context_1	L' éclaircie est généralement une co...
context_2	Améliorations utiles.
target_word	amélioration
pos	N
target_word_location_1	{'char_start': 41, 'char_end': 53}
target_word_location_2	{'char_start': 0, 'char_end': 13}
language	FR
label	1

Prompts

Input Template:

```
{{context_1}}  
{{context_2}}  
Sens similaire de {{target_word}} ?
```

Target Template:

```
{% if label != -1%}  
{{answer_choices[label]}}  
{% endif %}
```

Answer Choices Template:

```
Non ||| Oui
```

Input Template:

```
Devoirs  
  
Décidez si le mot "{{target_word}}" est utilisé avec le même sens dans les deux  
phrases suivantes. Répondez par oui ou non.  
{{context_1}}  
{{context_2}}
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

```
Non ||| Oui
```

Input Template:

```
Le mot "{{target_word}}" a-t-il le même sens dans ces deux phrases ? Oui Non?
{{context_1}}
{{context_2}}
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

```
Non ||| Oui
```

Input Template:

```
{{context_1}}
{{context_2}}
Question : Le mot '{{target_word}}' est-il utilisé dans le même sens dans les deux
phrases ci-dessus ? Oui Non?
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

```
Non ||| Oui
```

Input Template:

```
Le mot "{{target_word}}" a plusieurs significations. A-t-il le même sens dans les
phrases 1 et 2 ? Oui ou non?

Phrase 1 : {{context_1}}
Phrase 2 : {{context_2}}
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

```
Non ||| Oui
```

Input Template:

```
Phrase 1 : {{context_1}}
Phrase 2 : {{context_2}}

Déterminez si le mot "{{target_word}}" est utilisé dans le même sens dans les deux
phrases. Oui ou non?
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

```
Non ||| Oui
```

Input Template:

```
Le mot "{{target_word}}" a-t-il le même sens dans ces deux phrases ?
{{context_1}}
{{context_2}}
```

Target Template:

```
{% if label != -1%}
{{answer_choices[label]}}
{% endif %}
```

Answer Choices Template:

```
Non ||| Oui
```

Input Template:

Phrase A : `{{context_1}}`
Phrase B : `{{context_2}}`

"`{{target_word}}`" a une signification similaire dans les phrases A et B. Vrai ou faux ?

Target Template:

```
{% if label != -1%  
{{answer_choices[label]}}  
{% endif %}
```

Answer Choices Template:

Faux ||| Vrai

Input Template:

```
Déterminez si le mot '{{target_word}}' est utilisé de la même manière dans les  
deux phrases ci-dessous.  
{{context_1}}  
{{context_2}}
```

Target Template:

```
{% if label != -1%  
{{answer_choices[label]}}  
{% endif %}
```

Answer Choices Template:

Non ||| Oui

Input Template:

```
{{context_1}}  
{{context_2}}  
Question : Le mot '{{target_word}}' est-il utilisé dans le même sens dans les deux  
phrases ci-dessus ?
```

Target Template:

```
{% if label != -1%  
{{answer_choices[label]}}  
{% endif %}
```

Answer Choices Template:

Non ||| Oui

1.11 PARAPHRASE IDENTIFICATION

1.11.1 PAWS-X EN

Dataset from Yang et al. (2019). Used in training.

Data Example

Key	Value
id	2
sentence1	The NBA season of 1975 -- 76 was th...
sentence2	The 1975 -- 76 season of the Nation...
label	1

Prompts

Notes: Generalized prompt format, task_description-input.

Input Template:

```
Determine if the following two sentences paraphrase each other or not.  
Sent 1: {{sentence1}}  
Sent 2: {{sentence2}}
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Yes
```

Notes: Natural question.

Input Template:

```
Sentence 1: {{sentence1}}  
Sentence 2: {{sentence2}}  
Question: Do Sentence 1 and Sentence 2 express the same meaning? Yes or No?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

No ||| Yes

Notes: Generalized prompt format, context-question without any label.

Input Template:

```
{{sentence1}}  
Is that a paraphrase of the following sentence?  
{{sentence2}}?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

No ||| Yes

Notes: Natural Question without label.

Input Template:

```
Sentence 1: {{sentence1}}  
Sentence 2: {{sentence2}}  
Question: Can we rewrite Sentence 1 to Sentence 2?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

No ||| Yes

Notes: Generalized prompt format, context-question.

Input Template:

```
{{sentence1}}  
Is that a paraphrase of the following sentence?  
{{sentence2}}?  
Yes or No.
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Yes
```

Notes: Concatenation of sentence 1 and sentence 2.

Input Template:

```
Sentence 1: {{sentence1}}
Sentence 2: {{sentence2}}
Question: Does Sentence 1 paraphrase Sentence 2? Yes or No?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Yes
```

Note: the prompt does not correspond to the original task intended by the dataset authors.

Notes: Create a generative paraphrase task.

Input Template:

```
{% if label == 1 %}
Paraphrase the sentence: {{sentence1}}
```

Target Template:

```
{{sentence2}}
{% endif %}
```

Notes: Concatenation of sentence 1 and sentence 2 without any label.

Input Template:

```
Sentence 1: {{sentence1}}
Sentence 2: {{sentence2}}
Question: Does Sentence 1 paraphrase Sentence 2?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Yes
```

Notes: Natural question without label.

Input Template:

```
Sentence 1: {{sentence1}}  
Sentence 2: {{sentence2}}  
Question: Do Sentence 1 and Sentence 2 express the same meaning?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Yes
```

Prompt from Brown et al. (2020) **Notes:** ANLI prompt format from Table G7 in the GPT3 paper Brown et al. (2020)

Input Template:

```
{{sentence1}} Question: {{sentence2}} True or False?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
False ||| True
```

Notes: Natural Question.

Input Template:

```
Sentence 1: {{sentence1}}  
Sentence 2: {{sentence2}}  
Question: Can we rewrite Sentence 1 to Sentence 2? Yes or No?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Yes
```

Prompt from Brown et al. (2020) **Notes:** ANLI prompt format from Table G7 in the GPT3 paper Brown et al. (2020). Additionally added task information without any label.

Input Template:

```
{{sentence1}} Question: {{sentence2}} Paraphrase or not?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Yes
```

1.11.2 PAWS-X ES

Data Example

Key	Value
id	2
sentence1	La temporada de la NBA de 1975: 76 ...
sentence2	La temporada 1975 - 76 de la Asocia...
label	1

Prompts

Input Template:

```
Oración 1: {{sentence1}}  
Oración 2: {{sentence2}}  
Pregunta: ¿La oración 1 parafrasea la oración 2? ¿Si o no?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Sí
```

Input Template:

```
{{sentence1}} Pregunta: {{sentence2}} ¿Parafrasear o no?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Sí
```

Input Template:

```
{{sentence1}}  
¿Es una paráfrasis de la siguiente oración?  
{{sentence2}}?  
Sí o no.
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Sí
```

Input Template:

```
Oración 1: {{sentence1}}  
Oración 2: {{sentence2}}  
Pregunta: ¿La Oración 1 y la Oración 2 expresan el mismo significado?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Sí
```

Input Template:

```
{% if label == 1 %}  
Parafrasea la oración: {{sentence1}}
```

Target Template:

```
{{sentence2}}  
{% endif %}
```

Input Template:

```
{{sentence1}} Pregunta: {{sentence2}} ¿Verdadero o falso?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
Falso ||| Verdadero
```

Input Template:

```
Oración 1: {{sentence1}}  
Oración 2: {{sentence2}}  
Pregunta: ¿La oración 1 parafrasea la oración 2?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Sí
```

Input Template:

```
Determina si las siguientes dos oraciones se parafrasean entre sí o no.  
Enviado 1: {{sentence1}}  
Enviado 2: {{sentence2}}
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Sí
```

Input Template:

Oración 1: `{{sentence1}}`
Oración 2: `{{sentence2}}`
Pregunta: ¿La Oración 1 y la Oración 2 expresan el mismo significado? ¿Si o no?

Target Template:

`{{answer_choices[label]}}`

Answer Choices Template:

No ||| Sí

Input Template:

Oración 1: `{{sentence1}}`
Oración 2: `{{sentence2}}`
Pregunta: ¿Podemos reescribir la Oración 1 a la Oración 2? ¿Si o no?

Target Template:

`{{answer_choices[label]}}`

Answer Choices Template:

No ||| Sí

Input Template:

`{{sentence1}}`
¿Es una paráfrasis de la siguiente oración?
`{{sentence2}}`?

Target Template:

`{{answer_choices[label]}}`

Answer Choices Template:

No ||| Sí

Input Template:

Oración 1: `{{sentence1}}`
Oración 2: `{{sentence2}}`
Pregunta: ¿Podemos reescribir la Oración 1 a la Oración 2?

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

```
No ||| Sí
```

1.12 SENTENCE COMPLETION

1.12.1 XCOPA VI

Dataset from Ponti et al. (2020). Used in evaluation.

Data Example

Key	Value
premise	Cô gái tìm thấy con bọ trong ngũ cốc...
choice1	Cô đổ sữa vào bát.
choice2	Cô mất cảm giác ngon miệng.
question	effect
label	1
idx	1
changed	False

Prompts

Input Template:

```
{{ premise }}
```

```
Tôi đang lưỡng lự giữa hai lựa chọn. Giúp tôi chọn nguyên nhân {% if question ==  
"cause" %} có khả năng xảy ra cao hơn: {% else %} effect: {% endif %}  
- {{choice1}}  
- {{choice2}}
```

Target Template:

```
{% if label != -1 %} {{ answer_choices[label] }} {% endif %}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

Input Template:

```
{{ premise }}
```

Lựa chọn tốt nhất là gì?

- {{choice1}}
- {{choice2}}

Chúng tôi đang tìm kiếm {% if question == "cause" %} một nguyên nhân {% else %} một ảnh hưởng {% endif %}

Target Template:

```
{% if label != -1 %} {{answer_choices[label]}} {% endif %}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

Input Template:

```
{{ premise }} {% if question == "cause" %} Điều này xảy ra vì ... {% else %} Do đó ... {% endif %}  
Giúp tôi chọn tùy chọn hợp lý hơn:  
- {{choice1}}  
- {{choice2}}
```

Target Template:

```
{% if label != -1 %} {{ answer_choices[label] }} {% endif %}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

Input Template:

```
"{{ answer_choices[0] }}" hay "{{ answer_choices[1] }}"? {{ premise }} {% if question == "cause" %} bởi vì {% else %} nên {% endif %}
```

Target Template:

```
{% if label != -1 %} {{ answer_choices[label] }} {% endif %}
```

Answer Choices Template:

```
{{choice1 }} ||| {{choice2}}
```

Input Template:

```
{{ premise }}
```

```
Chọn nguyên nhân {% if question == "cause" %} hợp lý nhất: {% else %} effect: {% endif %}
- {{choice1}}
- {{choice2}}
```

Target Template:

```
{% if label != -1 %} {{ answer_choices[label] }} {% endif %}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

1.12.2 XCOPA ZH

Data Example

Key	Value
premise	这个女孩在麦片粥中发现了一个虫子。
choice1	她向碗里倒了牛奶。
choice2	她没了食欲。
question	effect
label	1
idx	1
changed	False

Prompts

Input Template:

```
{{ premise }} {% if question == "cause" %} 这是因为... {% else %} 结果... {% endif %}
帮助我选择更合理的选项:
- {{choice1}}
- {{choice2}}
```

Target Template:

```
{% if label != -1 %}{{ answer_choices[label] }}{%endif%}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

Input Template:

```
{{ premise }}
```

```
选择最合理的 {% if question == "cause" %} 原因: {% else %} 效果: {% endif %}
- {{choice1}}
- {{choice2}}
```

Target Template:

```
{% if label != -1 %}{{ answer_choices[label] }}{%endif%}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

Input Template:

```
"{{ answer_choices[0] }}" 还是 "{{ answer_choices[1] }}" ? {{ premise }} {% if
question == "cause" %} 因为 {% else %} 所以 {% endif %}
```

Target Template:

```
{% if label != -1 %}{{ answer_choices[label] }}{% endif %}
```

Answer Choices Template:

```
{{choice1 }} ||| {{choice2}}
```

Input Template:

```
{{ premise }}
```

```
最好的选择是什么？
- {{choice1}}
- {{choice2}}
```

```
我们正在寻找 {% if question == "cause" %} 一个原因 {% else %} 一个结果 {% endif %}
```

Target Template:

```
{% if label != -1 %}{{answer_choices[label]}}{%endif%}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

Input Template:

```
{{ premise }}
```

我在两个选项之间犹豫不决。帮我选择更有可能的 `{% if question == "cause" %}` 原因: `{% else %}` 效果: `{% endif %}`

- `{{choice1}}`
- `{{choice2}}`

Target Template:

```
{% if label != -1 %}{{ answer_choices[label] }}{%endif%}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

Input Template:

```
{{ premise }}
```

我正在考虑两个选项。请帮我最有可能的`{% if question == "cause" %}`导因: `{% else %}`后果:
`{% endif %}`

- `{{choice1}}`
- `{{choice2}}`

Target Template:

```
{% if label != -1 %}{{ answer_choices[label] }}{%endif%}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

Input Template:

```
{{ premise }} {% if question == "cause" %}这个会发生是因为... {% else %}结果是...  
{% endif %}  
帮我挑选合适的选项:  
- {{choice1}}  
- {{choice2}}
```

Target Template:

```
{% if label != -1 %}{{ answer_choices[label] }}{%endif%}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

Notes: Adapted from Perez et al. (2021) and Schick and Schütze (2020).

Input Template:

```
"{{ answer_choices[0] }}" 还是"{{ answer_choices[1] }}"? {{ premise }} {% if
question == "cause" %}因为{% else %}所以{% endif %}
```

Target Template:

```
{% if label != -1 %}{{ answer_choices[label] }}{% endif %}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

Input Template:

```
{{ premise }}
哪个是最好的答案?
- {{choice1}}
- {{choice2}}

我们正在考虑{% if question == "cause" %}起因{% else %}后果 {% endif %}
```

Target Template:

```
{% if label != -1 %}{{answer_choices[label]}}{%endif%}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

Input Template:

```
{{ premise }}
请选择最贴切的答案: {% if question == "cause" %}导因:{% else %}结果: {% endif %}
- {{choice1}}
- {{choice2}}
```

Target Template:

```
{% if label != -1 %}{{ answer_choices[label] }}{%endif%}
```

Answer Choices Template:

```
{{choice1}} ||| {{choice2}}
```

1.13 NATURAL LANGUAGE INFERENCE

1.13.1 XNLI EN

Dataset from Conneau et al. (2018). Used in evaluation.

Data Example

Key	Value
premise	you know during the season and i gu...
hypothesis	You lose the things to the followin...
label	0

Prompts

Notes: Sanh et al. (2022)

Input Template:

```
Take the following as truth: {{premise}}
Then the following statement: "{{hypothesis}}" is {"true"}, {"false"}, or
{"inconclusive"}?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
True ||| Inconclusive ||| False
```

Notes: Sanh et al. (2022)

Input Template:

```
{{premise}}
Question: Does this imply that "{{hypothesis}}"? Yes, no, or maybe?
```

Target Template:

```
{{answer_choices[label]}}
```

Answer Choices Template:

Yes ||| Maybe ||| No

Notes: Same as reported in Figure G7 of Brown et al. (2020), except that there is no task identifying tokens like "anli R1: ".

Input Template:

```
{{premise}}  
Question: {{hypothesis}} True, False, or Neither?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
True ||| Neither ||| False
```

Notes: Sanh et al. (2022)

Input Template:

```
Given that {{premise}} Does it follow that {{hypothesis}} Yes, no, or maybe?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Yes ||| Maybe ||| No
```

Notes: Adapted from the BoolQ prompts in Schick and Schütze (2020).

Input Template:

```
{{premise}} Based on the previous passage, is it true that "{{hypothesis}}"? Yes,  
no, or maybe?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Yes ||| Maybe ||| No
```

Notes: Webson and Pavlick (2021)

Input Template:

```
Given {{premise}} Is it guaranteed true that "{{hypothesis}}"? Yes, no, or maybe?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Yes ||| Maybe ||| No
```

Notes: Webson and Pavlick (2021)

Input Template:

```
Given {{premise}} Should we assume that "{{hypothesis}}" is true? Yes, no, or maybe?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Yes ||| Maybe ||| No
```

Notes: Sanh et al. (2022)

Input Template:

```
Given that {{premise}} Therefore, it must be true that "{{hypothesis}}"? Yes, no, or maybe?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Yes ||| Maybe ||| No
```

Notes: Webson and Pavlick (2021)

Input Template:

Suppose `{{premise}}` Can we infer that "`{{hypothesis}}`"? Yes, no, or maybe?

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Yes ||| Maybe ||| No
```

Notes: Webson and Pavlick (2021)

Input Template:

```
{{premise}} Are we justified in saying that "{{hypothesis}}"? Yes, no, or maybe?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Yes ||| Maybe ||| No
```

Notes: Sanh et al. (2022)

Input Template:

```
{{premise}} Based on that information, is the claim: "{{hypothesis}}" {{"true"}},  
{{"false"}}, or {{"inconclusive"}}?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
True ||| Inconclusive ||| False
```

Notes: Sanh et al. (2022)

Input Template:

```
{{premise}}  
Keeping in mind the above text, consider: {{hypothesis}} Is this {{"always"}},  
{{"sometimes"}}, or {{"never"}} correct?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Always ||| Sometimes ||| Never
```

Notes: Sanh et al. (2022)

Input Template:

```
Suppose it's true that {{premise}} Then, is "{{hypothesis}}" {{"always"}},  
{{"sometimes"}}, or {{"never"}} true?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Always ||| Sometimes ||| Never
```

Notes: Sanh et al. (2022)

Input Template:

```
Assume it is true that {{premise}}  
  
Therefore, "{{hypothesis}}" is {{"guaranteed"}}, {{"possible"}}, or  
{{"impossible"}}?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Guaranteed ||| Possible ||| Impossible
```

Notes: Adapted from Williams et al. (2018) instructions to crowdsourcing workers.

Input Template:

```
{{premise}} Using only the above description and what you know about the world,  
"{{hypothesis}}" is definitely correct, incorrect, or inconclusive?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Correct ||| Inconclusive ||| Incorrect
```

1.13.2 XNLI ES

Data Example

Key	Value
premise	Usted sabe durante la temporada y s...
hypothesis	Pierdes las cosas al siguiente nive...
label	0

Prompts

1.13.2.1 Human-translated prompts

Notes: Same as reported in Figure G7 of Brown et al. (2020), except that there is no task identifying tokens like "anli R1: ".

Input Template:

```
{{premise}}  
Pregunta: {{hypothesis}} Verdadero, Falso, o Ninguno?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Verdadero ||| Ninguno ||| Falso
```

Notes: Webson and Pavlick (2021)

Input Template:

```
Supongamos {{premise}} Podemos inferir que "{{hypothesis}}"? Si, no, o tal vez?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Sí ||| Tal vez ||| No
```

Notes: Webson and Pavlick (2021)

Input Template:

```
{{premise}} Estamos justificados en decir que "{{hypothesis}}"? Si, no, o tal vez?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Sí ||| Tal vez ||| No
```

Notes: Sanh et al. (2022)

Input Template:

```
Spongamos que es cierto que {{premise}}  
por lo tanto, "{{hypothesis}}" es {"garantizado"}, {"posible"}, o  
{"imposible"}?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Garantizado ||| Posible ||| Imposible
```

Notes: Adapted from Williams et al. (2018) instructions to crowdsourcing workers.

Input Template:

```
{{premise}} Usando solo la descripción anterior y lo que sabe sobre el mundo,  
"{{hypothesis}}" es definitivamente correcto, incorrecto o no concluyente?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Correcto ||| No concluyente ||| Incorrecto
```

1.13.2.2 Machine-translated prompts

Input Template:

```
{{premise}} ¿Estamos justificados al decir que &quot;{{hypothesis}}&quot;;? ¿Sí, no o tal vez?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Sí ||| Quizás ||| No
```

Input Template:

```
{{premise}} Pregunta: {{hypothesis}} ¿Verdadero, falso o ninguno?
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Verdadero ||| Ninguno de los dos ||| Falso
```

Input Template:

```
{{premise}} Usando solo la descripción anterior y lo que sabe sobre el mundo, &quot;{{hypothesis}}&quot; es definitivamente correcta, incorrecta o no concluyente.
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Correcto ||| Poco concluyente ||| Incorrecto
```

Input Template:

Supongamos `{{premise}}` ¿Podemos inferir que `"{{hypothesis}}"`? ¿Sí, no o tal vez?

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
Sí ||| Quizás ||| No
```

Input Template:

```
Supongamos que es cierto que {{premise}} Por lo tanto, &quot;{{hypothesis}}&quot;; es {{"guaranteed"}}, {{"possible"}} o {{"impossible"}}.
```

Target Template:

```
{{ answer_choices[label] }}
```

Answer Choices Template:

```
garantizado ||| Posible ||| Imposible
```

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Abstract & Section 1.
- A2. Did you discuss any potential risks of your work?
We did not find our work to have major potential risks.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract & Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 3 & Appendix D

- B1. Did you cite the creators of artifacts you used?
Section 3 & Appendix D
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Appendix D
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
The artifacts used did not detail their intended use except for license compliance, which we adhere to. The intended use of our artifacts is clear from their accompanying licenses.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
As we reuse existing datasets (ROOTS, XCOPIA, etc), we refer to the original works for such information.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
We explain the languages covered by our artifacts in Section 3.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Dataset sizes are reported in Figures 1 and 2, as well as the accompanying text.

C Did you run computational experiments?

Sections 3 & 4.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Model parameters and infrastructure details are reported in Section 3.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Hyperparameters are reported in Section 3.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Results are reported in Figures 4 - 8 and Tables 1 - 2 as well as in the Appendix.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

We report parameter settings in Section 3 and will provide code for reproduction of all our experiments.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Not applicable. Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Not applicable. Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Not applicable. Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Not applicable. Left blank.