

# Benchmarking and Improving Long-Text Translation with Large Language Models

Longyue Wang<sup>1\*</sup>, Zefeng Du<sup>1,2\*</sup>, Wenxiang Jiao<sup>1\*</sup>, Chenyang Lyu<sup>1</sup>, Jianhui Pang<sup>1,2</sup>,  
Leyang Cui<sup>1</sup>, Kaiqiang Song<sup>1</sup>, Derek F. Wong<sup>2</sup>, Shuming Shi<sup>1</sup>, Zhaopeng Tu<sup>1</sup>

<sup>1</sup>Tencent AI Lab, <sup>2</sup>University of Macau

vinnylywang@tencent.com, seele.du@connect.um.edu.mo

## Abstract

Recent studies have illuminated the promising capabilities of large language models (LLMs) in handling long texts. However, their performance in machine translation (MT) of long documents remains underexplored. This paper aims to shed light on how LLMs navigate this complex task, offering a comprehensive evaluation of their capabilities and limitations in long-text MT. First, we collect and construct an instruction-based benchmark dataset, specifically designed for the finetuning and evaluation of LLMs, encompassing multilingual, multi-domain, and document-level parallel data. Second, we conduct a comprehensive comparison between MT and LLM models concerning document-level translation. Our analysis uncovers that LLMs exhibit shortcomings in long-text domains, and their performance diminishes as document size escalates. By exploiting various extrapolation strategies, we enhance the capacity of LLMs to translate longer texts. We release data, code, and models at <https://github.com/longyuewangdcu/Document-MT-LLM>.

## 1 Introduction

Recently, significant progress has been made in the field of large language models (LLMs) such as ChatGPT (Ouyang et al., 2022; OpenAI, 2023). These models have demonstrated remarkable capabilities in processing extensive texts, opening up new avenues for a variety of natural language processing (NLP) tasks (Li et al., 2023). Specifically, research on long-text LLMs primarily focuses on two aspects: 1) *Maximum Capacity*: how long text a model can handle (Chen et al., 2023b; Ding et al., 2023); 2) *Efficient Utilization*: how efficiently that model can model long texts (Liu et al., 2023). Even though LLMs can process long texts, it does not

necessarily imply they can effectively model such texts. The effectiveness and necessity of extrapolating LLMs for various NLP tasks remain ambiguous. In this paper, we predominantly delve into the use of LLMs for long-text machine translation (MT).

Diverging from other long-text tasks such as text summarization and completion, which involve either long input or output, long-text MT presents unique challenges. It necessitates not only comprehending extensive input but also preserving context and coherence over extended text sequences (Yang et al., 2019, 2020). This amalgamation of requirements intensifies the complexity of the task, positioning long-text MT as a rigorous benchmark for evaluating the potential and prowess of LLMs.

To facilitate this study, we first construct an instruction-based benchmark dataset designed for finetuning and evaluating the document-level translation capabilities LLMs. This dataset comprises a total of 410K documents with a document length 160~54K tokens,<sup>1</sup> spanning across seven domains (i.e. news, TED, subtitles, Parliamentary, social, Q&A, novels) and three language pairs (i.e. Chinese-English, German-English, Russian-English). Leveraging this dataset, we conduct a comprehensive comparison between MT and LLM models (e.g. DeepL, GPT-4, Llama). The experiments reveal that 1) all models experience a performance decline in the novels domain due to its exceedingly long nature; 2) the performance of LLMs drastically deteriorates as the document length increases from 2K to 8K. To further investigate the impact of enlarging document length on translation quality, we examine a variety of extrapolation strategies during finetuning stage. Our findings indicate that a mixture of document lengths with multi-task learning can effectively enhance LLMs' performance. The **main contributions** are:

- To facilitate future research on long-text transla-

\* Longyue Wang, Zefeng Du and Wenxiang Jiao contribute equally. This work was completed during the internships of Zefeng Du, Jianhui Pang, and Chenyang Lyu at Tencent AI Lab.

<sup>1</sup>Here it means subword, referring a basic processing unit after sentencepiece.

ID	Instruction	Instruction-Based Example
1	Give the TGT translation for these SRC sentences: <b>D</b>	<p>Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request.</p> <p><b>### Instruction:</b> Give the English translation for these Chinese sentences:</p> <p><b>### Input:</b> 狂风吹过峡谷，在草原上掀起了阵阵波澜。正在俯首啃食着青嫩草叶的野牛抬起头来，望向天空。</p> <p><b>### Response:</b> As the wind howled through the valley, ripples spread across the grassland. The wild ox, with its head still buried in the green leaves of grass, raised its head and looked up at the sky.</p> <p>Each instance consists of 4 parts: system prompt, instruction, input and response.</p>
2	Please provide the TGT translation for these sentences: <b>D</b>	
3	Provide the TGT equivalent for the following SRC sentences: <b>D</b>	
4	Translate the following SRC sentences into TGT: <b>D</b>	
5	Translate the following sentences from SRC to TGT: <b>D</b>	
6	Convert the following sentences from SRC to TGT: <b>D</b>	
7	Render the following SRC sentences into TGT: <b>D</b>	
8	What is the TGT version of these SRC sentences? <b>D</b>	
9	What do the following sentences mean in TGT? <b>D</b>	
10	What are the translations of these SRC sentences in TGT? <b>D</b>	
11	How do these SRC sentences translate to TGT? <b>D</b>	
12	Give the TGT translation for this SRC <b>C</b> : <b>D</b>	
13	Please provide the TGT translation for the <b>C</b> : <b>D</b>	
14	Provide the TGT equivalent for the following SRC <b>C</b> : <b>D</b>	
15	Translate this <b>C</b> from SRC to TGT: <b>D</b>	
16	Translate the following SRC <b>C</b> into TGT: <b>D</b>	
17	Convert the following <b>C</b> from SRC to TGT: <b>D</b>	
18	Render the following SRC <b>C</b> into TGT: <b>D</b>	
19	What is the TGT version of this SRC <b>C</b> ? <b>D</b>	
20	What do the following <b>C</b> mean in TGT? <b>D</b>	
21	What is the translation of this SRC <b>C</b> in TGT? <b>D</b>	
22	How do this SRC <b>C</b> translate to TGT? <b>D</b>	

Table 1: The instructions used for document-level translation. SRC and TGT denote source and target languages, respectively. **D** represents the source sentences. In #12~22 prompts, we randomly select one word from the candidates  $C \in [\text{“document”, “text”, “paragraph”, “passage”}]$ .


tion, we release a large instruction-based dataset. This dataset is comprehensive, covering various domains and languages, and notably includes exceptionally long texts (Section 3).

- We establish a long-text MT benchmark by evaluating advanced MT and LLM models. This not only highlights the limitations of current LLMs but also provides insights for improving their performance in translating long texts (Section 4).
- We investigate simple finetuning strategies, and identify effective methods to enhance the LLMs’ proficiency in translating long texts (Section 5).

## 2 Preliminary

**Background** Recent years have seen a growing interest in document-level translation by leveraging traditional neural models (Wang et al., 2017; Voita et al., 2019a; Wang et al., 2018), pre-training (Liu et al., 2020) as well as LLMs (Wang et al., 2023a; Wu et al., 2024a). Surprisingly, LLMs have shown potentials in modeling exceptionally long texts with complex discourse structures for document-level MT (Wang et al., 2023b; Pang et al., 2024a; Lyu et al., 2024). More recently, there have been sev-

eral length extrapolation methods aimed at efficiently extending the pre-trained context length of LLMs (Chen et al., 2023a,b). However, they were rarely tested on challenging generation tasks such as long-text translation. Our preliminary experiments on document translation using simple extrapolation methods (Pang et al., 2024b; Chen et al., 2023b) show a significant performance degradation on WMT2023 English-German testset (Llama-2 vs. +Extrapolation: 39.1 vs. 35.6 d-BLEU), which is consistent with the findings in multi-document QA task (Liu et al., 2024). Thus, this work mainly investigate finetuning strategies.



**Pseudo-Document** Given a full length of document-level parallel corpus, previous works usually concatenate  $N$  consecutive sentences into a pseudo-document (Wang et al., 2017; Voita et al., 2019b; Wu et al., 2023) until it matches either the

	Domain	Source	Language	D	S	T	T / D
Training	News	WMT08~21 News Devsets	Zh-En	1.0K	12.8K	876.2K/47.6K	648
			De-En	1.5K	32.6K	1.4M/997.2K	795
			Ru-En	1.0K	21.8K	986.3K/645.0K	820
		WMT22 News Commentary	Zh-En	8.0K	313.7K	19.5M/10.0M	1,845
			De-En	9.8K	388.5K	17.9M/12.3M	1,544
			Ru-En	8.8K	331.5K	18.3M/11.0M	1,671
	TED	IWSLT2015 IWSLT2017 IWSLT2014	Zh-En	1.7K	213.4K	9.3M/5.1M	4,181
			De-En	1.7K	212.9K	6.4M/5.0M	3,348
			Ru-En	1.4K	181.0K	6.2M/4.3M	3,729
	Subtitle	OpenSubtitles 2018	Zh-En	12.0K	11.2M	189.2M/109.5M	12,456
			De-En	29.2K	22.5M	251.4M/218.5M	8,057
			Ru-En	34.5K	25.9M	338.3M/249.7M	8,511
	Parliamentary	United Nations v1.0 Europarl v7	Zh-En	91.0K	15.9M	983.5M/564.0M	8,501
			Ru-En	133.0K	23.2M	1.3B/798.4M	7,990
			De-En	11.3K	1.9M	87.7M/62.1M	6,633
Novels	Par3	Zh-En	3.7K	23.5K	906.7K/534.8K	196	
		De-En	8.9K	35.4K	1.2M/938.8K	120	
		Ru-En	34.0K	104.7K	3.2M/2.4M	82	
		<b>Ours</b>	Zh-En	22.6K	1.9M	79.6M/48.8M	2,759
Testing	News Social	WMT2022	Zh⇒En	38	505	41.2K/26.4K	889
				25	478	32.6K/19.7K	1,046
	Novels Q&A	mZPRT	Zh⇒En	12	857	40.7K/23.6K	2,679
				182	1,171	34.6K/24.7K	163
	TED	IWSLT2015 IWSLT2017	Zh⇒En	62	6,047	257.9K/136.6K	3,181
			En⇒De	23	2,271	51.9K/64.4K	2,527
	News Parliamentary	WMT2022 Europarl v7	En⇒De	31	511	16.9K/23.4K	511
				360	5,134	167.9K/236.8K	562
	Subtitle	<b>Ours</b>	En⇒Ru	13	22,602	193.8K/232.2K	16.4K
Novels-XL	<b>Ours</b>	Zh⇒En	12	16,742	824.4K/474.4K	54.1K	

Table 2: Statistics of our benchmark dataset. We count the number of documents |D|, sentences |S|, tokens |T| in terms of source/target language, and the average length of a document |T|/|D| (K thousand and M million).

*maximum length limitation* or the *actual document length*. We follow this method to preprocess all types of data in our paper.

**Evaluation Methods** The d-BLEU (document-level BLEU) metric (Liu et al., 2020) has emerged as the de facto standard for evaluating document-level translation models such as MCN (Zheng et al., 2021), G-Trans (Bao et al., 2021), and MR-Doc2Doc (Sun et al., 2022). Besides, Wang et al. (2023a) demonstrates that d-BLEU are more attuned to discourse than sentence-level BLEU. In this work, we utilize the commonly-used d-BLEU as the default instead of sentence-level metrics since all translations are at the document level. As a supplement, we also report results using human evaluation (Wang et al., 2023a) and the BlonDe automatic metric (Jiang et al., 2022) in Section 5.4

and the Appendix §A.3, respectively. Note that, evaluating document-level translation is another challenging issue, which potentially beyond the scope of this work.

### 3 Instruction-Based Benchmark Dataset

According to Taori et al. (2023), we construct instruction-based datasets for long-text translation training and testing, spanning various domains and languages. In addition to leveraging existing parallel datasets, we create two challenging testsets.

#### 3.1 Overview and Statistics

Following Wang et al. (2023a), we ask GPT-4 to suggest a variety of instructions and manually refine them for the long-text translation task (as enumerated in Table 1). We then generate

instruction-based instances by integrating the original document-level sentence pairs with these instructions. As seen, each instance consists of four parts: 1) *System Prompt*, a predefined text to set the context or frame for the model’s response (fixed length); 2) *Instruction*, describing the specific task (around 8 words); 3) *Input*, sentences in the source language; and 4) *Response*, the corresponding translation in the target language. During instruction data construction, we combine  $N$  continuous sentences as a pseudo-document, of which source and target sides are respectively added into the `###Input` and `###Response` to form one instance of instruction-based MT data.

Table 2 outlines the languages, domains, and statistics of our dataset. We choose three language pairs: Chinese-English (Zh-En), German-English (De-En) and Russian-English (Ru-En). We select texts from seven domains, of which five are consistent across both training and testing (i.e. News, TED, Subtitle, Parliamentar and Novels), while the remaining two are exclusive to the testing set (i.e. Social and Q&A):

- *News*. We used WMT2008~2021 News Devsets and WMT2022 News Commentary for training, and WMT2022 News Testset for testing.<sup>2</sup>
- *TED*. We collected TED talks in the IWSLT2015, 2017 and 2014 corpora for training, and the standard testsets for testing.<sup>3</sup>
- *Subtitle*. We collected movie subtitles in the OpenSubtitles2018 corpus for training.<sup>4</sup> As the existing Ru-En testset contains limited length of context (Voita et al., 2019a), we employed professional translators to create a new one.
- *Parliamentary*. We employed the United Nations v1.0 and Europarl v7 training and testing sets.<sup>5</sup>
- *Novels*. The Par3 corpus is organized at the paragraph level.<sup>6</sup> Note that, Par3 does not reflect the natural length since the publicly available data has already been processed into pseudo-documents (without original data). In addition, we release a copyrighted corpus: 23K continuous chapters from 152 web fictions, covering 14 genres such as fantasy science and romance. The texts are originally written in Chinese and were

translated into English by professional translators. Regarding testsets, we utilize both mZPRT and our specially curated long-text. More details are discussed in Appendix §A.2.

- *Q&A and Social*. We used subsets in Baidu-Knows and website domains from WMT2022 and mZPRT<sup>7</sup> corpora, respectively.

The average document length ( $|T|/|D|$ ) can be considered as an indicator of the complexity inherent in a document-level dataset. Extended documents introduce a two-fold challenge: they not only test the model’s capacity to process extensive sequences but also necessitate the understanding and generation of nuanced discourse properties across extended contexts. As seen, the document lengths of our constructed testsets are 16K and 54K, currently making them the longest document-level translation test sets available.

### 3.2 Analysis and Discussion

**Data Distribution** Figure 1 provides a visualization of the benchmark datasets, illustrating the natural distribution of data size and document length across various domains. As seen, 1) the data scale across different domains is imbalanced. For instance, there is an abundance of data in the Parliamentary, while spoken language domains such as TED are scarce; 2) the natural length of documents varies across different domains. For example, Novels and Subtitles tend to be quite lengthy, whereas News are typically shorter; 3) The distribution between the training and testing sets is not exactly identical, which mirrors real-world challenges.

**Data Contamination** There exists a risk of data contamination as publicly available data are often leveraged during different stages of LLM training (e.g., pre-training, supervised fine-tuning (SFT), or reinforcement learning from human feedback (RLHF)). To mitigate this risk, we have compiled three new datasets for training and testing models:

- *Novels Zh-En training set*. We crawl monolingual web novels and manually conduct cross-lingual alignment. The collected Par3 dataset is organized at the paragraph level, while ours is more extensive, organized at the chapter level.
- *Novels Zh-En testing set*. We collect latest Chinese web novels and then employ professional translators to translate them into English. We

<sup>2</sup><https://www.statmt.org/wmt22>.

<sup>3</sup><https://wit3.fbk.eu>.

<sup>4</sup><https://opus.nlpl.eu/OpenSubtitles-v2018.php>.

<sup>5</sup><https://conferences.unite.un.org/uncorpus> and <https://www.statmt.org/europarl>.

<sup>6</sup><https://github.com/katherinethai/par3>.

<sup>7</sup><https://github.com/longyuewangdcu/mZPRT>.




Figure 1: A visual representation of our dataset (left: training, right: testing). The x-axis denotes the average document length on a logarithmic scale, while the y-axis means the total number of documents in a dataset (we also adjust the circle sizes to intuitively showcase the data scale). Different colors correspond to distinct domains. Datasets from the same domain but in the same language pairs are combined.

combine 20 continuous chapters into one document, resulting in a super-long document.

- *Subtitle Ru-En testing set.* We collect latest movie subtitles in English and engage professional translators to translate them into Russian.

## 4 Benchmark Results

We provide a comprehensive view of how MT and LLM systems handle the complexities of long-text translation across various languages and domains.

### 4.1 Experimental Setup

**Models** We utilize *testing sets* in the benchmark dataset. We conduct a systematic comparison of:

- *Commercial MT Systems.* The commercial translation products always perform best in translation performance. However, they typically preprocess documents into individual sentences and then perform translation on a sentence-by-sentence basis. We mainly compare Google Translate,<sup>8</sup> DeepL Translate,<sup>9</sup> and Tencent Translate,<sup>10</sup> as adopted in (Jiao et al., 2023b).
- *Commercial LLM Systems.* ChatGPT are known for their extensive context modeling capabilities and can achieve state-of-the-art results across various languages and tasks. Between August 1st

and 30th, 2023, we used GPT-3.5 (4K) and GPT-4 (8K) APIs to translate documents.<sup>11</sup> In our preliminary experiments, we employed GPT-4 (128K) but observed frequent instances of omitted translations.

- *Open-sourced LLM Models.* Llama and Llama-2-chat are widely-used, open-source LLMs, pre-trained and fine-tuned with a maximum length of 4,096 tokens (Touvron et al., 2023a,b). This enables to handle a relatively longer context window size, making it suitable for document-level general tasks. About Llama-mt, we enhance Llama on document-level MT by using the training set of our benchmark for SFT. We randomly select 5K instances from each subset in Table 2, resulting in approximately 200K instances.

**Pre-processing Pseudo-Document** In testing commercial and open-source LLM systems, we concatenate  $N$  consecutive sentences into a pseudo-document (Wang et al., 2017; Voita et al., 2019b; Wu et al., 2023) until it matches either the *maximum LLM input* or the *actual document length*. Assuming a 1:1.5 input to output ratio, we can calculate the maximum size of the *maximum LLM input* based on the total length  $\in \{2K, 4K, 8K\}$ . We transform these pseudo-document parallel data into instruction formats as Jiao et al. (2023a).

<sup>8</sup><https://translate.google.com>.

<sup>9</sup><https://www.deepl.com>.

<sup>10</sup><https://transmart.qq.com>.

<sup>11</sup>API: gpt-3.5-turbo-0125 and gpt-4-1106-preview.

Model	#	Zh⇒En						En⇒De			En⇒Ru	Avg.
		News	Soc.	Nov.	Q&A	TED	Nov. <sub>XL</sub>	TED	News	Par.	Sub.	
Google	<i>sent</i>	27.7	35.4	16.0	12.0	28.1	26.6	32.6	51.9	32.2	29.1	29.2
DeepL	<i>sent</i>	30.3	33.4	16.1	11.9	29.1	23.3	33.5	48.8	34.7	27.7	28.9
Tencent	<i>sent</i>	29.3	38.8	20.7	15.0	30.4	31.4	32.3	48.8	31.9	24.3	30.3
GPT-3.5	4K	31.6	36.1	19.4	19.5	27.3	28.2	33.5	49.2	30.9	29.1	30.5
GPT-4	8K	32.6	41.9	19.8	21.0	29.0	23.9	33.8	50.8	31.9	28.9	31.4
Llama-2-chat	4K	9.3	7.4	4.4	9.7	10.8	10.3	8.0	18.7	11.0	4.4	9.4
Llama-mt	2K	29.1	30.2	18.4	7.3	28.7	26.3	32.5	46.9	31.8	27.8	27.9
Llama-mt	4K	29.0	26.2	16.8	5.6	27.9	23.0	32.7	46.6	31.9	25.8	26.6

Table 3: The comparison of MT and LLM systems on our benchmark *testing* set, spanning various domains and languages. We report performances using the case-sensitive d-BLEU metric. We evaluated the GPT-3.5-4K and GPT-4-8K in August 2023. The Llama-2-chat is a 7B model, which is pre-trained and fine-tuned with 4,096 max length. The Llama-mt are Llama 7B models, which are pre-trained with 2,048 max length, and finetuned on *training* set with pseudo-document length of 2,048 and 4,096. The numbers in color represent the best results within each category, while the highlighted background indicates the best performance across all models.

### Finetuning LLMs for Document Translation

We finetune Llama-7B for 3 epochs with a batch size of 128. Our experiments are conducted on 16 NVIDIA A100 GPUs with DeepSpeed ZeRO stage 3 facilitating model parallelism.<sup>12</sup> To conserve GPU memory, we integrate Flash Attention mechanism (Dao et al., 2022). The SFT task is akin to the casual language modeling task, where the loss is computed solely based on the *output* portion of the input sequence.

**Evaluation** According to Wang et al. (2023a), the document-level sacreBLEU (d-BLEU) (Liu et al., 2020) are consistent with performances of discourse-specific phenomena (e.g. consistency of terminology translation and accuracy zero pronoun translation). We employ case-sensitive d-BLEU as our metric, which is computed by matching n-grams in the whole document.

## 4.2 Main Results

Table 3 presents a comprehensive comparison of various MT and LLM models on *benchmark testing set* across different domains and languages. When evaluating individual testsets, we notice discernible differences between models. Taking Zh-En testsets for example, Tencent Translate excels in Novels, while GPT-4 is more adept with web-based content like Social and Q&A. Concerning language pairs, Google and DeepL models show proficiency in tasks targeting non-English languages, whereas GPT-4 predominates in tasks with English as the

target. These variances can largely be attributed to differences in training data. Consequently, we will draw our conclusions holistically rather than relying solely on individual test sets.

On average across all testsets, GPT-4 outperforms all other models (i.e. 31.36 BLEU), particularly excelling in the Zh-En News, Social, Q&A, and En-De TED. The performance of GPT-3.5 is comparable to that of sentence-level MT systems, such as Tencent Translate (e.g. GPT-3.5 30.29 vs. Tencent Translate 30.48 BLEU). In contrast, Llama-2-chat registers a significantly lower score of 9.4 BLEU, attributable to its limited MT data during SFT and coverage of other languages (e.g. Russian). Addressing the need for long-text MT, we developed Llama-mt models by finetuning Llama using our benchmark training data. This adaptation enhances the translation quality of Llama by almost 200% (Llama-mt-2K 27.9 vs. Llama-2-chat 9.4 BLEU), placing it on par with commercial sentence-level MT systems (Llama-mt-2K 27.9 vs. DeepL 28.9 BLEU). Upon closer comparison between Llama-mt 2K and 4K, we observed a significant drop in performance, particularly on the Novel-XL testset, with a decrease of -3.3 BLEU. We cannot enhance the long-text translation capability of Llama by simply finetuning and decoding it with longer texts. This highlights the complexities and challenges inherent in attempting to bolster LLMs’ proficiency in long-text translation.

When comparing the performance of ChatGPT reported in Wang et al. (2023a), we notice a general upward trend across the overlapping testsets.

<sup>12</sup><https://github.com/microsoft/DeepSpeed>.




Figure 2: The effects of translation quality on document length. We draw *line chart* by extending the pseudo-document length during decoding on Novels-XL.

While this might indicate performance enhancements over time with model evolution, it’s crucial to acknowledge the different experimental parameters: 1) Their evaluations were conducted in March 2023, whereas ours took place in August 2023. 2) They accessed ChatGPT via its web interface, whereas we used its API. 3) Their methodology involved pseudo-documents of predetermined lengths (4 or 10 sentences), taking into account multi-turn contexts. Conversely, we employed the maximum lengths allowed by the API (2K or 4K) and disregarded multi-turn contexts.

### 4.3 Effects of Document Lengths

We further investigate the impacts of directly extending document length at decoding phase on the  $\text{Nov}_{\text{XL}}$  dataset (more analysis are discussed in Appendix §A.1). For comparison, we select two representative systems: DeepL and Llama-2-chat with 4K length of texts in pre-trained and SFT. Figure 2 shows the translation performance as the document length increases  $\in \{256, 1K, 2K, 4K, 8K\}$  during the decoding phase. While the performance of sentence-level DeepL remains unaffected by decoding length, Llama-2-chat paints a different picture. It initially achieves its peak d-BLEU score at a document length of 1K but sees a sharp decline from 2K to 8K. This underscores that even though LLMs can process long texts, it does not necessarily imply they can effectively model such texts. This motivates us to exploit extrapolation methods to enhance the ability of LLMs to translate texts that extend beyond their pretraining length. Another interesting observation is that the MT performance of Llama-2-chat does not reach its zenith at 4K, diverging from the lengths used in pre-training and

Finetune	Decode	News	Novels
SENT	sent	24.7	16.6
	doc (256)	23.8	15.9
	doc (512)	24.4	14.1
DOC (2048)	sent	27.5	16.1
	doc (1024)	<b>29.1</b>	18.3
	doc (2048)	<b>29.1</b>	<b>18.4</b>
	doc (4096)	22.6	9.2
DOC (4096)	sent	27.5	17.2
	doc (2048)	29.0	17.7
	doc (4096)	29.0	16.8
	doc (8192)	23.4	6.6

Table 4: The performance of document-level translation using different supervised finetuning (Finetune) and decoding (Decode) strategies in terms of d-BLEU scores. We conducted experiments on Llama 7B model on Chinese $\Rightarrow$ English News and Novels datasets. The “DOC/doc ( $n$ )” denotes the documents with the size of context windows  $n$  used in finetuning/decoding while “sent” means the sentence level strategies.

SFT. We suspect that this may be caused by various training signals at SFT phase. To focus specifically on the MT task, without being influenced by signals from other tasks, we primarily employ the Llama-mt settings for further experiments (in Section 5).

## 5 Context Extrapolation on Translation

In practice, it is unrealistic to keep increasing the context window of LLMs in the training process due to resource limits and unsatisfied performance (Chen et al., 2023b). Thus, it is interesting to investigate the potential of LLMs in translating documents that are noticeably longer than the pre-defined context window. In this section, we focus on curating the training data of long texts and explore various strategies for context extrapolation.

We employ Llama (7B) as testbed, which is pre-trained and fine-tuned with a maximum length of 2,048 token (Touvron et al., 2023a). Following the Llama-mt configurations detailed in Section 4.1, we pre-process each training subset into pseudo-documents with defined length. Then we randomly select 5K instances from each training subset (consider language directions), culminating  $\sim 200K$  instances for SFT phase. About evaluation, we utilize Chinese $\Rightarrow$ English WMT2022 News and mZPRT Novels as representative datasets and report d-BLEU score. We have examined different methods for utilizing data to improve the effective-

Scaling		Setting		News	Novels
FT	DEC	CO	MDS		
2048	2048	X	X	29.1	18.4
		✓	X	26.9	16.8
4096	4096	X	X	29.0	16.8
		✓	X	24.3	17.8
		X	✓	<b>29.5</b>	18.1
		✓	✓	26.8	17.7
4096	4608	X	X	29.3	17.3
		X	✓	29.0	18.3
		✓	✓	26.6	17.1
		X	X	29.1	16.0
		X	✓	28.9	<b>18.7</b>
	5120	✓	✓	26.2	17.0

Table 5: Effects of content overlapping (CO) and mixture of document lengths (MDL) on longer-text translation (2048→4096+) in terms of d-BLEU scores. The scaling document sizes is listed in Table 4, which is represented as finetuning (FT) and decoding (DEC).

ness of modeling text at the document level. We maintain the overall bilingual token count used during instruction finetuning. Note that three methods in Section 5.1, 5.2, and 5.3 build upon each other.

### 5.1 Effects of Scaling Document Length

First of all, we present how LLMs perform when directly scaling the document size for training and decoding. For comparison, we also include the performance of the sentence-to-sentence approach like typical MT systems. At the SFT phase (FT), we employ two context windows (2048 and 4096), which are used to split the full documents with sliding block techniques. In the decoding phase (DEC), we assess LLMs on a variety of document lengths: shorter than the context window, matching the context window, and exceeding it. For example, when finetuning LLMs with the context window of 2048, we test on documents with context windows of 1024, 2048, and 4096, respectively. Here, 2048 corresponds to the context window employed during Llama’s pre-training and finetuning, while 4096 represents an extrapolated context window.

Table 4 lists the results. Compared with the sentence-by-sentence approach (SENT-sent), increasing the context window during training can lead to much better translation results, which corroborates the findings of paragraph-level translation demonstrated in Karpinska and Iyyer (2023). Both approaches highlight the importance of

context-aware translation techniques in achieving better results for long texts. Besides, we have two specific observations concerning the context windows: (1) With a context window of 2048 for SFT, the test performance on documents with the context window of 4096 degrades significantly, indicating the poor ability of LLMs in context extrapolation. (2) When conducting SFT with a context window of 4096, which is larger than that used in llama pre-training, we have not observed further improvements over the context window of 2048. Specifically, both DOC(4096)-doc(2048) and DOC(4096)-doc(4096) underperform DOC(2048)-doc(2048), especially on *Novels*, which indicates the ineffectiveness of merely increasing context window. Therefore, it becomes crucial to explore other methods for context extrapolation.

### 5.2 Effects of Content Overlapping and Mixture of Document Lengths

As an attempt, we explore the compositional generalization ability of LLMs in learning from short documents and generalizing to long documents. The default method for constructing training data has two problems that may limit the generalization ability of LLMs: (1) The full documents are split into blocks with no overlap; (2) The length of input documents is fixed rather than following the length distribution in practice. For the first problems, we propose shift block to allow the content overlap between adjacent document blocks (CO, Content Overlapping). With a sliding window size of  $N$ , the overlap is  $N/2$ , indicating that the window advances by  $N/2$  sentences to extract the next document. For the second, we propose to mix document blocks with varied sizes (MDS, Mixture of Document Sizes). For example, if the base  $N$  is 2048, we randomly select one of the numbers from the set 256, 512, 1024, 2048 (all predefined values smaller than 2048 are considered) to serve as  $N$  for extracting the current pseudo-document. The process is then repeated in subsequent cycles.

We present the results in Table 5. Unexpectedly, content overlapping leads to worse performance which may result from the less unique tokens during finetuning. In contrast, the model trained with the mixture of document sizes achieves noticeable improvement, compared to that with a fixed document size. The improvement is especially enlarged when we adopt a context window exceeding the one used in training, e.g., 5120 vs. 4096. It demon-



Scaling		+TE	News	Novels
DOC	doc			
4096	4096	X	29.5	18.1
		✓	<b>29.9</b>	<b>19.2</b>
4096	4608	X	29.0	18.3
		✓	29.7	19.1
	5120	X	28.9	18.7
		✓	26.9	17.4

Table 6: Effects of multi-task finetuning (TE, text completion) on translating longer texts (2048→4096+) in terms of d-BLEU scores.

strates that learning on various sizes of documents indeed improves the generalization of LLMs to longer documents.

### 5.3 Effects of Multi-task Finetuning

Unlike sentence-level translation, document-level translation tasks suffer from the scarcity issue of bilingual parallel corpora severely, which may limit the performance of long-text finetuning. We exploit the advantages of monolingual novel data and combine it with our document-level translation corpora in the form of text completion (TE). Similar to Section 3, we construct an instruction-based dataset for TE by combining instructions (in Table 7) with Chinese and English monolingual Novels (different from parallel data in Table 2). We employ the data split by MDS (in Section 5.2) and the scale of the data is equivalent to that of the Chinese-English document translation corpus.

As shown in Table 6, incorporating the monolingual text completion tasks can improve the performance considerably. However, it can also degrade the translation performance slightly when extrapolating the context, making more hallucinations due to the relatively free generation style of text completion. Nonetheless, monolingual text completion tasks can improve the modeling ability of Llama models in non-English languages, and also ameliorate the language imbalance issue of English-centric parallel data. We will explore this direction more extensively in the future.

### 5.4 Results of Human Evaluation

We randomly selected a paragraph of Novels subset (around 50 sentences) from three system outputs: Table 4 DOC (2048)-doc (2048), Table 5 DOC (4096)-doc (4096) +MDS, Table 6 DOC (4096)-doc (4096) +TE. We follow Wang et al. (2023a)’s

ID	Instruction
1	Given the story in the input, please provide a continuation in the output. <b>D</b>
2	Based on the narrative provided, kindly offer a follow-up in the resulting text. <b>D</b>
3	Taking into account the story presented, please supply a subsequent part in the response. <b>D</b>
4	Considering the tale in the input, please deliver an extension in the output. <b>D</b>
5	With the story mentioned, please generate a progression in the resulting content. <b>D</b>
6	In light of the story shared, kindly produce a sequel in the output. <b>D</b>
7	Given the plot in the input, please create a further development in the response. <b>D</b>
8	Keeping the story in mind, please contribute a continuation in the output. <b>D</b>
9	Reflecting on the narrative given, please furnish an additional segment in the response. <b>D</b>
10	Acknowledging the story in the input, please extend it in the output. <b>D</b>
11	Upon reviewing the story provided, kindly add a follow-up in the resulting text. <b>D</b>

Table 7: The prompts suggested by ChatGPT for story continuation. Chinese continuation and English continuation tasks use the same instruction. **D** represents the story that needs to be continued.

general guideline to conduct a human evaluation. The results are 1.6, 2.1, 2.8, respectively (d-BLEU scores are 16.6, 18.1, 19.2), which confirms our findings based on d-BLEU scores. However, it is still an open question whether human and automatic evaluation metrics are complementary or mutually exclusive in measuring the document-level translation quality. We will explore more human evaluation methods in the future.

## 6 Conclusion

We provided valuable insights into the performance of LLMs in long-text translation. By creating a comprehensive, multi-domain, and multilingual document-level corpus for training and testing, we have been able to analyze the impact of various extrapolation methods on translation performance and propose a simple yet effective approach. By releasing data, code, and models, we aim to significantly promote further research in this field, ultimately improving the performance of LLMs and MT systems in handling long translation tasks.

## Limitations

We list the main limitations of this work as follows:

1. *Extrapolation Methods*: This study focuses on enhancing long-text translation through data manipulation and training strategies. Nonetheless, a wide array of alternative approaches merit exploration for boosting model performance. Delving into these methodologies may reveal more efficacious means of bolstering the robustness and precision of predictive models.
2. *Evaluation Methods*: In this paper, we mainly use the commonly-used automatic metric. The d-BLEU allows for a more nuanced assessment of translation quality, reflecting the unique challenges and requirements of long-text translation. However, it is still an open question whether human and automatic evaluation metrics are complementary or mutually exclusive in measuring the document-level translation quality (Wang et al., 2020, 2023a; Wu et al., 2024b).

## Ethics Statement

We take ethical considerations very seriously, and strictly adhere to the ACL Ethics Policy. Most datasets used in this paper are publicly available. Besides, we are the copyright owners of the newly proposed dataset. We ensure that the findings and conclusions of this paper are reported accurately and objectively.

## Acknowledgements

This work was supported in part by the Science and Technology Development Fund, Macau SAR (Grant Nos. FDCT/060/2022/AFJ, FDCT/0070/2022/AMJ), National Natural Science Foundation of China (Grant No. 62261160648), Ministry of Science and Technology of China (Grant No. 2022YFE0204900), the Multi-year Research Grant from the University of Macau (Grant No. MYRG-GRG2023-00006-FST-UMDF), and the Tencent AI Lab Rhino-Bird Gift Fund (Grant No. EF2023-00151-FST).

## References

Guangsheng Bao, Yue Zhang, Zhiyang Teng, Boxing Chen, and Weihua Luo. 2021. G-transformer for document-level machine translation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3442–3455.

Guangzheng Chen, Xin Li, Zaiqiao Meng, Shangsong Liang, and Lidong Bing. 2023a. Clx: Continuous length extrapolation for large language models. *arXiv preprint arXiv:2310.16450*.

Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. 2023b. Extending context window of large language models via positional interpolation. *arXiv preprint arXiv:2306.15595*.

Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. FlashAttention: Fast and memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing Systems*.

Jiayu Ding, Shuming Ma, Li Dong, Xingxing Zhang, Shaohan Huang, Wenhui Wang, and Furu Wei. 2023. Longnet: Scaling transformers to 1,000,000,000 tokens. *arXiv preprint arXiv:2307.02486*.

Yuchen Jiang, Tianyu Liu, Shuming Ma, Dongdong Zhang, Jian Yang, Haoyang Huang, Rico Sennrich, Ryan Cotterell, Mrinmaya Sachan, and Ming Zhou. 2022. BlonDe: An automatic evaluation metric for document-level machine translation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1550–1565, Seattle, United States. Association for Computational Linguistics.

Wenxiang Jiao, Jen-tse Huang, Wenxuan Wang, Xing Wang, Shuming Shi, and Zhaopeng Tu. 2023a. Parrot: Translating during chat using large language models. *arXiv preprint arXiv:2304.02426*.

Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. 2023b. Is chatgpt a good translator? a preliminary study. *arXiv preprint arXiv:2301.08745*.

Marzena Karpinska and Mohit Iyyer. 2023. Large language models effectively leverage document-level context for literary translation, but critical errors persist. *arXiv preprint arXiv:2304.03245*.

Dacheng Li, Rulin Shao, Anze Xie, Ying Sheng, Lianmin Zheng, Joseph E. Gonzalez, Ion Stoica, Xuezhe Ma, and Hao Zhang. 2023. How long can open-source llms truly promise on context length?

Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. Lost in the middle: How language models use long contexts. *arXiv preprint arXiv:2307.03172*.

Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12.

- 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.
- Chenyang Lyu, Zefeng Du, Jitao Xu, Yitao Duan, Minghao Wu, Teresa Lynn, Alham Fikri Aji, Derek F Wong, and Longyue Wang. 2024. A paradigm shift: The future of machine translation lies with large language models. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1339–1352.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Jianhui Pang, Fanghua Ye, Longyue Wang, Dian Yu, Derek F Wong, Shuming Shi, and Zhaopeng Tu. 2024a. Salute the classic: Revisiting challenges of machine translation in the age of large language models. *arXiv preprint arXiv:2401.08350*.
- Jianhui Pang, Fanghua Ye, Derek F Wong, and Longyue Wang. 2024b. Anchor-based large language models. *arXiv preprint arXiv:2402.07616*.
- Zwei Sun, Mingxuan Wang, Hao Zhou, Chengqi Zhao, Shujian Huang, Jiajun Chen, and Lei Li. 2022. Rethinking document-level neural machine translation. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3537–3548.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models](#).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019a. When a good translation is wrong in context: Context-aware machine translation improves on deixis, ellipsis, and lexical cohesion. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1198–1212.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019b. [When a good translation is wrong in context: Context-aware machine translation improves on deixis, ellipsis, and lexical cohesion](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Longyue Wang, Chenyang Lyu, Tianbo Ji, Zhirui Zhang, Dian Yu, Shuming Shi, and Zhaopeng Tu. 2023a. Document-level machine translation with large language models. *arXiv preprint arXiv:2304.02210*.
- Longyue Wang, Zhaopeng Tu, Yan Gu, Siyou Liu, Dian Yu, Qingsong Ma, Chenyang Lyu, Liting Zhou, Chao-Hong Liu, Yufeng Ma, et al. 2023b. Findings of the wmt 2023 shared task on discourse-level literary translation: A fresh orb in the cosmos of llms. In *Proceedings of the Eighth Conference on Machine Translation*, pages 55–67.
- Longyue Wang, Zhaopeng Tu, Shuming Shi, Tong Zhang, Yvette Graham, and Qun Liu. 2018. [Translating pro-drop languages with reconstruction models](#). In *Proceedings of the 2018 AAAI Conference on Artificial Intelligence*, volume 32, pages 4937–4945.
- Longyue Wang, Zhaopeng Tu, Xing Wang, Li Ding, Liang Ding, and Shuming Shi. 2020. Tencent ai lab machine translation systems for wmt20 chat translation task. In *Proceedings of the fifth conference on machine translation*, pages 483–491.
- Longyue Wang, Zhaopeng Tu, Andy Way, and Qun Liu. 2017. Exploiting cross-sentence context for neural machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2826–2831.
- Minghao Wu, George Foster, Lizhen Qu, and Gholamreza Haffari. 2023. Document flattening: Beyond concatenating context for document-level neural machine translation. *arXiv preprint arXiv:2302.08079*.
- Minghao Wu, Thuy-Trang Vu, Lizhen Qu, George Foster, and Gholamreza Haffari. 2024a. Adapting large language models for document-level machine translation. *arXiv preprint arXiv:2401.06468*.
- Minghao Wu, Yulin Yuan, Gholamreza Haffari, and Longyue Wang. 2024b. (perhaps) beyond human translation: Harnessing multi-agent collaboration for translating ultra-long literary texts. *arXiv preprint arXiv:2405.11804*.
- Baosong Yang, Longyue Wang, Derek F Wong, Lidia S Chao, and Zhaopeng Tu. 2019. Convolutional self-attention networks. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4040–4045.
- Yilin Yang, Longyue Wang, Shuming Shi, Prasad Tadepalli, Stefan Lee, and Zhaopeng Tu. 2020. On the sub-layer functionalities of transformer decoder. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4799–4811.

Zaixiang Zheng, Xiang Yue, Shujian Huang, Jiajun Chen, and Alexandra Birch. 2021. Towards making the most of context in neural machine translation. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 3983–3989.

## A Appendix

### A.1 Effects of Different Domains and Languages

As shown in Figure 3, the Social and Q&A domains exhibit a more informal structure and length, leading to varied performance across models. Besides, most models consistently excel in the News and Parliamentary domains. However, they underperform in the Novel and Subtitle domains, which are characterized by its extended context. This underscores the challenges LLMs face in long-text translation and motivates our pursuit to further explore translations of longer texts. Concerning languages, most models exhibit consistent performance: they shine in En-De translation while showing relatively subdued results in Zh-En and En-Ru. Therefore, we will delve deeper into the impact of document lengths using the Zh-En Novels dataset.

### A.2 About Novels Dataset

The novels were translated from one language to another by professional translators, with each novel chapter serving as the translation unit. Due to the absence of sentence-level alignment information—having only chapter-level data—we employed a method that combines automatic alignment with manual proofreading to achieve high-quality sentence-by-sentence alignment. For more details, please check GuoFeng Webnovel<sup>13</sup> and mZPRT<sup>14</sup>.

### A.3 Results of BlonDe Metric

We report results of different methods in Section 5 using the BlonDe automatic metric (Jiang et al., 2022) in Table 8, Table 9, Table 10. BlonDe takes discourse coherence into consideration by categorizing discourse-related spans and calculating the similarity-based F1 measure of categorized spans. There is a certain degree of discrepancy between d-BLEU and BlonDe metrics, which potentially provide complementary reference points when measuring the quality of document-level translations.

<sup>13</sup><https://github.com/longyuewangdcu/GuoFeng-Webnovel>.

<sup>14</sup><https://github.com/longyuewangdcu/mZPRT>.

Finetune	Decode	News	Novels
SENT	sent	41.35	38.52
	doc (256)	37.27	17.88
	doc (512)	<b>42.04</b>	<b>38.70</b>
DOC (2048)	sent	37.07	31.05
	doc (1024)	37.91	34.89
	doc (2048)	40.49	34.79
	doc (4096)	37.48	33.92
DOC (4096)	sent	35.88	26.66
	doc (2048)	39.71	36.18
	doc (4096)	41.53	34.54
	doc (8192)	40.86	36.68

Table 8: The performance of methods introduced in Section 5.1 in terms of BlonDe scores. The caption is similar to Table 4.

Scaling		Setting		News	Novels
FT	DEC	CO	MDS		
2048	2048	✗	✗	38.89	29.67
		✓	✗	38.67	28.17
4096	4096	✗	✗	40.50	34.74
		✓	✗	38.13	30.59
		✗	✓	41.53	34.54
		✓	✓	<b>42.04</b>	<b>38.70</b>
4096	4608	✗	✗	37.85	36.44
		✗	✓	37.65	35.74
	✓	✓	41.88	36.56	
	✗	✗	38.23	29.91	
	5120	✗	✓	41.67	35.46
		✓	✓	40.39	34.78

Table 9: The performance of methods introduced in Section 5.2 in terms of BlonDe scores. The caption is similar to Table 5.

Scaling		+TE	News	Novels
DOC	doc			
4096	4096	✗	41.67	35.46
		✓	37.58	33.89
4096	4608	✗	<b>41.88</b>	<b>36.56</b>
		✓	41.53	34.54
	5120	✗	38.23	34.04
✓		39.08	34.10	

Table 10: The performance of methods introduced in Section 5.3 in terms of BlonDe scores. The caption is similar to Table 6.

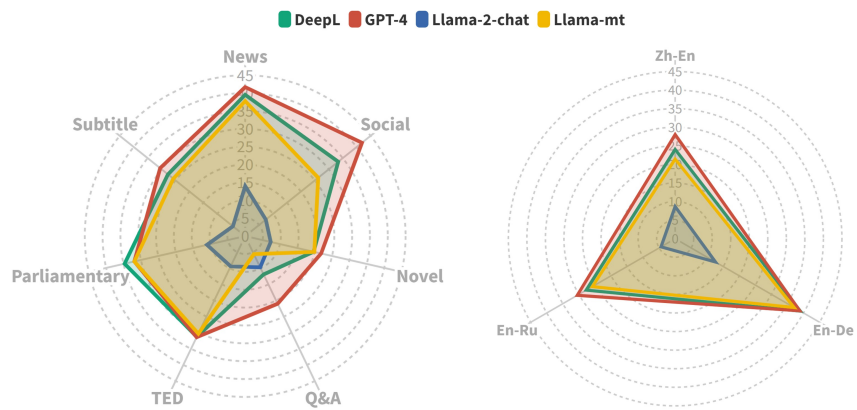


Figure 3: The effects of translation performance on domain and language. We draw *radar chart* by combining results in Table 3.