

# Adapting Large Language Models for Document-Level Machine Translation

Anonymous ACL submission

## Abstract

Large language models (LLMs) have significantly advanced various natural language processing (NLP) tasks. Recent research indicates that moderately-sized LLMs often outperform larger ones after task-specific fine-tuning. This study focuses on adapting LLMs for document-level machine translation (DOCMT) for specific language pairs. We first investigate the impact of prompt strategies on translation performance and then conduct extensive experiments using two fine-tuning methods, three LLM backbones, and 18 translation tasks across nine language pairs. Our results show that specialized models can sometimes surpass GPT-4 in translation performance but still face issues like *off-target translation* due to error propagation in decoding. We provide an in-depth analysis of these LLMs tailored for DOCMT, examining translation errors, discourse phenomena, strategies for training and inference, the data efficiency of parallel documents, recent test set evaluations, and zero-shot crosslingual transfer. Our findings highlight the strengths and limitations of LLM-based DOCMT models and provide a foundation for future research.

## 1 Introduction

Large language models (LLMs) demonstrate impressive proficiency in a wide range of applications (Ouyang et al., 2022; Wei et al., 2022a; Sanh et al., 2022; Chung et al., 2022; OpenAI, 2023; Anil et al., 2023; Touvron et al., 2023a,b; Jiang et al., 2023). However, in the realm of translation tasks, only few very large models, such as GPT-3.5-TURBO and GPT-4-TURBO, can match or surpass the performance of state-of-the-art supervised encoder-decoder models like NLLB (Costa-jussà et al., 2022), while they still under-perform in translating low-resource languages (Robinson et al., 2023; Jiao et al., 2023; Hendy et al., 2023). Consequently, a number of recent works attempt to bridge the gap between LLMs and supervised encoder-decoder

models in translation tasks (Zhu et al., 2023; Yang et al., 2023; Zhang et al., 2023; Moslem et al., 2023; Xu et al., 2023; Kudugunta et al., 2023). Recently, research suggests that smaller, specialized models can outperform larger, general-purpose models in specific tasks (Gunasekar et al., 2023; Luo et al., 2023; Azerbayev et al., 2023). Therefore, we explore adapting LLMs for document-level machine translation (DOCMT) in this study.

In this study, we explore the capabilities of moderately-sized Large Language Models (LLMs) with 7 billion parameters across 18 translation tasks involving nine language pairs. We conduct experiments under both *bilingual* and *multilingual* settings, employing Parameter-Efficient Fine-Tuning (PEFT) and Fully Fine-Tuning (FFT) techniques on three LLM backbones. Our findings reveal that, after fine-tuning, these LLMs demonstrate superior translation performance compared to state-of-the-art models, as evidenced by metrics such as *sBLEU*, *dBLEU*, and COMET. However, a significant challenge identified is the issue of *off-target translations*, which persists even after exclusive fine-tuning on bilingual corpora. Our analysis attributes this high rate of off-target translations to error propagation during the inference. Additionally, we provide a comprehensive analysis of our LLM-based DOCMT models from various perspectives. This includes examining translation error distribution, discourse phenomena, training strategies, the data efficiency of parallel documents, and additional evaluations on the recent WMT2023 test sets. We also investigate zero-shot cross-lingual transfer, aiming to enhance the understanding and efficacy of LLMs in DOCMT tasks.

We present extensive empirical evidence that highlights both the superior translation capabilities and limitations of the LLM-based DOCMT models in this study, making several significant discoveries. Here are the main takeaways:

- **Selective Excellence in Translation Tasks:**

Our findings show that our moderately-sized LLMs outperform GPT-4-TURBO in certain translation tasks, but struggle in others due to the *off-target translation* issue due to error propagation in decoding. Despite this, our DOCMT models exhibit better context awareness and fewer errors, while maintaining comparable performance.

- **Fine-Tuning Strategies:** Our research indicates that the PEFT approach outperforms the FFT approach overall. However, the FFT approach shows greater data efficiency, needing only about 1% of the total dataset to reach the performance level of models trained on the entire dataset. In contrast, the PEFT approach requires 10% of the total dataset for comparable results.
- **Evaluation on Recent Test Sets:** We evaluate our models on recent test sets between English and German from WMT2023 (Koehn et al., 2023). Our empirical results show that, when the data leakage risks are mitigated, the LLM-based DOCMT models generalize better on out-of-domain text, compared to the conventional DOCMT models.
- **Advantage of Base LLMs for Task-Specific Supervised Fine-Tuning:** Our study shows that base LLMs, when used as the backbone for task-specific supervised fine-tuning, perform better than instruction-tuned LLMs. They demonstrate more effective zero-shot cross-lingual transfer.

## 2 Related Work

**Document-Level Machine Translation** In recent years, numerous approaches have been proposed for document-level machine translation (DOCMT). There exist other approaches to DOCMT, including document embedding (Macé and Servan, 2019; Huo et al., 2020), multiple encoders (Wang et al., 2017; Bawden et al., 2018; Voita et al., 2018; Zhang et al., 2018), attention variations (Miculicich et al., 2018; Zhang et al., 2020; Maruf et al., 2019; Wong et al., 2020; Wu et al., 2023), and translation caches (Maruf and Haffari, 2018; Tu et al., 2018; Feng et al., 2022). Furthermore, Maruf et al. (2022) present a comprehensive survey of DOCMT.

**Large Language Models** Large language models (LLMs) have demonstrated remarkable proficiency across a wide range of Natural Lan-

guage Processing (NLP) tasks (Brown et al., 2020; Chowdhery et al., 2022; Scao et al., 2022; Anil et al., 2023; Touvron et al., 2023a,b). Furthermore, recent research has shown that supervised fine-tuning (SFT) and Reinforcement Learning from Human Feedback (RLHF) can significantly enhance their performance when following general language instructions (Mishra et al., 2022; Wang et al., 2022; Shen et al., 2023; Li et al., 2023; Wu and Aji, 2023; Wu et al., 2024a). More recently, there is a growing body of work exploring the translation capabilities of LLMs (Lu et al., 2023; Zhang et al., 2023; Xu et al., 2023; Robinson et al., 2023; Wu et al., 2024d). However, it is important to note that these efforts have primarily focused on sentence-level machine translation (SENMT) and have not delved into document-level machine translation (DOCMT). A noteworthy study in DOCMT is conducted by Wang et al. (2023b), where they investigate the document-level translation capabilities of GPT-3.5-TURBO, making it the most closely related work to our work.

**Ours** In contrast to the work of Wang et al. (2023b), who primarily investigate the use of GPT-3.5-TURBO for DOCMT through prompting techniques, our study concentrates on analyzing the effectiveness of parameter-efficient fine-tuning (PEFT) and full fine-tuning (FFT) methods on moderately-sized LLMs in the context of DOCMT.

## 3 Experimental Setup

In this section, we detail our experimental setup, covering the training strategy (Section 3.1), datasets (Section 3.2), models (Section 3.3), and evaluation (Section 3.4).

### 3.1 Two-Stage Training

DOCMT approaches typically start by pre-training the translation model on parallel sentences and then fine-tuning it on parallel documents (Voita et al., 2019; Ma et al., 2020; Wu et al., 2023). Recently, Xu et al. (2023) proposed a two-stage training strategy. This involves initially fine-tuning a LLM on monolingual text, followed by a second fine-tuning phase on parallel text. Following Xu et al. (2023), our approach begins with fine-tuning an LLM on monolingual documents, followed by fine-tuning on parallel documents.

**Fine-tuning on Monolingual Documents** Existing LLMs are typically pre-trained on English-centric corpora. Recent research highlights that

these LLMs often exhibit sub-optimal performance on multilingual benchmarks (Li et al., 2023; Chen et al., 2023; Scao et al., 2022). To address this limitation, our initial step involves fine-tuning all the parameters of LLMs using monolingual data from the target languages.

**Fine-tuning on Parallel Documents** We fine-tune the model on document-level parallel corpora in this stage. Following Wang et al. (2023a), we condition each sentence pair on its context, consisting of the three preceding consecutive sentence pairs. As demonstrated by Wang et al. (2023b), the prompting strategy plays a significant role in translating documents using LLMs. However, they only investigate how the prompting strategies affect GPT-3.5-TURBO and GPT-4-TURBO at the inference stage. In our study, we first delve into how these prompting strategies impact the fine-tuning process, as shown in Figure 1, and we present our findings in Section 4.

### 3.2 Datasets

**Parallel Documents** Following Zhang et al. (2022), we conduct experiments on IWSLT2017 translation tasks (Cettolo et al., 2017). IWSLT2017 comprises translation datasets sourced from TED talks, encompassing translations between English and nine other languages, including Arabic, German, French, Italian, Japanese, Korean, Dutch, Romanian, and Chinese. There are approximately 1.9K sentence-aligned parallel documents with about 240K sentences for each language pair. The dataset statistics can be found in Appendix A.

**Monolingual Documents** We gather monolingual documents for all the target languages in our translation tasks, totaling ten languages. To manage computational limitations and address concerns about catastrophic forgetting that might result from excessive continued training, we leverage the data pruning technique suggested by Marion et al. (2023) to select 100M tokens for each language, including English, from the CulturaX corpus (Nguyen et al., 2023), totaling 1B tokens.

### 3.3 Models

**Baselines** The baseline models in this study can be classified into three categories, including state-of-the-art LLMs and SENMT models, and our re-implemented DOCMT models:

- **State-of-the-art SENMT models:** Our selection includes models such as NLLB, which

are available with three different sets of parameters: 600M, 1.3B, and 3.3B.<sup>1</sup> We also incorporate the widely-used commercial translation system, Google Translate.

- **State-of-the-art LLMs:** For our baseline LLMs in the context of DOCMT, we utilize GPT-3.5-TURBO and GPT-4-TURBO.<sup>2</sup> We use the Prompt 4 as detailed in Figure 1d during the translation process.
- **Our re-implemented DOCMT models:** We conduct full fine-tuning on the concatenation-based DOCMT model (Tiedemann and Scherrer, 2017), as well as several recent DOCMT baselines (Sun et al., 2022; Wu et al., 2023, 2024b), initialized with MT5 (Xue et al., 2021). These models are available with parameters of 300M, 580M, and 1.2B, representing the strong DOCMT baseline.

**Ours** In this work, we utilize LLAMA2-7B, BLOOM-7B, and VICUNA-7B, as our backbones.<sup>3</sup> The LLAMA2 models are predominantly pre-trained on English text, while the BLOOM models are pre-trained on multilingual text. The use of VICUNA models allows us to compare the differences between base models and instruction-tuned models (LLAMA2 vs. VICUNA). We denote those fully fine-tuned models as L-7B-FFT, B-7B-FFT, and V-7B-FFT. We denote those models fine-tuned with LORA (Hu et al., 2022) as L-7B-LORA, B-7B-LORA, and V-7B-LORA. The optimization details can be found in Appendix B.

### 3.4 Evaluation

**Evaluation Metrics** We evaluate the translation on the IWSLT2017 test sets (Cettolo et al., 2017) using sentence-level BLEU (Papineni et al., 2002) and document-level BLEU (Liu et al., 2020) using SacreBLEU (Post, 2018), denoted as *s*BLEU and *d*BLEU.<sup>4</sup> Furthermore, as conventional MT metrics like BLEU demonstrate poor correlation to human judgments (Freitag et al., 2022), we also evaluate the translation quality with the state-of-the-

<sup>1</sup>Model signatures: facebook/nllb-200-distilled-600M, facebook/nllb-200-1.3B, and facebook/nllb-200-3.3B.

<sup>2</sup>Model signatures: gpt-3.5-turbo-1106 and gpt-4-1106-preview.

<sup>3</sup>LLAMA2 signature: meta-llama/llama-2-7b-hf, BLOOM signature: bigscience/bloom-7b1, and VICUNA signature: lmsys/vicuna-7b-v1.5. Note that VICUNA-v1.5 models are fine-tuned from LLAMA2.

<sup>4</sup>BLEU signature: nrefs:1|case:mixed|eff:no|tok:[13a|ja-mecab-0.996-IPA|ko-mecab-0.996/ko-0.9.2-K0|zh]|smooth:exp|version:2.3.1.

```
[<src_lang> Context]: <src1> <src2> <src3>
[<tgt_lang> Context]: <tgt1> <tgt2> <tgt3>
[<src_lang> Sentence]: <src4>
[<tgt_lang> Sentence]: <tgt4>
```

(a) Prompt 1

```
[<src_lang> Context]: <src1> <src2> <src3>
[<tgt_lang> Context]: <tgt1> <tgt2> <tgt3>
Given the provided parallel context, translate the following
↪ <src_lang> sentence to <tgt_lang>:
[<src_lang> Sentence]: <src4>
[<tgt_lang> Sentence]: <tgt4>
```

(c) Prompt 3

```
[<src_lang>]: <src1> [<tgt_lang>]: <tgt1>
[<src_lang>]: <src2> [<tgt_lang>]: <tgt2>
[<src_lang>]: <src3> [<tgt_lang>]: <tgt3>
[<src_lang>]: <src4> [<tgt_lang>]: <tgt4>
```

(b) Prompt 2

```
[<src_lang>]: <src1> [<tgt_lang>]: <tgt1>
[<src_lang>]: <src2> [<tgt_lang>]: <tgt2>
[<src_lang>]: <src3> [<tgt_lang>]: <tgt3>
Given the provided parallel sentence pairs, translate the following
↪ <src_lang> sentence to <tgt_lang>:
[<src_lang>]: <src4> [<tgt_lang>]: <tgt4>
```

(d) Prompt 4

Figure 1: Prompt types used in the preliminary study. <src\_lang> and <tgt\_lang> indicate the source and target languages. <src\*> and <tgt\*> indicate the source and target sentences. **Note that the target sentences <tgt\*> are only used during training and are replaced with the hypotheses <hyp\*> generated by the model during inference.** Concrete examples for each prompt variation can be found in [Appendix C](#).

art neural evaluation metric COMET (Rei et al., 2020).<sup>5</sup> Moreover, we use the average sentence-level BLEU  $\mu_{sBLEU}$ , the average document-level BLEU  $\mu_{dBLEU}$ , and the average COMET  $\mu_{COMET}$  for the overall performance.

**Settings** We evaluate our models in two settings: *Bilingual English-from/to-Many Translations* and *Multilingual English-from/to-Many Translations*. For the bilingual translations, we train our models on the training set for one specific language pair. For the multilingual translations, we combine all the available training sets for training.

**Inference** In this work, we explore two inference strategies. The first strategy involves using previous translations as context for the current translation, translating test examples in their original order. This begins with the first sentence, which is free from context, as illustrated in Figure 1d. We refer to this strategy as **REUSE**. Alternatively, sentences within the context can be re-generated individually for each document translation to prevent error propagation during the translation process. We call this strategy **REGEN**. For reasons of inference efficiency, REUSE is used as the default inference strategy throughout this study, unless specified otherwise. A comparison of these two inference strategies is provided in Section 6.

## 4 A Preliminary Study on Prompts

The prompt plays a crucial role in LLM research. Recent studies show that an optimal prompt can greatly enhance model performance and reveal unexpected model capabilities (Kojima et al., 2022;

	PID	$\mu_{sBLEU}$	$\mu_{dBLEU}$	$\mu_{COMET}$
L-7B-LoRA	1	15.5	18.2	67.5
	2	19.0	21.9	70.7
	3	15.8	18.3	69.8
	4	<b>20.2</b>	<b>23.4</b>	<b>72.7</b>
B-7B-LoRA	1	19.3	20.5	70.5
	2	20.6	23.5	73.6
	3	19.8	20.8	73.9
	4	<b>23.1</b>	<b>27.3</b>	<b>76.8</b>
V-7B-LoRA	1	19.0	22.4	74.2
	2	20.4	23.5	71.6
	3	18.3	21.4	70.0
	4	<b>22.4</b>	<b>25.7</b>	<b>76.2</b>

Table 1: Overall performance given by L-7B-LoRA, B-7B-LoRA, and V-7B-LoRA on different prompt variations, across four English-centric translation tasks involving German and Chinese. PID indicates the prompt ID in Figure 1. Best results are highlighted in **bold**.

Wei et al., 2022b). Hence, our initial focus is on investigating the prompt’s impact during fine-tuning.

**Prompt Variations** Displayed in Figure 1, our preliminary study features four prompt types. These designs aim to tackle two research questions: *How does context structure impact translation quality?* (Prompt 1 vs. Prompt 2) and *How do natural language instructions influence translation quality?* (Prompt 1 vs. Prompt 3). We also investigate the combined effect of these aspects in Prompt 4.

**Results** Our investigation analyzes prompt variations using three PEFT models (L-7B-LoRA, B-7B-LoRA, and V-7B-LoRA) on four English-centric translation tasks involving German and Chinese. Overall results are presented in Table 1. Comparing Prompt 1 (Figure 1a) and Prompt 2 (Figure 1b), we find that models fine-tuned with Prompt 2 generally outperform those with Prompt 1, indicating Prompt 2’s effectiveness in enhancing

<sup>5</sup>COMET signature: Unbabel/wmt22-comet-da.

	# of param.	# of train. param.	En-X			X-En		
			$\mu_{sBLEU}$	$\mu_{dBLEU}$	$\mu_{COMET}$	$\mu_{sBLEU}$	$\mu_{dBLEU}$	$\mu_{COMET}$
<i>State-of-the-art SENMT baselines</i>								
NLLB	600M	—	23.6	27.3	82.3	18.2	22.1	72.8
	1.3B	—	25.7	29.5	83.5	25.0	28.7	78.1
	3.3B	—	<u>26.8</u>	<u>30.5</u>	<u>84.3</u>	<u>25.8</u>	<u>29.4</u>	78.9
GOOGLETRANS	—	—	24.5	28.4	81.6	25.0	28.5	<u>81.2</u>
<i>State-of-the-art LLMs</i>								
GPT-3.5-TURBO	—	—	26.3	30.1	85.3	30.7	34.1	85.5
GPT-4-TURBO	—	—	<b><u>27.0</u></b>	<b><u>30.7</u></b>	<b><u>86.3</u></b>	<b><u>31.7</u></b>	<b><u>35.1</u></b>	<b><u>86.0</u></b>
<i>LLM backbones</i>								
LLAMA2-7B	7B	—	2.7	3.5	40.1	4.2	4.4	52.2
BLOOM-7B	7B	—	2.5	2.9	35.5	6.7	7.3	49.4
VICUNA-7B	7B	—	<u>10.2</u>	<u>12.4</u>	<u>64.7</u>	<u>9.5</u>	<u>9.8</u>	<u>62.9</u>
<i>Re-implemented DOCMT baselines</i>								
Doc2Doc-MT5 (2017)	300M	300M	17.2	20.2	75.1	19.4	21.2	75.1
	580M	580M	18.6	21.5	78.3	20.7	22.5	77.4
	1.2B	1.2B	18.4	21.4	79.2	21.5	23.4	78.7
MR-Doc2Sen-MT5 (2022)	1.2B	1.2B	18.8	21.9	79.9	22.0	23.8	79.3
MR-Doc2Doc-MT5 (2022)	1.2B	1.2B	—	<u>22.5</u>	—	—	24.0	—
DocFLAT-MT5 (2023)	1.2B	1.2B	19.2	22.4	80.2	<u>22.2</u>	<u>24.3</u>	79.3
IADA-MT5 (2024c)	1.2B	1.2B	<u>19.3</u>	22.4	<u>80.4</u>	22.1	24.0	<u>79.5</u>
<i>Bilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>								
L-7B-LoRA	7B	8M	17.2	20.2	<u>70.8</u>	23.8	25.7	73.7
L-7B-FFT	7B	7B	13.7	16.2	67.4	22.4	24.1	74.0
B-7B-LoRA	7B	8M	<u>17.7</u>	<u>20.5</u>	68.5	<u>29.9</u>	<u>33.6</u>	<u>81.4</u>
B-7B-FFT	7B	7B	12.0	13.8	59.6	22.3	24.5	69.9
V-7B-LoRA	7B	8M	15.8	18.6	68.8	21.6	23.3	71.4
V-7B-FFT	7B	7B	14.3	16.8	65.0	21.8	23.5	74.3
<i>Multilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>								
L-7B-LoRA	7B	8M	13.9	16.9	67.2	17.1	18.4	62.0
L-7B-FFT	7B	7B	<u>17.1</u>	<u>20.4</u>	<u>73.8</u>	18.3	19.5	69.6
B-7B-LoRA	7B	8M	10.3	12.3	58.1	23.5	25.7	75.8
B-7B-FFT	7B	7B	16.0	18.6	67.8	<u>27.2</u>	<u>30.5</u>	78.7
V-7B-LoRA	7B	8M	12.9	15.7	66.4	17.2	18.6	60.6
V-7B-FFT	7B	7B	13.8	16.8	69.3	18.4	19.4	65.6

Table 2: Overall performance on IWSLT2017. # of param. indicates the number of parameters of the model. # of train. param. indicates the number of trainable parameters of the model. All the LLM approaches use Prompt 4 (Figure 1d) during inference. Best results are highlighted in **bold**. Best results in each group are underlined.

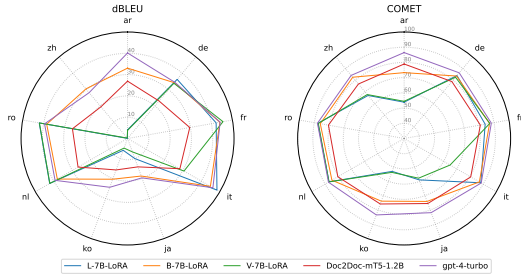
LLM performance. Regarding our second research question (Figure 1a vs. Figure 1c), we observe varied performance. L-7B-LoRA and B-7B-LoRA perform better with Prompt 3, while V-7B-LoRA performs better with Prompt 1. These results highlight varying impacts of prompt variations across models and suggest natural language instructions are less effective when using instruction-tuned language models as model backbones. Finally, LLMs with Prompt 4 (Figure 1d) achieve the best overall performance, suggesting a positive compound effect of context structure and instructions.

**Conclusion** As expected, the prompt plays a significant role in LLM performance. A well-structured prompt, combining appropriate context and clear instructions, can significantly boost

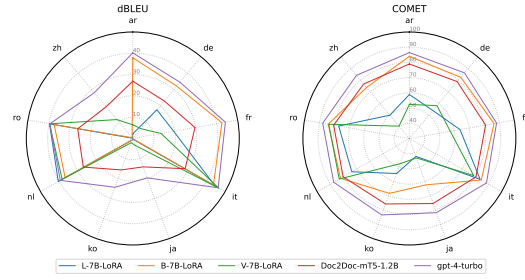
model performance. In this work, we use Prompt 4 (Figure 1d) in our experiments.

## 5 Main Results

**Overall Performance** In our results presented in Table 2, we observe that GPT-4-TURBO and GPT-3.5-TURBO significantly outperform all other models. Notably, the NLLB variants also demonstrate superior performance among specialized machine translation (MT) models. In the context of DOCMT, conventional DOCMT models still outperform our LLM-based DOCMT models for translations from English to other languages when evaluated using standard MT metrics. Conversely, for translations from other languages to English, our LLM-based DOCMT models perform on par or better than conventional DOCMT models in  $\mu_s$ BLEU



(a) Bilingual English-from-Many translations.



(b) Multilingual English-from-Many translations.

Figure 2: Breakdown results on  $sBLEU$ ,  $dBLEU$ , and COMET given by L-7B-LoRA, V-7B-LoRA, B-7B-LoRA, DOC2DOC-MT5-1.2B, and GPT-4-TURBO for the bilingual (Figure 2a) and multilingual (Figure 2b) translation tasks from other languages to English.

and  $\mu_{dBLEU}$  metrics, while conventional DOCMT models maintain superior performance in  $\mu_{COMET}$ .

**LLM-based DOCMT Models** As shown in Table 2, models using LoRA generally outperform fully fine-tuned (FFT) LLMs in bilingual translations, while fully fine-tuned LLMs surpass LoRA models in multilingual translations. The LoRA approach, which updates only a limited number of parameters, allows the model to retain the rich linguistic knowledge of the pre-trained LLM while adapting its capabilities to a particular bilingual context. However, in multilingual translation tasks, this parameter-efficient fine-tuning is insufficient to capture the complexities of multiple languages simultaneously. Conversely, fully fine-tuning all the parameters of LLMs enables better capture of multilingual complexity but makes the model prone to overfitting in bilingual translations.

**Breakdown Performance** We present the results for translation tasks from other languages to English in Figure 2. For clarity, only the results from our models using LoRA are shown. We observe that LLM-based DOCMT models sometimes even surpassing GPT-4-TURBO in certain bilingual and multilingual translation tasks, though they fail completely in others. Our manual review indicates that the primary cause of these failures is *off-target translation*. An in-depth analysis of this issue is provided in Section 6, and a complete breakdown of the results is available in Appendix E.

## 6 Analysis

In this section, we examine off-target issues with decoding strategy and use GPT-4 to analyze translation errors. We also explore discourse phenomena, training strategies, the data efficiency of parallel documents, and additional evaluations on recent

	$\mu\%$	Ar	Ja	Ko	Zh
L-7B-LoRA	29.2	87.9	25.5	44.2	93.1
L-7B-FFT	40.2	87.9	75.5	92.3	93.6
B-7B-LoRA	2.8	2.9	4.0	8.4	1.6
B-7B-FFT	28.0	54.1	43.8	70.4	76.4
V-7B-LoRA	32.3	88.2	40.4	35.7	90.5
V-7B-FFT	44.7	94.1	98.3	96.6	94.6

Table 3: Off-target rate (%) provided by our LLM-based DOCMT models for translation tasks from selective languages to English.  $\mu\%$  indicates the average off-target rate across all nine language pairs. **A lower off-target rate indicates better performance.** The complete results are presented in Appendix F.

WMT2023 test sets. Additionally, we evaluate the zero-shot crosslingual transfer capabilities of our models and present the results in Appendix I.

**Off-Target Translation** In Figure 2, our LLM-based DOCMT models excel in some translation tasks but struggle in others due to off-target translation issues. We investigate this problem using the *fasttext* library (Bojanowski et al., 2017) to identify translation languages and quantify off-target rates, which represent the proportion of translations that are off-target. Results are presented in Table 3, with off-target rates reaching up to 98.3% in failing tasks. Notably, only B-7B-LoRA consistently maintains low off-target rates, likely due to BLOOM-7B’s multilingual pre-training. Detailed off-target rates are provided in Appendix F.

**Inference Strategy** As described in Section 3.4, previous translations provide context for the current translation. We hypothesize that the high off-target translation rate in Table 3 is due to error propagation during the decoding stage. Therefore, we refer to the decoding strategy in Section 3.4 as REUSE, and introduce an alternative strategy where all con-

	sBLEU	dBLEU	COMET
L-7B-LoRA + REUSE	3.9	4.1	53.9
L-7B-LoRA + REGEN	17.5	19.6	69.7
L-7B-FFT + REUSE	2.5	2.6	51.6
L-7B-FFT + REGEN	15.9	17.4	65.0

Table 4: Results on Arabic-English translation given by different decoding strategies.

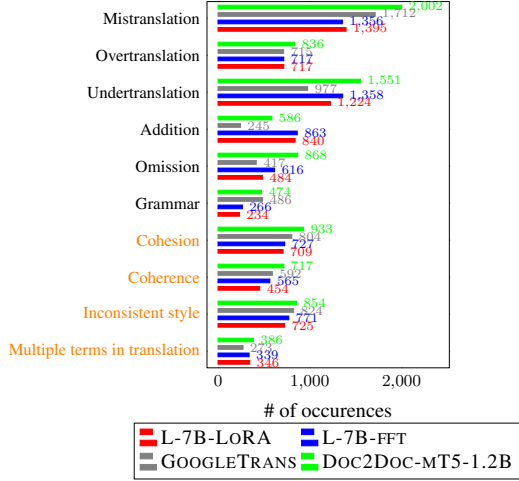


Figure 3: Error type analysis given by GPT-4-TURBO for translations from English to German, Romanian, and Chinese. The error types in orange are context-dependent. We omit those error types that are rare.

text translations are re-generated, called REGEN. As shown in Table 4, the REGEN strategy significantly improves translation quality, confirming our hypothesis on decoding error propagation. However, the inference cost of REGEN is four times higher than that of REUSE. These findings highlight the main reason for translation failures in LLM-based DOCMT models.

**Translation Errors** To understand the translation capabilities of our LLM-based DOCMT models, we select specific error types from the Multi-dimensional Quality Metrics (MQM) framework (Burchardt, 2013). Kocmi and Federmann (2023) demonstrate that GPT-4 can identify error spans and achieve state-of-the-art MT evaluation accuracy, so we use GPT-4-TURBO to analyze translation errors in texts translated by these models. Due to resource constraints, we focus on four models: L-7B-LoRA, L-7B-FFT, DOC2DOC-MT5-1.2B, and GOOGLETRANS, assessing translations from English to German, Romanian, and Chinese. The error identification prompt is detailed in Appendix D, and we present the frequency of error

	En-De	En-Fr
DOC2DOC-MT5-1.2B	77.0	89.9
L-7B-LoRA	83.1	<b>95.1</b>
L-7B-FFT	81.1	94.6
B-7B-LoRA	75.5	91.9
B-7B-FFT	68.3	90.8
V-7B-LoRA	<b>84.9</b>	94.8
V-7B-FFT	84.4	94.8

Table 5: Accuracy (in %) on the English-German and English-French contrastive test sets. Best results are highlighted in **bold**.

types in Figure 3. Most errors are limited to individual sentences. Despite similar scores in metrics such as sBLEU, dBLEU, and COMET among the models, our LLM-based DOCMT models (L-7B-LoRA and L-7B-FFT) exhibit fewer context-independent and context-dependent errors. This suggests that current evaluation metrics may not sufficiently assess document-level translations and indicates that fine-tuning LLMs holds promise for enhancing DOCMT performance.

**Discourse Phenomena** To evaluate our LLM-based DOCMT model’s ability to leverage contextual information, we assessed it using the English-German contrastive test set by Müller et al. (2018) and the English-French contrastive test set by Lopes et al. (2020). This evaluation tests the model’s accuracy in selecting the correct pronoun from multiple translation options. Results, shown in Table 5, reveal that models initialized with LLAMA2-7B and VICUNA-7B outperform DOC2DOC-MT5-1.2B, while BLOOM-7B-initialized models perform worse, indicating that contextual understanding is mostly acquired during pre-training, as detailed by Scao et al. (2022) due to the lack of German text in BLOOM pre-training. The generative accuracy results following Post and Junczys-Dowmunt (2023) are presented in Appendix G.

**Training Strategy** In this study, we follow the two-stage approach of Xu et al. (2023). Unlike traditional DOCMT methods that typically start with parallel sentence training, we investigate the effectiveness of this strategy on LLM-based DOCMT models. Specifically, we attempt to directly fine-tune the LLMs on parallel documents (one-stage) and add an extra fine-tuning stage using parallel sentences to the training strategy in Section 3.1 (three-stage). The results in Table 6 indicate that both the one-stage and three-stage training strate-

	sBLEU	dBLEU	COMET
<i>One-Stage</i>			
Nl-En	30.2	32.4	71.2
Ro-En	28.3	30.1	72.3
Ar-En	1.3	2.1	50.1
Zh-En	0.1	0.1	60.9
<i>Two-Stage</i>			
Nl-En	38.9	41.9	87.0
Ro-En	38.2	41.4	87.3
Ar-En	2.5	2.6	51.6
Zh-En	0.1	0.1	67.1
<i>Three-Stage</i>			
Nl-En	39.1	42.1	87.0
Ro-En	38.4	41.6	87.3
Ar-En	2.3	2.4	52.4
Zh-En	0.3	0.3	67.4

Table 6: Results from one-stage, two-stage, and three-stage training strategies. The one-stage strategy involves directly fine-tuning LLAMA2-7B on parallel documents. The two-stage results are produced by L-7B-FFT. In the three-stage strategy, all model parameters of LLAMA2-7B are fine-tuned across all three stages.

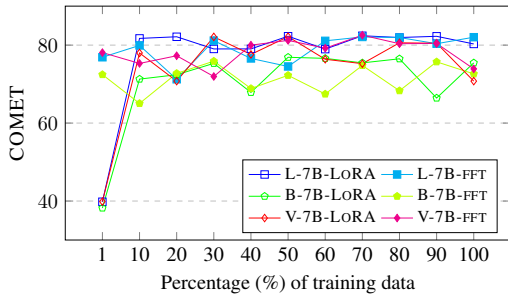


Figure 4: COMET-Percentage (%) of parallel documents for the translations from English to German.

gies are sub-optimal for both high-performing languages (Dutch and Romanian) and low-performing languages (Arabic and Chinese) with LLM-based DOCMT models.

**Data Efficiency of Parallel Documents** We explore the data efficiency of parallel documents. Results for English-German translation are presented in Figure 4, and for English-Romanian and English-Chinese in Appendix H. While LLMs typically excel with minimal training data, different fine-tuning strategies show distinct scaling behaviors. Our LoRA models match full training set performance with just 10% of the data (around 20K examples), while fully fine-tuned models achieve near-equivalent performance with only about 1% of the data (approximately 2K examples). These are crucial for low-resource languages, as recent LLMs are mainly pre-trained on English text.

	En-De		De-En	
	dBLEU	COMET	dBLEU	COMET
Doc2Doc-MT5-1.2B	20.2	74.4	20.0	76.5
MR-Doc2Sen-MT5	20.5	74.9	21.0	76.5
MR-Doc2Doc-MT5	21.2	75.6	21.5	76.5
DocFlat-MT5	20.9	75.1	21.8	76.5
IADA-MT5	21.2	75.4	22.0	76.5
L-7B-LoRA	28.9	76.4	35.5	83.2
L-7B-FFT	<b>29.0</b>	<b>77.0</b>	<b>36.1</b>	<b>84.0</b>
B-7B-LoRA	23.7	73.0	30.5	80.8
B-7B-FFT	21.0	69.0	30.0	80.5
V-7B-LoRA	20.5	63.8	33.9	81.8
V-7B-FFT	27.8	75.0	34.7	83.1

Table 7: dBLEU and COMET on WMT2023 test sets. Best results are highlighted in **bold**.

**Evaluation on Recent Test Sets** Given their extensive pre-training on large text corpora, LLMs are susceptible to data leakage risks. We evaluate our models using recent test sets from WMT2023 (Koehn et al., 2023). These tests, conducted between English and German, assess the out-of-domain generalization of our models and help mitigate data leakage risks. We use spaCy to segment documents and discard any parallel documents where the source and target sides have a differing number of sentences. Our findings, presented in Table 7, reveal that while MT5-based models outperform LLM-based models in Table 2, LLM-based models excel in translating out-of-domain text on the WMT2023 test sets. These results highlight the ability of LLM-based DOCMT models to generalize well to out-of-domain translation tasks.

## 7 Conclusion

In this study, we adapt large language models (LLMs) for document-level machine translation (DOCMT) across specific language pairs and found that moderately-sized, fine-tuned models can sometimes surpass larger models like GPT-4 in translation performance. By investigating prompt strategies and conducting extensive experiments with different fine-tuning methods and LLM backbones, we highlight both the potential and challenges of using LLMs for DOCMT. Despite achieving promising results, issues like off-target translations due to error propagation persist. Our analysis of translation errors, discourse phenomena, strategies for training and inference, and data efficiency underscores the strengths and limitations of LLM-based DOCMT models, providing a foundation for future research to enhance their effectiveness.

## 8 Limitations

**Constraints on Model Scale** Our research is confined to language models of a moderate size, specifically those with 7B parameters. This limitation is due to the constraints of our available resources. Consequently, it is crucial to acknowledge that the outcomes of our study might vary if conducted with larger models.

**Instability in Training** The process of supervised fine-tuning for LLMs shows instability in our observations. As detailed in Figure 4, there are noticeable inconsistencies in performance. These variations are too significant to attribute solely to the randomness inherent in training. In some cases, the fine-tuning of LLMs fails to reach convergence. Unfortunately, our limited resources restrict us from investigating these failures in depth or devising potential remedies.

**Influence of Prompting Techniques** Section 4 of our study highlights the significant role of prompting methods in fine-tuning. We experiment with four different prompting techniques. It is important to note that the prompt we recommend may not be the most effective, potentially leading to suboptimal performance of our models.

We acknowledge these limitations and leave them to the future work.

## References

Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernández Ábrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan A. Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vladimir Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, and et al. 2023. [Palm 2 technical report](#). *CoRR*, abs/2305.10403.

Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen McAleer, Albert Q. Jiang, Jia Deng, Stella Biderman, and Sean Welleck. 2023. [Llemma: An open language model for mathematics](#). *CoRR*, abs/2310.10631.

Rachel Bawden, Rico Sennrich, Alexandra Birch, and Barry Haddow. 2018. [Evaluating discourse phenomena in neural machine translation](#). 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 1304–1313, New Orleans, Louisiana. Association for Computational Linguistics.

Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.

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.

Aljoscha Burchardt. 2013. [Multidimensional quality metrics: a flexible system for assessing translation quality](#). In *Proceedings of Translating and the Computer 35*, London, UK. Aslib.

Mauro Cettolo, Marcello Federico, Luisa Bentivogli, Jan Niehues, Sebastian Stüker, Katsuhito Sudoh, Koichiro Yoshino, and Christian Federmann. 2017. [Overview of the IWSLT 2017 evaluation campaign](#). In *Proceedings of the 14th International Conference on Spoken Language Translation*, pages 2–14, Tokyo, Japan. International Workshop on Spoken Language Translation.

Pinzhen Chen, Shaoxiong Ji, Nikolay Bogoychev, Barry Haddow, and Kenneth Heafield. 2023. [Monolingual or multilingual instruction tuning: Which makes a better alpaca](#). *CoRR*, abs/2309.08958.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira,

641	Rewon Child, Oleksandr Polozov, Katherine Lee,	Young Jin Kim, Mohamed Afify, and Hany Has-	700
642	Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark	san Awadalla. 2023. <a href="#">How good are GPT models</a>	701
643	Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy	<a href="#">at machine translation? A comprehensive evaluation.</a>	702
644	Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov,	<i>CoRR</i> , abs/2302.09210.	703
645	and Noah Fiedel. 2022. <a href="#">Palm: Scaling language mod-</a>		
646	<a href="#">eling with pathways.</a> <i>CoRR</i> , abs/2204.02311.		
647	Hyung Won Chung, Le Hou, Shayne Longpre, Barret	Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan	704
648	Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang,	Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and	705
649	Mostafa Dehghani, Siddhartha Brahma, Albert Web-	Weizhu Chen. 2022. <a href="#">Lora: Low-rank adaptation of</a>	706
650	son, Shixiang Shane Gu, Zhuyun Dai, Mirac Suz-	<a href="#">large language models.</a> In <i>The Tenth International</i>	707
651	gun, Xinyun Chen, Aakanksha Chowdhery, Sharan	<i>Conference on Learning Representations, ICLR 2022,</i>	708
652	Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao,	<i>Virtual Event, April 25-29, 2022.</i> OpenReview.net.	709
653	Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav	Jingjing Huo, Christian Herold, Yingbo Gao, Leonard	710
654	Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam	Dahlmann, Shahram Khadivi, and Hermann Ney.	711
655	Roberts, Denny Zhou, Quoc V. Le, and Jason Wei.	2020. <a href="#">Diving deep into context-aware neural ma-</a>	712
656	2022. <a href="#">Scaling instruction-finetuned language models.</a>	<a href="#">chine translation.</a> In <i>Proceedings of the Fifth Confer-</i>	713
657	<i>CoRR</i> , abs/2210.11416.	<i>ence on Machine Translation</i> , pages 604–616, Online.	714
		Association for Computational Linguistics.	715
658	Marta R. Costa-jussà, James Cross, Onur Çelebi,	Albert Q. Jiang, Alexandre Sablayrolles, Arthur Men-	716
659	Maha Elbayad, Kenneth Heafield, Kevin Heffer-	sch, Chris Bamford, Devendra Singh Chaplot, Diego	717
660	nan, Elahe Kalbassi, Janice Lam, Daniel Licht,	de Las Casas, Florian Bressand, Gianna Lengyel,	718
661	Jean Maillard, Anna Sun, Skyler Wang, Guillaume	Guillaume Lample, Lucile Saulnier, Léo Ren-	719
662	Wenzek, Al Youngblood, Bapi Akula, Loïc Bar-	nard Lavaud, Marie-Anne Lachaux, Pierre Stock,	720
663	rault, Gabriel Mejia Gonzalez, Prangthip Hansanti,	Teven Le Scao, Thibaut Lavril, Thomas Wang, Timo-	721
664	John Hoffman, Semarley Jarrett, Kaushik Ram	thée Lacroix, and William El Sayed. 2023. <a href="#">Mistral</a>	722
665	Sadagopan, Dirk Rowe, Shannon Spruit, Chau	<a href="#">7b.</a> <i>CoRR</i> , abs/2310.06825.	723
666	Tran, Pierre Andrews, Necip Fazil Ayan, Shruti	Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing	724
667	Bhosale, Sergey Edunov, Angela Fan, Cynthia	Wang, and Zhaopeng Tu. 2023. <a href="#">Is chatgpt A</a>	725
668	Gao, Vedanuj Goswami, Francisco Guzmán, Philipp	<a href="#">good translator? A preliminary study.</a> <i>CoRR</i> ,	726
669	Koehn, Alexandre Mourachko, Christophe Rop-	abs/2301.08745.	727
670	pers, Safiyyah Saleem, Holger Schwenk, and Jeff	Tom Kocmi and Christian Federmann. 2023. <a href="#">GEMBA-</a>	728
671	Wang. 2022. <a href="#">No language left behind: Scal-</a>	<a href="#">MQM: Detecting translation quality error spans with</a>	729
672	<a href="#">ing human-centered machine translation.</a> <i>CoRR</i> ,	<a href="#">GPT-4.</a> In <i>Proceedings of the Eighth Conference</i>	730
673	abs/2207.04672.	<i>on Machine Translation</i> , pages 768–775, Singapore.	731
674	Yukun Feng, Feng Li, Ziang Song, Boyuan Zheng, and	Association for Computational Linguistics.	732
675	Philipp Koehn. 2022. <a href="#">Learn to remember: Trans-</a>	Philipp Koehn, Barry Haddow, Tom Kocmi, and	733
676	<a href="#">former with recurrent memory for document-level</a>	Christof Monz, editors. 2023. <a href="#">Proceedings of the</a>	734
677	<a href="#">machine translation.</a> In <i>Findings of the Association</i>	<a href="#">Eighth Conference on Machine Translation.</a> Associa-	735
678	<i>for Computational Linguistics: NAACL 2022</i> , pages	tion for Computational Linguistics, Singapore.	736
679	1409–1420, Seattle, United States. Association for		
680	Computational Linguistics.		
681	Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo,	Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yu-	737
682	Craig Stewart, Eleftherios Avramidis, Tom Kocmi,	taka Matsuo, and Yusuke Iwasawa. 2022. <a href="#">Large lan-</a>	738
683	George Foster, Alon Lavie, and André F. T. Martins.	<a href="#">guage models are zero-shot reasoners.</a> In <i>NeurIPS</i> .	739
684	2022. <a href="#">Results of WMT22 metrics shared task: Stop</a>	Sneha Kudugunta, Isaac Caswell, Biao Zhang, Xavier	740
685	<a href="#">using BLEU – neural metrics are better and more</a>	Garcia, Christopher A. Choquette-Choo, Katherine	741
686	<a href="#">robust.</a> In <i>Proceedings of the Seventh Conference</i>	Lee, Derrick Xin, Aditya Kusupati, Romi Stella,	742
687	<i>on Machine Translation (WMT)</i> , pages 46–68, Abu	Ankur Bapna, and Orhan Firat. 2023. <a href="#">MADLAD-</a>	743
688	Dhabi, United Arab Emirates (Hybrid). Association	<a href="#">400: A multilingual and document-level large audited</a>	744
689	for Computational Linguistics.	<a href="#">dataset.</a> <i>CoRR</i> , abs/2309.04662.	745
690	Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio	Haonan Li, Fajri Koto, Minghao Wu, Alham Fikri Aji,	746
691	César Teodoro Mendes, Allie Del Giorno, Sivakanth	and Timothy Baldwin. 2023. <a href="#">Bactrian-x : A multi-</a>	747
692	Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo	<a href="#">lingual replicable instruction-following model with</a>	748
693	de Rosa, Olli Saarikivi, Adil Salim, Shital Shah,	<a href="#">low-rank adaptation.</a> <i>CoRR</i> , abs/2305.15011.	749
694	Harkirat Singh Behl, Xin Wang, Sébastien Bubeck,	Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey	750
695	Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and	Edunov, Marjan Ghazvininejad, Mike Lewis, and	751
696	Yuanzhi Li. 2023. <a href="#">Textbooks are all you need.</a> <i>CoRR</i> ,	Luke Zettlemoyer. 2020. <a href="#">Multilingual denoising pre-</a>	752
697	abs/2306.11644.	<a href="#">training for neural machine translation.</a> <i>Transac-</i>	753
698	Amr Hendy, Mohamed Abdelrehim, Amr Sharaf,	<i>tions of the Association for Computational Linguis-</i>	754
699	Vikas Raunak, Mohamed Gabr, Hitokazu Matsushita,	<i>tics</i> , 8:726–742.	755

- António Lopes, M. Amin Farajian, Rachel Bawden, Michael Zhang, and André F. T. Martins. 2020. [Document-level neural MT: A systematic comparison](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 225–234, Lisboa, Portugal. European Association for Machine Translation.
- Hongyuan Lu, Haoyang Huang, Dongdong Zhang, Hao-ran Yang, Wai Lam, and Furu Wei. 2023. [Chain-of-dictionary prompting elicits translation in large language models](#). *CoRR*, abs/2305.06575.
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023. [Wizardcoder: Empowering code large language models with evol-instruct](#). *CoRR*, abs/2306.08568.
- Shuming Ma, Dongdong Zhang, and Ming Zhou. 2020. [A simple and effective unified encoder for document-level machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3505–3511, Online. Association for Computational Linguistics.
- Valentin Macé and Christophe Servan. 2019. [Using whole document context in neural machine translation](#). In *Proceedings of the 16th International Conference on Spoken Language Translation*, Hong Kong. Association for Computational Linguistics.
- Max Marion, Ahmet Üstün, Luiza Pozzobon, Alex Wang, Marzieh Fadaee, and Sara Hooker. 2023. [When less is more: Investigating data pruning for pretraining llms at scale](#). *CoRR*, abs/2309.04564.
- Sameen Maruf and Gholamreza Haffari. 2018. [Document context neural machine translation with memory networks](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1275–1284, Melbourne, Australia. Association for Computational Linguistics.
- Sameen Maruf, André F. T. Martins, and Gholamreza Haffari. 2019. [Selective attention for context-aware neural machine translation](#). 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 3092–3102, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sameen Maruf, Fahimeh Saleh, and Gholamreza Haffari. 2022. [A survey on document-level neural machine translation: Methods and evaluation](#). *ACM Comput. Surv.*, 54(2):45:1–45:36.
- Lesly Miculicich, Dhananjay Ram, Nikolaos Pappas, and James Henderson. 2018. [Document-level neural machine translation with hierarchical attention networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2947–2954, Brussels, Belgium. Association for Computational Linguistics.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2022. [Cross-task generalization via natural language crowdsourcing instructions](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3470–3487, Dublin, Ireland. Association for Computational Linguistics.
- Yasmin Moslem, Rejwanul Haque, John D. Kelleher, and Andy Way. 2023. [Adaptive machine translation with large language models](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 227–237, Tampere, Finland. European Association for Machine Translation.
- Mathias Müller, Annette Rios, Elena Voita, and Rico Sennrich. 2018. [A large-scale test set for the evaluation of context-aware pronoun translation in neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 61–72, Brussels, Belgium. Association for Computational Linguistics.
- Thuat Nguyen, Chien Van Nguyen, Viet Dac Lai, Hieu Man, Nghia Trung Ngo, Franck Dernoncourt, Ryan A. Rossi, and Thien Huu Nguyen. 2023. [Culturax: A cleaned, enormous, and multilingual dataset for large language models in 167 languages](#). *CoRR*, abs/2309.09400.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *NeurIPS*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Matt Post and Marcin Junczys-Dowmunt. 2023. [Escaping the sentence-level paradigm in machine translation](#). *CoRR*, abs/2304.12959.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.



989	Document-level machine translation with large lan-	Dubrovnik, Croatia. Association for Computational	1047
990	guage models. In <i>Proceedings of the 2023 Confer-</i>	Linguistics.	1048
991	<i>ence on Empirical Methods in Natural Language Pro-</i>		
992	<i>cessing</i> , pages 16646–16661, Singapore. Association	Minghao Wu, Abdul Waheed, Chiyu Zhang, Muham-	1049
993	for Computational Linguistics.	ad Abdul-Mageed, and Alham Aji. 2024a. <a href="#">LaMini-</a>	1050
		<a href="#">LM: A diverse herd of distilled models from large-</a>	1051
994	Longyue Wang, Zhaopeng Tu, Andy Way, and Qun	<a href="#">scale instructions</a> . In <i>Proceedings of the 18th Confer-</i>	1052
995	Liu. 2017. <a href="#">Exploiting cross-sentence context for neu-</a>	<i>ence of the European Chapter of the Association for</i>	1053
996	<a href="#">ral machine translation</a> . In <i>Proceedings of the 2017</i>	<i>Computational Linguistics (Volume 1: Long Papers)</i> ,	1054
997	<i>Conference on Empirical Methods in Natural Lan-</i>	pages 944–964, St. Julian’s, Malta. Association for	1055
998	<i>guage Processing</i> , pages 2826–2831, Copenhagen,	Computational Linguistics.	1056
999	Denmark. Association for Computational Linguis-		
1000	tics.	Minghao Wu, Yufei Wang, George Foster, Lizhen Qu,	1057
		and Gholamreza Haffari. 2024b. <a href="#">Importance-aware</a>	1058
1001	Yizhong Wang, Swaroop Mishra, Pegah Alipoormo-	<a href="#">data augmentation for document-level neural ma-</a>	1059
1002	labashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva	<a href="#">chine translation</a> . In <i>Proceedings of the 18th Confer-</i>	1060
1003	Naik, Arjun Ashok, Arut Selvan Dhanasekaran,	<i>ence of the European Chapter of the Association for</i>	1061
1004	Anjana Arunkumar, David Stap, Eshaan Pathak,	<i>Computational Linguistics (Volume 1: Long Papers)</i> ,	1062
1005	Giannis Karamanolakis, Haizhi Lai, Ishan Puro-	pages 740–752, St. Julian’s, Malta. Association for	1063
1006	hit, Ishani Mondal, Jacob Anderson, Kirby Kuznia,	Computational Linguistics.	1064
1007	Krima Doshi, Kuntal Kumar Pal, Maitreya Patel,		
1008	Mehrad Moradshahi, Mihir Parmar, Mirali Purohit,	Minghao Wu, Yufei Wang, George Foster, Lizhen Qu,	1065
1009	Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma,	and Gholamreza Haffari. 2024c. <a href="#">Importance-aware</a>	1066
1010	Ravsehaj Singh Puri, Rushang Karia, Savan Doshi,	<a href="#">data augmentation for document-level neural ma-</a>	1067
1011	Shailaja Keyur Sampat, Siddhartha Mishra, Sujan	<a href="#">chine translation</a> . <i>arXiv preprint arXiv:2401.15360</i> .	1068
1012	Reddy A, Sumanta Patro, Tanay Dixit, and Xudong		
1013	Shen. 2022. <a href="#">Super-NaturalInstructions: Generaliza-</a>	Minghao Wu, Yulin Yuan, Gholamreza Haffari, and	1069
1014	<a href="#">tion via declarative instructions on 1600+ NLP tasks</a> .	Longyue Wang. 2024d. (perhaps) beyond human	1070
1015	In <i>Proceedings of the 2022 Conference on Empiri-</i>	translation: Harnessing multi-agent collaboration for	1071
1016	<i>cal Methods in Natural Language Processing</i> , pages	translating ultra-long literary texts. <i>arXiv preprint</i>	1072
1017	5085–5109, Abu Dhabi, United Arab Emirates. As-	<i>arXiv:2405.11804</i> .	1073
1018	sociation for Computational Linguistics.		
1019	Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin	Haoran Xu, Young Jin Kim, Amr Sharaf, and Hany Has-	1074
1020	Guu, Adams Wei Yu, Brian Lester, Nan Du, An-	san Awadalla. 2023. <a href="#">A paradigm shift in machine</a>	1075
1021	drew M. Dai, and Quoc V. Le. 2022a. <a href="#">Finetuned</a>	<a href="#">translation: Boosting translation performance of</a>	1076
1022	<a href="#">language models are zero-shot learners</a> . In <i>The Tenth</i>	<a href="#">large language models</a> . <i>CoRR</i> , abs/2309.11674.	1077
1023	<i>International Conference on Learning Representa-</i>		
1024	<i>tions, ICLR 2022, Virtual Event, April 25-29, 2022</i> .	Linting Xue, Noah Constant, Adam Roberts, Mihir Kale,	1078
1025	OpenReview.net.	Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and	1079
		Colin Raffel. 2021. <a href="#">mT5: A massively multilingual</a>	1080
1026	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten	<a href="#">pre-trained text-to-text transformer</a> . In <i>Proceedings</i>	1081
1027	Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le,	<i>of the 2021 Conference of the North American Chap-</i>	1082
1028	and Denny Zhou. 2022b. <a href="#">Chain-of-thought prompt-</a>	<i>ter of the Association for Computational Linguistics:</i>	1083
1029	<a href="#">ing elicits reasoning in large language models</a> . In	<i>Human Language Technologies</i> , pages 483–498, On-	1084
1030	<i>NeurIPS</i> .	line. Association for Computational Linguistics.	1085
1031	KayYen Wong, Sameen Maruf, and Gholamreza Haf-	Wen Yang, Chong Li, Jiajun Zhang, and Chengqing	1086
1032	fari. 2020. <a href="#">Contextual neural machine translation</a>	Zong. 2023. <a href="#">Bigtrans: Augmenting large language</a>	1087
1033	<a href="#">improves translation of cataphoric pronouns</a> . In <i>Pro-</i>	<a href="#">models with multilingual translation capability over</a>	1088
1034	<i>ceedings of the 58th Annual Meeting of the Asso-</i>	<a href="#">100 languages</a> . <i>CoRR</i> , abs/2305.18098.	1089
1035	<i>ciation for Computational Linguistics</i> , pages 5971–		
1036	5978, Online. Association for Computational Lin-	Biao Zhang, Ankur Bapna, Melvin Johnson, Ali Dabir-	1090
1037	guistics.	moghaddam, Naveen Arivazhagan, and Orhan Firat.	1091
		2022. <a href="#">Multilingual document-level translation en-</a>	1092
1038	Minghao Wu and Alham Fikri Aji. 2023. <a href="#">Style over sub-</a>	<a href="#">ables zero-shot transfer from sentences to documents</a> .	1093
1039	<a href="#">stance: Evaluation biases for large language models</a> .	In <i>Proceedings of the 60th Annual Meeting of the</i>	1094
1040	<i>CoRR</i> , abs/2307.03025.	<i>Association for Computational Linguistics (Volume</i>	1095
		<i>1: Long Papers)</i> , pages 4176–4192, Dublin, Ireland.	1096
1041	Minghao Wu, George Foster, Lizhen Qu, and Gho-	Association for Computational Linguistics.	1097
1042	lamreza Haffari. 2023. <a href="#">Document flattening: Be-</a>		
1043	<a href="#">yond concatenating context for document-level neu-</a>	Jiacheng Zhang, Huanbo Luan, Maosong Sun, Feifei	1098
1044	<a href="#">ral machine translation</a> . In <i>Proceedings of the 17th</i>	Zhai, Jingfang Xu, Min Zhang, and Yang Liu. 2018.	1099
1045	<i>Conference of the European Chapter of the Associa-</i>	<a href="#">Improving the transformer translation model with</a>	1100
1046	<i>tion for Computational Linguistics</i> , pages 448–462,	<a href="#">document-level context</a> . In <i>Proceedings of the 2018</i>	1101
		<i>Conference on Empirical Methods in Natural Lan-</i>	1102
		<i>guage Processing</i> , pages 533–542, Brussels, Bel-	1103
		gium. Association for Computational Linguistics.	1104

- 1105 Pei Zhang, Boxing Chen, Niyu Ge, and Kai Fan. 2020.  
1106 [Long-short term masking transformer: A simple but](#)  
1107 [effective baseline for document-level neural machine](#)  
1108 [translation](#). In *Proceedings of the 2020 Conference*  
1109 *on Empirical Methods in Natural Language Process-*  
1110 *ing (EMNLP)*, pages 1081–1087, Online. Association  
1111 for Computational Linguistics.
- 1112 Shaolei Zhang, Qingkai Fang, Zhuocheng Zhang, Zhen-  
1113 grui Ma, Yan Zhou, Langlin Huang, Mengyu Bu,  
1114 Shangdong Gui, Yunji Chen, Xilin Chen, and Yang  
1115 Feng. 2023. [Bayling: Bridging cross-lingual align-](#)  
1116 [ment and instruction following through interac-](#)  
1117 [tive translation for large language models](#). *CoRR*,  
1118 abs/2306.10968.
- 1119 Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu,  
1120 Lingpeng Kong, Jiajun Chen, Lei Li, and Shujian  
1121 Huang. 2023. [Multilingual machine translation with](#)  
1122 [large language models: Empirical results and analy-](#)  
1123 [sis](#). *CoRR*, abs/2304.04675.

	Train		Validation		Test	
	# of sen.	# of doc.	# of sen.	# of doc.	# of sen.	# of doc.
En-Ar	232K	1907	2453	19	1460	12
En-De	206K	1705	2456	19	1138	10
En-Fr	233K	1914	2458	19	1455	12
En-It	232K	1902	2495	19	1147	10
En-Ja	223K	1863	2420	19	1452	12
En-Ko	230K	1920	2437	19	1429	12
En-Nl	237K	1805	2780	19	1181	10
En-Ro	221K	1812	2592	19	1129	10
En-Zh	231K	1906	2436	19	1459	12

Table 8: Dataset statistics of parallel documents.

## A Statistics of Parallel Documents

We present the dataset statistics of parallel documents in Table 8.

## B Optimization and Hyperparameters

**Fine-tuning on Monolingual Documents** We fine-tune all the parameters of large language models (LLMs) using a learning rate of  $5 \times 10^{-5}$  and a batch size of 256. During the training process, we apply the linear learning rate schedule, which includes a warm-up phase comprising 10% of the total training steps.

**Fine-tuning on Parallel Documents** When fine-tuning L-7B-LORA and V-7B-LORA on parallel documents, we use a learning rate of  $5 \times 10^{-5}$  and a batch size of 64. We apply a linear learning rate schedule with a 10% warm-up phase. The LoRA rank is set to 16, affecting only 0.1% of the parameters (about 8M parameters). The same hyperparameters are used for fine-tuning DOC2DOC-MT5 models, except for a learning rate of  $5 \times 10^{-4}$ . L-7B-LORA and V-7B-LORA are fine-tuned for up to 3 epochs, and DOC2DOC-MT5 models for up to 10 epochs. Early stopping is based on validation loss.

## C Prompt Types

We present concrete examples of prompt variations in Figure 5.

## D GPT-4 Prompts

We present the prompts used for error type analysis in Figure 6.

## E Breakdown Results

We provide detailed breakdowns of the translation tasks from English to other languages, evaluated

using sBLEU, dBLEU, and COMET. These are presented in Table 9, Table 10, and Table 11, respectively. Additionally, we present similar breakdowns for translations from other languages to English, assessed using the same metrics. These results can be found in Table 12, Table 13, and Table 14.

## F Off-Target Translation

We present the complete results on the off-target translation problem in Table 15 and Table 16.

## G Discourse Phenomena

Post and Junczys-Dowmunt (2023) propose to evaluate the accuracy on the constrastive test sets in a generative way. Hence, we present the generative accuracy in Table 17.

## H Scaling Law of Parallel Documents from English to Romanian and Chinese

In Section 6, we find that our LLM-based DOCMT models are highly efficient in terms of the amount of training data. To confirm our findings in Section 6, we conduct additional experiments on the translation tasks from English to Romanian and Chinese. As shown in Figure 7, we can confirm the superiority of LLM-based DOCMT models with regard to data efficiency.

## I Zero-Shot Crosslingual Transfer

We also explore the transferability of translation capabilities acquired from one language pair to others. We assess our English-German LLM-based DOCMT models on English-to-other-language test sets, comparing their COMET scores to their base models in Table 18. Our results indicate that models with fine-tuned instructions (LLAMA2-7B and BLOOM-7B) consistently exhibit positive transfer effects across all language pairs, while those with instruction-tuned models (VICUNA-7B) benefits only few languages.

	$\mu_{sBLEU}$	Ar	De	Fr	It	Ja	Ko	Nl	Ro	Zh
<i>State-of-the-art SENMT baselines</i>										
NLLB-600M	23.6	14.2	22.2	38.5	36.0	11.4	16.2	30.2	25.6	17.6
NLLB-1.3B	25.7	16.2	27.6	40.6	37.7	12.6	18.6	32.2	27.3	18.3
NLLB-3.3B	26.8	17.4	28.8	41.3	39.2	14.1	19.5	33.7	28.1	18.7
GOOGLETRANS	24.5	14.2	25.3	38.0	35.0	11.6	16.5	29.6	24.0	26.4
<i>State-of-the-art LLMs</i>										
GPT-3.5-TURBO	26.3	14.9	27.2	40.5	36.6	13.2	15.9	31.5	26.6	30.4
GPT-4-TURBO	27.0	16.1	27.4	40.0	35.8	14.1	18.3	32.2	27.3	31.6
<i>LLM backbones</i>										
LLAMA2	2.8	0.4	6.6	4.5	1.0	1.0	1.9	0.2	1.9	7.4
BLOOM	2.5	1.0	1.0	12.1	1.4	0.1	3.1	0.7	0.1	3.4
VICUNA	10.2	4.5	6.4	6.4	8.6	10.2	9.8	13.9	6.8	25.4
<i>Re-implemented DOCMT baselines</i>										
Doc2Doc-MT5-300M	17.2	9.4	16.8	24.0	21.0	11.0	13.7	20.5	17.1	21.6
Doc2Doc-MT5-580M	18.6	10.8	18.2	24.9	23.0	12.9	15.2	21.7	17.8	22.9
Doc2Doc-MT5-1.2B	18.4	10.3	18.1	24.9	22.4	13.9	15.4	19.6	18.8	22.6
MR-Doc2SEN-MT5	18.8	10.2	18.8	25.6	22.3	14.5	16.2	19.6	19.3	22.8
MR-Doc2Doc-MT5	—	—	—	—	—	—	—	—	—	—
DocFLAT-MT5	19.2	11.0	19.2	25.7	22.6	14.7	16.5	20.3	19.2	23.8
IADA-MT5	19.7	11.7	19.4	26.3	23.9	15.2	16.9	20.9	19.6	23.4
<i>Bilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	17.2	13.0	25.1	34.9	6.8	8.7	13.0	3.7	22.7	27.3
L-7B-FFT	13.7	13.1	25.3	19.5	2.6	7.9	7.2	4.1	21.1	22.8
B-7B-LoRA	17.7	12.1	20.6	32.6	32.9	3.6	1.4	28.1	12.2	15.7
B-7B-FFT	12.0	10.1	19.6	38.5	0.1	1.9	2.4	1.5	19.9	14.5
V-7B-LoRA	16.4	13.3	20.1	20.7	13.6	9.1	14.3	5.5	23.0	28.1
V-7B-FFT	14.3	13.5	23.3	21.1	4.8	3.8	15.9	3.3	17.4	25.8
<i>Multilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	13.9	8.8	22.2	27.9	5.0	7.7	9.4	3.4	19.4	21.6
L-7B-FFT	17.1	7.8	22.0	32.5	17.9	12.5	11.5	4.2	20.3	25.1
B-7B-LoRA	10.3	12.5	3.8	36.0	12.8	5.8	1.6	3.7	3.4	12.9
B-7B-FFT	16.0	12.6	15.8	37.4	27.7	2.0	1.3	18.4	7.8	21.3
V-7B-LoRA	12.9	8.0	15.4	28.2	6.2	10.1	7.1	3.3	15.1	22.4
V-7B-FFT	13.8	10.4	15.2	11.3	8.4	11.1	12.2	12.4	17.0	26.3

Table 9: Breakdown  $sBLEU$  results for the translation tasks from English to other languages.

[English Context]: And it's truly a great  
 ↪ honor to have the opportunity to come to  
 ↪ this stage twice; I'm extremely grateful.  
 ↪ I have been blown away by this conference,  
 ↪ and I want to thank all of you for the  
 ↪ many nice comments about what I had to say  
 ↪ the other night. And I say that sincerely,  
 ↪ partly because I need that.  
 [German Context]: Es ist mir wirklich eine  
 ↪ Ehre, zweimal auf dieser Bühne stehen zu  
 ↪ dürfen. Tausend Dank dafür. Ich bin  
 ↪ wirklich begeistert von dieser Konferenz,  
 ↪ und ich danke Ihnen allen für die vielen  
 ↪ netten Kommentare zu meiner Rede  
 ↪ vorgestern Abend. Das meine ich ernst,  
 ↪ teilweise deshalb -- weil ich es wirklich  
 ↪ brauchen kann!  
 [English Sentence]: Put yourselves in my  
 ↪ position.  
 [German Sentence]: Versetzen Sie sich mal in  
 ↪ meine Lage!

[English]: And it's truly a great honor to  
 ↪ have the opportunity to come to this stage  
 ↪ twice; I'm extremely grateful.  
 [German]: Es ist mir wirklich eine Ehre,  
 ↪ zweimal auf dieser Bühne stehen zu dürfen.  
 ↪ Tausend Dank dafür.  
 [English]: I have been blown away by this  
 ↪ conference, and I want to thank all of you  
 ↪ for the many nice comments about what I  
 ↪ had to say the other night.  
 [German]: Ich bin wirklich begeistert von  
 ↪ dieser Konferenz, und ich danke Ihnen  
 ↪ allen für die vielen netten Kommentare zu  
 ↪ meiner Rede vorgestern Abend.  
 [English]: And I say that sincerely, partly  
 ↪ because I need that.  
 [German]: Das meine ich ernst, teilweise  
 ↪ deshalb -- weil ich es wirklich brauchen  
 ↪ kann!  
 [English]: Put yourselves in my position.  
 [German]: Versetzen Sie sich mal in meine  
 ↪ Lage!

(a) Prompt 1

[English Context]: And it's truly a great  
 ↪ honor to have the opportunity to come to  
 ↪ this stage twice; I'm extremely grateful.  
 ↪ I have been blown away by this conference,  
 ↪ and I want to thank all of you for the  
 ↪ many nice comments about what I had to say  
 ↪ the other night. And I say that sincerely,  
 ↪ partly because I need that.  
 [German Context]: Es ist mir wirklich eine  
 ↪ Ehre, zweimal auf dieser Bühne stehen zu  
 ↪ dürfen. Tausend Dank dafür. Ich bin  
 ↪ wirklich begeistert von dieser Konferenz,  
 ↪ und ich danke Ihnen allen für die vielen  
 ↪ netten Kommentare zu meiner Rede  
 ↪ vorgestern Abend. Das meine ich ernst,  
 ↪ teilweise deshalb -- weil ich es wirklich  
 ↪ brauchen kann!  
 Given the provided parallel context, translate  
 ↪ the following English sentence to German:  
 [English Sentence]: Put yourselves in my  
 ↪ position.  
 [German Sentence]: Versetzen Sie sich mal in  
 ↪ meine Lage!

(b) Prompt 2

[English]: And it's truly a great honor to  
 ↪ have the opportunity to come to this stage  
 ↪ twice; I'm extremely grateful.  
 [German]: Es ist mir wirklich eine Ehre,  
 ↪ zweimal auf dieser Bühne stehen zu dürfen.  
 ↪ Tausend Dank dafür.  
 [English]: I have been blown away by this  
 ↪ conference, and I want to thank all of you  
 ↪ for the many nice comments about what I  
 ↪ had to say the other night.  
 [German]: Ich bin wirklich begeistert von  
 ↪ dieser Konferenz, und ich danke Ihnen  
 ↪ allen für die vielen netten Kommentare zu  
 ↪ meiner Rede vorgestern Abend.  
 [English]: And I say that sincerely, partly  
 ↪ because I need that.  
 [German]: Das meine ich ernst, teilweise  
 ↪ deshalb -- weil ich es wirklich brauchen  
 ↪ kann!  
 Given the provided parallel sentence pairs,  
 ↪ translate the following English sentence  
 ↪ to German:  
 [English]: Put yourselves in my position.  
 [German]: Versetzen Sie sich mal in meine  
 ↪ Lage!

(c) Prompt 3

(d) Prompt 4

Figure 5: Prompt types used in the preliminary study. <src\_lang> and <tgt\_lang> indicate the language IDs. <src\*> and <tgt\*> indicate the source and target sentences. **Note that the target sentences <tgt\*> are only used during training and are replaced with the hypotheses <hyp\*> generated by the model during inference.**

```

[Context]:

[Source]: <src1>
[Reference]: <tgt1>
[Hypothesis]: <hyp1>
[Source]: <src2>
[Reference]: <tgt2>
[Hypothesis]: <hyp2>
[Source]: <src3>
[Reference]: <tgt3>
[Hypothesis]: <hyp3>

[Current Sentence]:

[Source]: <src4>
[Reference]: <tgt4>
[Hypothesis]: <hyp4>

[Error Types]:

- Mistranslation: Error occurring when the target content does not accurately represent the source
  ↪ content.
- Overtranslation: Error occurring in the target content that is inappropriately more specific than
  ↪ the source content.
- Undertranslation: Error occurring in the target content that is inappropriately less specific than
  ↪ the source content.
- Addition: Error occurring in the target content that includes content not present in the source.
- Omission: Error where content present in the source is missing in the target.
- Unjustified euphemism: Target content that is potentially offensive in some way in the source
  ↪ language, but that has been inappropriately "watered down" in the translation.
- Do not translate: Error occurring when a text segment marked "Do not translate!" is translated in
  ↪ the target text.
- Untranslated: Error occurring when a text segment that was intended for translation is omitted in
  ↪ the target content.
- Retained factual error: Untrue statement or an incorrect data value present in the source content
  ↪ and retained in the target content.
- Completeness: Source text incomplete, resulting in instances where needed content is missing in
  ↪ the source language.
- Grammar: Error that occurs when a text string (sentence, phrase, other) in the translation
  ↪ violates the grammatical rules of the target language.
- Punctuation: Punctuation incorrect according to target language conventions.
- Spelling: Error occurring when a word is misspelled.
- Duplication: Content (e.g., a word or longer portion of text) repeated unintentionally.
- Unclear reference: Relative pronouns or other referential mechanisms unclear in their reference.
- Cohesion: Portions of the text needed to connect it into an understandable whole (e.g., reference,
  ↪ substitution, ellipsis, conjunction, and lexical cohesion) missing or incorrect.
- Coherence: Text lacking a clear semantic relationship between its parts, i.e., the different parts
  ↪ don't hang together, don't follow the discourse conventions of the target language, or don't
  ↪ "make sense."
- Inconsistent style: Style that varies inconsistently throughout the text, e.g., One part of a text
  ↪ is written in a clear, "terse" style, while other sections are written in a more wordy style.
- Multiple terms in translation: Error where source content terminology is correct, but target
  ↪ content terms are not used consistently.

Considering the provided context, please identify the errors of the translation from the source to
  ↪ the target in the current sentence based on a subset of Multidimensional Quality Metrics (MQM)
  ↪ error typology.
You should pay extra attention to the error types related to the relationship between the current
  ↪ sentence and its context, such as "Unclear reference", "Cohesion", "Coherence", "Inconsistent
  ↪ style", and "Multiple terms in translation".
You should list all the errors you find in the sentence, and provide a justification for each error.
Your output should always be in JSON format, formatted as follows: {'justification': '...',
  ↪ 'error_types': [...]}.
```

Figure 6: Prompt used for analyzing translation error types.

	$\mu_{dBLEU}$	Ar	De	Fr	It	Ja	Ko	Nl	Ro	Zh
<i>State-of-the-art SENMT baselines</i>										
NLLB-600M	27.3	15.4	26.0	42.1	39.0	17.1	23.8	33.8	28.5	20.2
NLLB-1.3B	29.5	17.4	31.8	44.0	40.8	18.1	26.6	35.8	30.2	21.0
NLLB-3.3B	30.5	18.6	32.9	44.6	42.2	19.8	27.4	37.0	30.9	21.2
GOOGLETRANS	28.4	16.0	29.3	41.3	38.5	15.7	23.4	32.8	26.7	32.1
<i>State-of-the-art LLMs</i>										
GPT-3.5-TURBO	30.1	16.4	30.9	43.7	39.7	17.5	22.7	34.5	29.0	36.3
GPT-4-TURBO	30.7	17.4	31.1	43.2	39.0	18.4	25.3	35.3	29.8	37.2
<i>LLM backbones</i>										
LLAMA2	3.5	0.5	7.4	4.9	1.1	1.8	3.7	0.2	2.2	9.6
BLOOM	2.8	1.0	1.3	12.9	1.7	0.3	2.7	1.0	0.1	4.4
VICUNA	12.4	5.7	6.4	7.1	8.5	15.0	16.1	14.4	7.4	31.2
<i>Re-implemented DOCMT baselines</i>										
Doc2Doc-MT5-300M	20.2	10.3	18.8	26.1	21.9	16.8	21.5	21.5	18.4	26.6
Doc2Doc-MT5-580M	21.5	11.7	20.0	27.0	23.9	18.4	23.2	22.6	18.8	28.0
Doc2Doc-MT5-1.2B	21.4	11.2	20.1	27.1	23.4	19.7	23.0	20.7	20.2	27.1
MR-Doc2SEN-MT5	—	—	—	—	—	—	—	—	—	—
MR-Doc2Doc-MT5	21.9	11.9	20.7	27.9	23.8	19.8	23.3	21.5	20.7	27.9
DocFLAT-MT5	22.5	12.1	20.8	28.0	24.7	20.9	24.3	21.9	21.7	27.9
IADA-MT5	22.4	12.2	21.3	28.4	24.0	20.8	24.1	21.1	21.4	28.1
<i>Bilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	20.2	14.7	29.1	37.5	7.3	13.9	19.5	4.2	22.9	33.1
L-7B-FFT	16.2	14.7	29.4	20.6	2.7	12.5	12.3	4.5	21.6	27.6
B-7B-LoRA	20.5	13.7	24.8	36.1	36.3	6.9	2.6	32.2	12.2	19.7
B-7B-FFT	13.8	11.2	23.6	41.7	0.1	3.7	3.9	1.7	20.1	18.1
V-7B-LoRA	19.3	14.9	23.1	21.8	14.7	14.3	21.6	5.9	23.3	34.2
V-7B-FFT	16.8	15.2	26.9	22.3	4.9	6.2	23.6	3.7	17.5	31.1
<i>Multilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	16.9	10.2	26.0	30.4	5.6	12.8	15.7	3.8	20.2	27.6
L-7B-FFT	20.4	9.2	25.5	35.5	18.5	19.2	18.3	4.9	20.9	31.6
B-7B-LoRA	12.3	14.3	4.4	39.5	14.4	10.8	3.3	3.9	3.5	16.8
B-7B-FFT	18.6	14.2	19.5	40.9	30.9	3.8	2.6	20.9	8.5	26.5
V-7B-LoRA	15.7	9.5	18.1	31.0	6.5	16.0	11.8	3.9	16.1	28.1
V-7B-FFT	16.8	12.0	17.5	12.9	8.8	17.4	19.5	12.4	17.8	32.7

Table 10: Breakdown  $dBLEU$  results for the translation tasks from English to other languages.

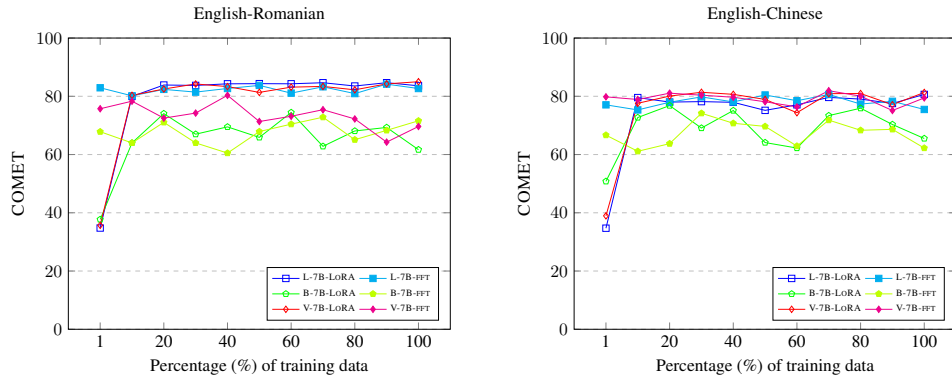


Figure 7: COMET-Percentage (%) of training data for the translations from English to Romanian, and Chinese.

	$\mu$ COMET	Ar	De	Fr	It	Ja	Ko	Nl	Ro	Zh
<i>State-of-the-art SENMT baselines</i>										
NLLB-600M	82.3	82.6	81.3	83.7	86.3	79.5	82.5	84.0	85.0	75.4
NLLB-1.3B	83.5	84.3	83.0	84.8	87.5	80.1	84.2	85.2	86.3	76.4
NLLB-3.3B	84.3	84.8	84.1	85.3	88.0	82.1	85.2	86.0	86.4	76.7
GOOGLETRANS	81.6	81.5	80.2	82.3	85.1	80.7	79.9	83.7	82.9	78.5
<i>State-of-the-art LLMs</i>										
GPT-3.5-TURBO	85.3	83.8	84.6	85.9	87.7	84.7	84.2	86.3	86.5	83.9
GPT-4-TURBO	86.3	85.4	85.4	86.2	87.8	86.0	86.4	87.0	87.4	84.8
<i>LLM backbones</i>										
LLAMA2	40.1	37.4	41.5	41.2	39.3	39.7	42.8	35.5	41.6	42.1
BLOOM	35.5	34.9	34.8	45.3	33.0	33.9	35.5	34.2	28.6	39.0
VICUNA	64.7	68.3	48.7	49.0	62.5	81.3	72.4	64.1	56.1	80.2
<i>Re-implemented DOCMT baselines</i>										
Doc2Doc-MT5-300M	75.1	77.2	71.0	72.4	74.1	78.5	77.8	74.3	74.0	77.0
Doc2Doc-MT5-580M	78.3	80.8	74.4	74.8	77.9	81.5	82.0	76.8	76.6	79.6
Doc2Doc-MT5-1.2B	79.2	81.1	75.9	75.9	78.9	82.9	82.4	76.5	79.3	80.3
MR-Doc2SEN-MT5	79.9	82.2	76.3	76.5	80.0	83.6	83.1	76.7	79.7	80.9
MR-Doc2Doc-MT5	—	—	—	—	—	—	—	—	—	—
DocFLAT-MT5	80.4	81.8	77.3	76.9	80.3	83.6	83.4	78.0	80.7	81.8
IADA-MT5	80.7	82.4	77.0	77.3	80.9	84.1	83.7	77.8	80.9	81.8
<i>Bilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	70.8	82.5	80.3	79.1	42.9	70.4	75.4	42.5	83.6	80.6
L-7B-FFT	67.4	83.1	82.0	59.4	38.8	65.8	69.7	49.2	82.7	75.5
B-7B-LoRA	68.5	77.0	75.4	76.8	85.1	51.4	40.4	82.9	61.7	65.5
B-7B-FFT	59.6	68.4	72.6	83.8	45.3	40.3	45.2	46.8	71.6	62.2
V-7B-LoRA	69.7	82.7	70.8	60.1	56.2	70.2	76.9	44.1	85.0	81.3
V-7B-FFT	65.0	83.1	73.9	58.8	41.5	54.8	81.1	42.4	69.6	79.4
<i>Multilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	67.2	78.8	74.1	69.3	40.9	65.9	74.7	44.9	79.3	77.1
L-7B-FFT	73.8	78.8	76.3	77.0	65.2	80.6	79.0	45.3	81.9	79.8
B-7B-LoRA	58.1	80.9	43.8	79.9	54.9	66.7	46.6	38.1	46.9	65.3
B-7B-FFT	67.8	82.2	69.6	83.8	81.7	46.3	45.7	70.5	54.5	76.0
V-7B-LoRA	66.4	78.2	65.5	70.8	46.9	74.5	66.4	41.9	76.0	77.6
V-7B-FFT	69.3	80.3	64.7	54.2	52.0	78.1	77.5	57.6	78.1	81.1

Table 11: Breakdown COMET results for the translation tasks from English to other languages.

	$\mu_{sBLEU}$	Ar	De	Fr	It	Ja	Ko	Nl	Ro	Zh
<i>State-of-the-art SENMT baselines</i>										
NLLB-600M	18.2	26.8	11.0	31.0	22.7	10.9	13.6	17.1	13.2	17.9
NLLB-1.3B	25.0	35.9	18.8	37.4	35.3	13.2	15.7	26.0	23.1	19.9
NLLB-3.3B	25.8	36.5	22.3	36.8	33.5	12.4	18.5	28.3	25.2	19.1
GOOGLETRANS	25.0	28.7	26.1	34.7	35.1	10.2	13.3	30.8	29.6	16.6
<i>State-of-the-art LLMs</i>										
GPT-3.5-TURBO	30.7	35.8	30.8	40.7	41.8	15.5	17.3	36.1	35.4	22.9
GPT-4-TURBO	31.7	37.2	31.2	41.4	42.3	15.9	19.8	36.6	36.5	24.4
<i>LLM backbones</i>										
LLAMA2	4.2	0.1	6.2	1.3	4.4	0.0	0.1	15.1	10.3	0.1
BLOOM	6.7	4.7	8.7	14.3	16.9	0.1	0.3	9.4	5.5	0.2
VICUNA	9.5	1.1	14.4	24.9	14.1	4.7	0.4	17.2	8.5	0.0
<i>Re-implemented DOCMT baselines</i>										
DOC2DOC-MT5-300M	19.4	23.0	19.5	26.6	25.4	9.5	11.4	22.0	22.6	14.5
DOC2DOC-MT5-580M	20.7	24.2	20.3	28.0	26.6	11.0	11.6	24.4	23.8	16.1
DOC2DOC-MT5-1.2B	21.5	25.7	21.0	28.7	27.3	11.0	12.8	25.4	25.0	16.8
MR-DOC2SEN-MT5	22.5	26.9	22.0	29.9	27.7	11.8	13.9	26.5	26.1	18.0
MR-DOC2DOC-MT5	—	—	—	—	—	—	—	—	—	—
DOCFLAT-MT5	22.5	26.6	22.3	29.7	28.4	11.7	13.9	26.4	25.9	17.6
IADA-MT5	23.1	26.9	22.7	30.5	28.5	12.8	14.5	27.1	26.1	18.7
<i>Multilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	23.8	3.9	33.1	40.3	45.2	8.3	5.0	39.2	39.0	0.1
L-7B-FFT	22.4	2.5	32.2	42.6	44.8	1.0	1.0	38.9	38.2	0.1
B-7B-LoRA	29.9	30.9	30.5	41.1	41.2	13.9	15.5	35.0	35.2	25.6
B-7B-FFT	22.3	17.0	29.8	41.3	40.7	0.4	1.1	35.6	34.3	1.0
V-7B-LoRA	21.6	3.8	30.8	43.1	29.4	5.5	4.0	39.1	38.7	0.3
V-7B-FFT	24.3	24.2	31.0	43.4	45.0	0.0	0.7	36.2	38.0	0.1
<i>Multilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	17.1	2.3	15.7	18.6	41.6	0.4	0.8	37.6	36.7	0.3
L-7B-FFT	18.3	15.8	10.3	22.0	40.7	2.1	6.0	29.7	30.8	7.8
B-7B-LoRA	23.5	34.4	28.5	39.5	40.0	0.1	0.6	33.7	34.1	0.6
B-7B-FFT	27.2	35.9	28.3	39.7	38.8	7.1	11.2	32.4	33.5	17.9
V-7B-LoRA	17.2	6.1	5.5	12.2	43.3	2.9	1.8	35.9	36.4	10.2
V-7B-FFT	18.4	13.0	17.7	16.7	37.6	2.5	3.7	36.2	35.2	3.0

Table 12: Breakdown  $sBLEU$  results for the translation tasks from other languages to English.

	$\mu_{dBLEU}$	Ar	De	Fr	It	Ja	Ko	Nl	Ro	Zh
<i>State-of-the-art SENMT baselines</i>										
NLLB-600M	22.0	30.5	14.7	34.1	26.3	14.8	18.1	21.0	16.8	22.2
NLLB-1.3B	28.6	39.2	22.6	40.1	38.7	17.1	20.3	29.4	26.7	23.7
NLLB-3.3B	29.4	39.7	26.1	39.6	37.0	16.5	23.2	31.3	28.8	22.8
GOOGLETRANS	28.5	32.0	29.8	37.8	38.9	13.3	17.7	33.7	33.1	20.4
<i>State-of-the-art LLMs</i>										
GPT-3.5-TURBO	34.0	38.8	34.0	43.4	44.8	19.5	22.0	38.8	38.4	26.9
GPT-4-TURBO	35.1	40.2	34.5	44.2	46.3	19.7	24.4	39.3	39.6	28.1
<i>LLM backbones</i>										
LLAMA2	4.4	0.1	6.9	1.5	4.7	0.0	0.1	15.7	10.9	0.1
BLOOM	7.3	5.3	9.8	15.2	17.6	0.1	0.5	10.6	6.6	0.2
VICUNA	9.8	1.1	14.5	24.6	14.4	5.9	0.5	18.5	8.9	0.0
<i>Re-implemented DOCMT baselines</i>										
DOC2DOC-MT5-300M	21.2	24.5	21.1	27.5	26.5	12.6	14.1	23.5	23.9	17.0
DOC2DOC-MT5-580M	22.5	25.5	22.1	28.9	27.8	14.0	14.6	25.8	25.2	18.5
DOC2DOC-MT5-1.2B	23.4	26.9	23.0	29.7	28.4	14.2	15.7	26.8	26.3	19.4
MR-DOC2SEN-MT5	24.0	27.4	24.2	30.3	29.4	14.9	16.1	27.5	26.8	19.8
MR-DOC2DOC-MT5	24.6	28.3	24.3	30.5	29.8	15.7	16.8	27.8	27.8	20.8
DOCFLAT-MT5	24.6	27.6	24.5	31.1	29.7	15.1	17.0	28.1	27.8	20.3
IADA-MT5	24.5	28.2	24.6	30.9	29.6	15.0	17.1	27.8	27.1	20.5
<i>Bilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	25.7	4.1	36.2	42.2	48.5	10.0	5.9	42.2	42.1	0.1
L-7B-FFT	24.1	2.6	35.3	45.1	48.1	1.0	1.1	41.9	41.4	0.1
B-7B-LoRA	33.6	33.0	34.2	44.0	44.9	18.8	20.4	38.1	38.6	30.4
B-7B-FFT	24.5	18.6	33.4	44.2	44.4	0.6	1.4	38.7	37.9	1.1
V-7B-LoRA	23.3	3.8	33.9	45.5	30.6	6.7	4.8	42.1	42.0	0.3
V-7B-FFT	23.5	2.3	34.1	46.0	48.3	0.0	0.7	39.2	41.0	0.0
<i>Multilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	18.4	2.2	17.6	19.5	44.6	0.4	0.8	40.4	39.7	0.2
L-7B-FFT	19.5	15.4	11.1	23.3	43.5	2.4	7.0	31.4	33.0	8.5
B-7B-LoRA	25.7	37.9	32.1	42.5	43.7	0.2	0.7	36.7	37.5	0.5
B-7B-FFT	30.5	39.5	31.9	42.6	41.9	10.0	16.0	35.3	37.0	20.6
V-7B-LoRA	18.6	6.4	6.0	13.5	46.6	3.5	2.1	38.3	39.7	11.6
V-7B-FFT	19.4	13.4	17.1	17.4	39.6	2.9	4.4	37.8	38.1	3.6

Table 13: Breakdown  $dBLEU$  results for the translation tasks from other languages to English.

	$\mu$ COMET	Ar	De	Fr	It	Ja	Ko	Nl	Ro	Zh
<i>State-of-the-art SENMT baselines</i>										
NLLB-600M	72.8	76.6	63.0	79.8	72.4	74.2	76.3	68.4	67.0	77.6
NLLB-1.3B	78.1	82.3	71.9	83.8	80.6	75.3	78.4	77.5	75.4	77.9
NLLB-3.3B	78.9	82.6	76.1	83.7	80.2	74.9	80.1	78.7	76.5	77.7
GOOGLETRANS	81.2	81.1	82.3	84.6	84.5	75.6	76.9	84.2	83.8	78.1
<i>State-of-the-art LLMs</i>										
GPT-3.5-TURBO	85.5	85.7	86.0	87.9	88.1	81.4	82.4	87.2	87.4	83.6
GPT-4-TURBO	86.0	86.5	86.3	88.2	88.5	81.9	83.5	87.5	87.9	84.2
<i>LLM backbones</i>										
LLAMA2	52.2	50.3	47.1	42.6	51.6	56.1	55.7	56.5	53.0	56.9
BLOOM	49.4	50.6	52.8	55.5	56.8	44.5	44.3	51.2	45.2	43.3
VICUNA	62.7	51.3	65.2	70.5	57.3	69.5	56.3	68.0	58.8	67.5
<i>Re-implemented DOCMT baselines</i>										
Doc2Doc-MT5-300M	75.1	75.0	75.2	78.0	77.5	71.8	72.4	75.3	76.6	73.9
Doc2Doc-MT5-580M	77.4	77.4	77.0	79.7	79.8	74.7	74.3	78.5	79.1	75.8
Doc2Doc-MT5-1.2B	78.7	79.0	78.8	80.9	80.5	75.5	75.8	80.3	80.5	76.8
MR-Doc2SEN-MT5	79.8	80.3	79.8	82.3	81.1	76.4	76.6	81.5	81.8	78.2
MR-Doc2Doc-MT5	—	—	—	—	—	—	—	—	—	—
DocFLAT-MT5	80.3	80.2	80.0	82.5	81.7	77.4	77.6	82.2	82.3	78.4
IADA-MT5	80.4	80.3	80.7	82.9	82.3	77.7	77.7	81.8	81.7	78.5
<i>Bilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	73.7	53.9	84.0	84.1	88.2	59.0	53.0	87.0	87.6	66.9
L-7B-FFT	74.0	51.6	81.9	86.5	88.3	63.6	52.9	87.0	87.3	67.1
B-7B-LoRA	81.4	73.3	83.6	87.0	87.1	74.0	73.8	84.8	86.0	82.6
B-7B-FFT	69.9	53.7	83.2	86.9	86.8	50.5	43.3	85.1	84.4	55.5
V-7B-LoRA	71.4	54.5	82.6	87.2	64.9	57.8	53.7	87.1	87.3	67.7
V-7B-FFT	77.0	77.0	81.3	87.0	88.2	65.6	55.5	84.2	87.2	67.0
<i>Multilingual English-from/to-Many LLM-based DOCMT models (Ours)</i>										
L-7B-LoRA	62.0	58.9	54.0	63.8	83.7	42.5	54.5	73.7	77.2	50.0
L-7B-FFT	69.6	52.9	63.2	63.2	85.1	65.3	57.0	86.3	86.3	66.8
B-7B-LoRA	78.7	84.0	82.5	86.5	84.6	62.4	68.4	82.3	84.2	73.8
B-7B-FFT	75.8	81.0	83.3	86.2	86.9	59.1	51.7	83.9	85.7	64.4
V-7B-LoRA	60.6	52.4	58.2	57.3	78.8	43.6	47.6	83.3	83.8	40.6
V-7B-FFT	65.6	50.9	47.4	53.6	87.2	61.4	56.8	83.3	86.3	63.5

Table 14: Breakdown COMET results for the translation tasks from other languages to English.

	$\mu\%$	Ar	De	Fr	It	Ja	Ko	Nl	Ro	Zh
L-7B-LoRA	6.2	0.4	4.6	1.2	0.5	17.1	0.4	0.8	1.0	29.7
L-7B-FFT	9.4	0.3	0.9	0.3	4.4	17.7	9.4	15.4	1.2	34.7
B-7B-LoRA	11.2	8.4	1.0	20.8	3.9	16.4	0.0	2.8	0.9	46.9
B-7B-FFT	31.8	36.6	15.8	2.7	90.1	10.7	0.1	82.0	0.2	47.7
V-7B-LoRA	10.6	0.2	15.4	0.3	13.3	15.9	0.5	20.3	0.9	28.9
V-7B-FFT	8.9	0.1	14.4	0.5	0.4	27.8	0.5	4.6	0.4	31.5

Table 15: Off-target rate (%) provided by our LLM-based DOCMT models for translation tasks from English to other languages.  $\mu\%$  indicates the average off-target rate. **A lower off-target rate indicates better performance.**

	$\mu\%$	Ar	De	Fr	It	Ja	Ko	Nl	Ro	Zh
L-7B-LoRA	29.2	87.9	2.0	4.9	1.4	25.5	44.2	1.9	1.8	93.1
L-7B-FFT	40.2	87.9	5.8	2.1	1.5	75.5	92.3	1.6	1.9	93.6
B-7B-LoRA	2.8	2.9	2.1	1.0	1.3	4.0	8.4	1.9	2.0	1.6
B-7B-FFT	28.0	54.1	2.0	1.0	1.1	43.8	70.4	1.6	1.9	76.4
V-7B-LoRA	32.3	88.2	2.6	1.2	28.0	40.4	35.7	1.9	1.9	90.5
V-7B-FFT	44.7	94.1	9.0	1.3	1.3	98.3	96.6	5.3	1.9	94.6

Table 16: Off-target rate (%) provided by our LLM-based DOCMT models for translation tasks from other languages to English.  $\mu\%$  indicates the average off-target rate. **A lower off-target rate indicates better performance.**

	En-De	En-Fr
Doc2Doc-MT5-1.2B	39.8	26.2
L-7B-LoRA	<b>64.4</b>	29.9
L-7B-FFT	62.4	28.4
B-7B-LoRA	58.3	23.6
B-7B-FFT	49.8	25.1
V-7B-LoRA	48.9	<b>30.2</b>
V-7B-FFT	43.5	27.2

Table 17: Generative accuracy (in %) on the English-German and English-French contrastive test sets. Best results are highlighted in **bold**.

	$\mu_{\Delta}$	Ar	De	Fr	It	Ja	Ko	Nl	Ro	Zh
L-7B-LoRA	+29.4	+36.3	+38.8	+37.2	+32.1	+15.9	+17.1	+21.7	+35.8	+29.5
L-7B-FFT	+29.0	+41.2	+40.5	+37.1	+18.0	+27.7	+29.4	+11.2	+18.5	+37.5
B-7B-LoRA	+20.3	+7.5	+40.7	+20.7	+21.9	+17.5	+15.9	+23.7	+25.3	+9.8
B-7B-FFT	+27.3	+14.8	+37.8	+28.9	+43.3	+13.1	+15.3	+38.5	+34.7	+19.5
V-7B-LoRA	-8.9	-12.6	+22.1	+18.9	-28.6	-27.8	-18.7	-11.8	+12.1	-34.1
V-7B-FFT	-1.4	+7.3	+25.2	+17.7	-14.6	-24.7	-5.3	-21.8	+7.6	-3.5

Table 18: The difference ( $\Delta$ ) in COMET scores on the test sets from English to other languages between our English-German LLM-based DOCMT models and their backbones.  $\mu_{\Delta}$  indicates the average difference across all the languages in this table.