

Fine-Grained and Multi-Dimensional Metrics for Document-Level Machine Translation

Yirong Sun¹, Dawei Zhu², Yanjun Chen^{1,3}, Erjia Xiao⁴, Xinghao Chen^{1,3}, Xiaoyu Shen^{1*}

¹ Digital Twin Institute, Eastern Institute of Technology, Ningbo, China

² Saarland University, Saarland Informatics

³ Department of Computing, The Hong Kong Polytechnic University

⁴ The Hong Kong University of Science and Technology (Guangzhou)

win1282467298@gmail.com

{xyshen}@eitech.edu.cn

Abstract

Large language models (LLMs) have excelled in various NLP tasks, including machine translation (MT), yet most studies focus on sentence-level translation. This work investigates the inherent capability of instruction-tuned LLMs for document-level translation (docMT). Unlike prior approaches that require specialized techniques, we evaluate LLMs by directly prompting them to translate entire documents in a single pass. Our results show that this method improves translation quality compared to translating sentences separately, even without document-level fine-tuning. However, this advantage is not reflected in BLEU scores, which often favor sentence-based translations. We propose using the LLM-as-a-judge paradigm for evaluation, where GPT-4 is used to assess document coherence, accuracy, and fluency in a more nuanced way than n-gram-based metrics. Overall, our work demonstrates that instruction-tuned LLMs can effectively leverage document context for translation. However, we caution against using BLEU scores for evaluating docMT, as they often provide misleading outcomes, failing to capture the quality of document-level translation.¹

1 Introduction

Large language models (LLMs) have demonstrated exceptional performance across a wide range of natural language processing tasks (Radford et al., 2019; Brown et al., 2020; Touvron et al., 2023; Dubey et al., 2024). In the realm of machine translation (MT), recent findings also suggest that LLM-based models rival dedicated commercial systems like Google Translate, particularly in translating high-resource languages (Hendy et al., 2023; Peng et al., 2023; Jiao et al., 2023; Zhu et al., 2024a,b). Nonetheless, most research has focused only on

sentence-level translation. While some studies have begun to explore document-level translation (docMT) with LLMs, there is a prevailing belief that directly applying instruction-tuned LLMs to docMT performs poorly without specialized training and prompting techniques, largely due to the limited availability of document-level content in instruction-tuning datasets (Wu et al., 2024; Cui et al., 2024; Li et al., 2024). However, their conclusions are frequently drawn from n-gram-based metrics without thorough analysis to substantiate the models' true performance.

In this work, we conduct an in-depth investigation into the *inherent* capabilities of instruction-tuned LLMs in handling docMT tasks. Unlike previous studies that explore special tricks, such as multi-turn inference (Wang et al., 2023), we directly prompt LLMs to translate entire documents in a single pass. Comparing this method to a simpler baseline that translates individual sentences separately and then stitches them together, we can evaluate whether instruction-tuned LLMs can leverage their inherent ability to incorporate document-level context and improve translation quality.

A key challenge in our research is the evaluation of document-level machine translation (docMT). Traditional metrics² like BLEU³, ChrF, and TER (Papineni et al., 2002; Popović, 2015; Snover et al., 2006), though widely used, often poorly correlate with human judgment (Freitag et al., 2022), especially in docMT, where maintaining coherence and logical flow across a document is essential—something n-gram overlap struggles to capture. Metrics like cTT, AZPT, and BLONDE (Jiang et al., 2021; Wang et al., 2023)

²While COMET (Rei et al., 2020) is more reliable than BLEU for sentence-level translation, it is trained exclusively on sentence-level data. As a result, using COMET to evaluate docMT can be unreliable, since out-of-distribution.

³Although we do not want to use BLEU based metric, it remains a common metric in existing/recent research, despite its limitations.

*Corresponding authors.

¹Our code and the outputs from GPT4-as-a-judge are available at https://github.com/EIT-NLP/BLEUless_DocMT

address specific aspects such as terminology consistency and zero-pronoun accuracy, but still rely heavily on word matching and symbolic statistics. We argue that an ideal docMT metric should be (1) context-aware—capturing document-level coherence and accuracy, (2) structured—evaluating aspects such as fluency, accuracy, and coherence separately, and (3) interpretable—explicitly identifying translation errors for clear, objective human evaluation. To this end, we design a GPT-4-based evaluation pipeline to provide deeper insights into the docMT capabilities of LLMs.

- We show that translating entire documents yields better results than translating sentences independently then merging them, even without document-level fine-tuning.
- We propose using the LLM-as-a-judge paradigm with multiple prompts that assess different aspects of translated text to achieve a more targeted and accurate evaluation.
- We recommend against using d-BLEU scores for docMT, as they fail to capture discourse-level phenomena and can often provide misleading results.

2 Problem Settings

Given a document containing l source sentences $\mathbf{X} = \{x^1, \dots, x^l\}$, the goal of docMT is to generate its translation $\mathbf{Y} = \{y^1, \dots, y^l\}$ as a sequence of sentences in the target language. In this work, we explore two approaches for generating translations using instruction-tuned LLMs:

- **ST[k]**: We concatenate k source sentences into a chunk, input each chunk into the LLM for translation, and then concatenate the translated chunks together to form the full document translation.
- **DOC**: We instruct the LLM to directly translate the entire document in one pass.

The DOC approach is designed to capture inter-sentence dependencies by considering the full document context, potentially leading to more coherent and accurate translations. However, this approach requires the LLM to process and generate longer sequences of text, which can increase the risk of cumulative errors, especially if the model has not been explicitly optimized for document-level translation.

3 BLEU-based Evaluation

Document-level BLEU (d-BLEU, Liu et al., 2020) is widely used for evaluating translations in DocMT. However, we notice that it is sensitive to overly lengthy generation, which can be problematic as LLMs sometimes overgenerate. We find that even minor overgeneration can significantly affect the final d-BLEU score.⁴ We argue that documents are generally independent units, so they should be weighted equally in the evaluation. We, therefore, propose an alternative, AvgBLEU, defined as:

$$\text{AvgBLEU} = \frac{1}{N} \sum_{i=1}^N \text{BLEU} \left(Y_i^{\text{ref}}, Y_i^{\text{pred}} \right)$$

Here, N is the number of documents, and \mathbf{Y}^{ref} and \mathbf{Y}^{pred} represent the reference document translations and the predicted translations, respectively. This allows us to calculate the average BLEU score (AvgBLEU) for the entire dataset, providing a comprehensive measure of translation quality.

	Number of Sentences	Avg. Document Length
zh-en	1142	252
en-zh	1696	219
de-en	1899	204
en-de	1780	231
Total	6517	225

Table 1: Statistics of our test set. The document length is measured by the token count using Vicuna’s tokenizer.

Evaluation Setup. For evaluation, we use the test set from WMT22 (Kocmi et al., 2022), which includes sentence-level reference translations along with annotated document boundaries. Document-level references are obtained by concatenating the corresponding sentence translations. We cover four translation directions in our evaluation: German (de) and Chinese (zh) translated to and from English (en). Specific dataset statistics are presented in Table 1. We evaluate five instruction-tuned LLMs: Vicuna-7B/13B (Zheng et al., 2023), their -16K versions and Mistral-instruct-7B (Jiang et al., 2023), all of which have very limited document-level content in their instruction-tuning datasets.

Results. Table 2 presents the comparison between the two document-level translation approaches. ST[k] consistently achieves higher AvgBLEU scores across all models and nearly all translation directions, with zh-en using Vicuna-7B and

⁴For completeness, we report results using the standard d-BLEU in Appendix B.

Model	Eval Type	Translation Direction			
		zh-en	en-zh	de-en	en-de
Vicuna-7B	ST1	19.70	30.97	29.42	20.82
	ST2	19.69	31.65	29.56	22.10
	ST3	19.62	32.14	29.22	22.53
	DOC	20.50	31.70	29.15	21.94
Vicuna-7B-16K	ST1	20.26	28.08	28.16	21.11
	ST2	20.05	31.17	28.78	22.99
	ST3	19.99	31.64	28.89	22.93
	DOC	20.20	30.77	28.65	21.57
Vicuna-13B	ST1	22.40	36.22	30.50	25.03
	ST2	21.01	35.82	30.89	25.46
	ST3	21.13	36.24	30.84	25.66
	DOC	21.83	34.93	30.60	25.59
Vicuna-13B-16K	ST1	21.07	35.55	29.87	25.22
	ST2	20.97	36.76	30.47	24.87
	ST3	20.79	36.46	30.71	25.58
	DOC	21.07	34.97	30.62	25.14
Mistral-7B	ST1	19.82	26.24	29.23	21.28
	ST2	18.89	26.84	29.86	21.44
	ST3	18.78	26.87	29.82	21.74
	DOC	18.61	24.31	28.98	21.09

Table 2: AvgBLEU scores with different translation approaches across four translation directions. The best scores are in bold, with red/blue shading indicating the highest score paradigm, respectively. In most cases, merged sentence translations yield higher BLEU scores than direct document translations.

Vicuna-13B-16K as the only two exceptions. The specific value of k that yields the highest AvgBLEU score varies depending on the translation direction, however, on average, ST3 achieves the highest score overall. While independently translated sentences yield better AvgBLEU scores than document translations done in one pass by LLMs, manual inspection reveals that ST[k] translations often contain more redundancy, literal translations, and disjointed phrasing. While these translations may achieve higher AvgBLEU scores, we find that DOC translations result in more fluent, readable, and cohesive output. This raises concerns about how much AvgBLEU can be trusted as a metric for evaluating docMT.

4 LLM-as-a-judge Evaluation

Maruf et al. (2021) outlines various discourse phenomena that should be considered when evaluating document-level translations, such as cohesion and the use of discourse connectives. In the past, automatic evaluation of these aspects was difficult due to the need for deep semantic understanding, and evaluations typically focused on one aspect at a time using specialized test sets (Hardmeier and Federico, 2010; Gong et al., 2015; Jwalapuram et al., 2019). Inspired by the “LLM-as-a-judge” approach (Zheng et al., 2023), we aim to assess mul-

tiples aspects simultaneously using a strong LLM.

Evaluation Setup. We design four (sub) metrics: (1) **Fluency**, (2) Content Errors (**CE**), (3) Lexical Cohesion Errors (**LE**), and (4) Grammatical Cohesion Errors (**GE**). All metrics are measured using prompts provided to GPT-4. See Appendix C for details on prompt design.

Fluency is rated on a scale of 1 to 5, with higher being better. Since fluency can be evaluated solely based on the translated text, we present only the model’s outputs to GPT-4 for this assessment, decoupling fluency from metrics that require consideration of source and reference texts.

Content Errors refer to translation mistakes such as mistranslations, omissions, or additions. We instruct GPT-4 (gpt-4-0613) to output a list containing all identified mistakes. The CE score is determined by the length of this list, and report the average CE score over the test set.⁵

Cohesion Errors are further divided into two subcategories: lexical (LE) and grammatical (GE), which affect text connection and the logic flow, respectively. LE includes incorrect vocabulary usage, missing synonyms, or overuse of certain terms that disrupt the flow. GE includes pronouns, conjunctions, and sentence-linking structure mistakes. Similar to CE, we prompt GPT-4 to generate a list of identified errors, with the score corresponding to the length of the list.

Other settings, such as translation directions and the models of interest, remain consistent with Section 3. Due to the cost associated with using GPT-4, we sample 70 documents per translation direction from the WMT22 dataset for our evaluation.

Results. The results with en-zh are shown in Table 3. Although ST3 scores higher than DOC on AvgBLEU, DOC consistently outperforms ST3 in Fluency. Additionally, DOC generally exhibits fewer CE, also known as content errors. For cohesion errors, the results are mixed: DOC shows better LE with vicuna-7B and its -16K version, and Mistral-7B, while Vicuna-13B and its -16K version yield higher LE. As for GE, DOC performs better with -16K models and Mistral-7B while others are mostly comparable. We also observe that the -16K versions perform similarly to their original counterparts in fluency but demonstrate notable improvements in CE reduction. This pattern is

⁵For simplicity, all mistake types are equally weighted, but our approach is flexible and can easily use different weights if certain types are considered more severe than others.

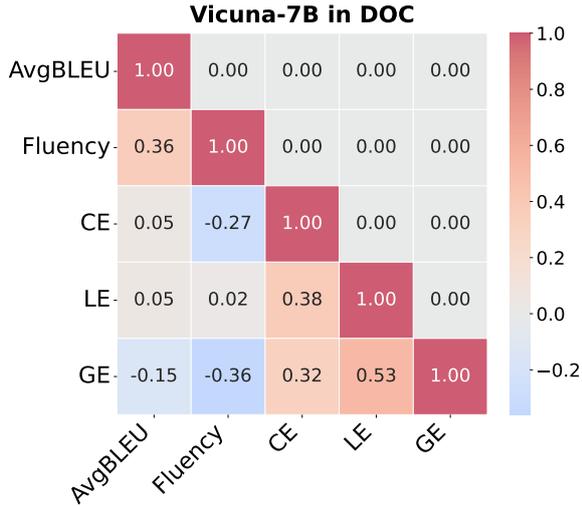


Figure 1: PCC Heatmaps among AvgBLEU, Fluency, CE, LE and GE for Vicuna-7B under DOC evaluation type in the en-zh translation direction.

consistent across all translation directions, with full results provided in Appendix D. Overall, our approach enables a more detailed evaluation of translation quality in DocMT. It clearly shows that instruction-tuned LLMs, even without fine-tuning for document-level MT tasks, are effective at capturing long-context information for DocMT.

To gain a deeper understanding of how these metrics correlates with each other, we compute the Pearson Correlation Coefficients (PCC) among those metrics and visualize them in Figure 1, as well as translation directions, showing that BLEU score has poor correlation with those discourse-level phenomena metrics. Other translation directions also exhibit low correlation results in Appendix E. Therefore, we suggest not using BLEU score for docMT since it fails to account for discourse-level phenomena, and even worse, it often produces misleading results—such as suggesting that sentence translations are better.

Human Agreement. While some judgments by the LLM-as-a-judge may appear reasonable, certain nuances may still be misinterpreted due to unique human perspectives. To validate the alignment between our LLM-as-a-judge paradigm and human evaluations, we conducted experiments to assess agreement. For each model in both ST3 and DOC, we used 10 samples per translation direction and asked human evaluators to respond with a simple “yes” or “no” regarding their agreement with the LLM-as-a-judge’s assessments according to our metrics.

Model	Eval Type	AvgBLEU↑	Fluency↑	CE↓	LE↓	GE↓
Vicuna-7B	ST3	33.44	3.64	4.97	2.55	1.21
	DOC	28.48	4.04	4.40	2.31	1.25
Vicuna-7B-16K	ST3	31.30	3.08	5.30	2.22	1.71
	DOC	30.80	3.97	4.72	2.17	1.15
Vicuna-13B	ST3	37.44	3.78	4.82	1.70	1.14
	DOC	35.58	4.12	4.87	2.02	1.14
Vicuna-13B-16k	ST3	38.66	2.98	4.21	1.84	1.02
	DOC	34.25	4.10	4.15	2.04	0.95
Mistral-7B	ST3	26.82	2.80	6.77	4.08	2.62
	DOC	23.27	3.11	5.98	3.71	2.51

Table 3: Evaluation results (en-zh) by GPT-4 for Vicuna-7B, Vicuna-13B, their -16K versions and Mistral-7B under ST3 and DOC, showing metrics AvgBLEU, fluency, content errors, lexical cohesion errors, and grammatical cohesion errors. Best performances are in bold, with red/blue shading indicating the winning paradigm, respectively.

Our manual evaluation confirmed a strong alignment between human judgments and the LLM-as-a-judge paradigm. As shown in Table 4, GPT-4-as-a-judge achieved approximately 95% agreement with human evaluations across all languages and evaluation types (ST3 and DOC), indicating robust concordance with human judgment across translation directions and metrics. This high level of agreement further validates GPT-4-as-a-judge as a reliable metric for document-level translation quality.

	AFluency↑	ACE↑	ALE↑	AGE↑
zh-en	0.96	0.95	0.94	0.96
en-zh	0.97	0.98	0.96	0.96
de-en	0.98	0.96	0.94	0.95
en-de	0.96	0.96	0.95	0.97

Table 4: Human agreement percentage on GPT4-as-a-judge with our metrics in WMT22. Each judgment is independently reviewed three times by different annotators and consensus results are recorded. AFluency, ACE, ALE, and AGE denote human agreement on the metrics of Fluency, CE, LE, and GE.

Case Study. To inspect the advantages of LLMs in docMT, we present two pairs of samples from Vicuna-7B and Vicuna-7B-16K(zh-en), covering beginning, middle, and end of each sample.

On the right side of first case in Box 2, the translation of “Hunan” remains consistent throughout the document, illustrating the LLM’s capability to leverage context and capture inter-sentence dependencies. Conversely, on the left side, we see an

Comparison Cases in ST3 and DOC

Model: Vicuna-7B

Prediction in ST3:

The 13th Provincial Tea Expo opened today... This morning, the 13th **Hunan** Tea Industry Expo and... were held in the **Hunan** International Convention...

At the opening ceremony, the provincial leaders awarded the fourth batch of enterprises with the right to use the **Lake South** Red Tea trademark. New Hope held the **Lake** Red Source. The **Chaozhou** tea industry warmly welcomed the arrival of the new spring.

Prediction in DOC:

The 13th **Hunan** Tea Industry Expo opened today at the **Hunan** International Convention...

At the opening ceremony, Hunan’s provincial leaders awarded the fourth batch of enterprises with the right to use the **Hunan** Red Tea trademark... The **Hunan** Red Tea is red in color and has injected the cultural connotation of tea into it, making it popular and lively, and the **Hunan** tea industry is looking forward to a new spring.

Model: Vicuna-7B-16K

Prediction in ST3:

Color: As shown in the picture (**please avoid shooting to avoid color difference**...

Therefore, girls who can’t **drive** should not complain about their clothes being old...

2021.6.11部分圈中售出。看好编号下单，古董物品售出不退不换。购买须知The products sold at this store are non-refundable... **The store does not accept styles that are different from what is imagined**, and size and style cannot be used as reasons for refunds or exchanges...

Prediction in DOC:

Color: As shown in the picture (**Please note that the color difference may not be avoided due to shooting**...

So, some girls who **can’t drive vintage clothing** should not say that the clothes are old-fashioned, but that you are not suitable for it!...

If the item is not suitable for personal reasons, such as not fitting or not liking it, you can ask the store owner to transfer it to the shelf, and **once it is sold, it cannot be exchanged or refunded. Part of the circle in the middle was sold on June 11, 2021**...

Box 2: Comparison of Vicuna-7B and Vicuna-7B-16K translations under ST3 and DOC evaluation types in the en-zh translation direction.

erroneous translation where “Hunan” is rendered as “Lake South” in “Lake South Red Tea” and simply as “Lake” in “Lake Red Source”. Notably, the model in ST3 correctly translates “Hunan” in other parts of the text. In this case, although ST3 achieves a BLEU score approximately 11.88 points higher than DOC, it is evident that DOC provides more coherent wording and aligns better with natural human expression.

We present another case in Box 2: ST3 translates “color difference is inevitable in the photos” as “please avoid shooting to avoid color difference,” resulting in a significant change in meaning. Additionally, a description about some girls’ struggles with a style is mistranslated as “girls who can’t drive” where “drive” is incorrectly used as an intransitive verb. In contrast, DOC accurately translates this as “some girls who can’t drive vintage clothing” preserving the intended meaning while employing the same words in different contexts. Furthermore, the statement “once it is sold, it cannot be exchanged or refunded. Part of the circle in the middle was sold on June 11, 2021” is correctly translated in DOC, while ST3 reject translating this segment entirely. These cases explicitly demonstrate that instruction-tuned LLMs can effectively capture inter-sentence dependencies by considering

the entire document context, leading to a deeper understanding of the text and fewer content errors.

Thus, we advocate against using BLEU as an evaluation metric for docMT, as it fails to detect the true advantages of LLMs in this context and can yield misleading results.

5 Conclusion

In this work, we investigate the performance of instruction-tuned LLMs in document-level machine translation (docMT), comparing the translation of entire documents in a single pass to translating individual sentences that are then concatenated. Our findings show that translating entire documents yields better results, as the model can capture inter-sentence dependencies and maintain discourse coherence, even without explicit fine-tuning for docMT tasks. However, evaluating these improvements is challenging. Traditional metrics like d-BLEU fail to consider discourse-level phenomena, often favoring sentence-level translations and producing misleading results. To address this limitation, we propose the LLM-as-a-judge approach, utilizing GPT-4 to assess specific aspects of discourse through tailored prompts. This method enhances interpretability and can be adapted for evaluating translation quality in other domains.

Limitations

Translation Directions. We evaluate only high-resource language pairs, which limits the generalizability of our findings for low-resource languages. Due to data availability constraints, our experiments focus on well-resourced translation directions. Future research should explore whether instruction-tuned LLMs translating entire documents yield better results than translating sentences independently in low-resource languages.

Model Size and Diversity. We focus exclusively on small-scale LLMs. Future work should investigate larger models to observe whether instruction-tuned LLMs continue to perform better in docMT, and whether BLEU would work.

Max Length. A small fraction ($\sim 2\%$) of documents in WMT22, including both their source texts and translations, exceed 2048 tokens. Thus, we focus solely on samples within the model’s context length (2048 tokens), as these instruction-tuned LLMs are primarily trained on text within this limit. In future work, we will evaluate LLMs with longer context lengths, examine -16K models, and investigate whether long conversation instruction-tuned will help and whether those phenomena persist when translating text that exceeds the models’ context length.

Ethical Considerations

Our study aims to investigate the docMT reliability of instruction-tuned LLMs without fine-tuning for docMT, concerned by the potential for accumulating errors during decoding, which may lead to increased hallucinations. We expect minimal social risks associated with our efforts.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*.
- Menglong Cui, Jiangcun Du, Shaolin Zhu, and Deyi Xiong. 2024. [Efficiently exploring large language models for document-level machine translation with in-context learning](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 10885–10897, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and André F. T. Martins. 2022. [Results of WMT22 metrics shared task: Stop using BLEU – neural metrics are better and more robust](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 46–68, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Zhengxian Gong, Min Zhang, and Guodong Zhou. 2015. [Document-level machine translation evaluation with gist consistency and text cohesion](#). In *Proceedings of the Second Workshop on Discourse in Machine Translation*, pages 33–40, Lisbon, Portugal. Association for Computational Linguistics.
- Michael Alexander Kirkwood Halliday and Ruqaiya Hasan. 2014. *Cohesion in english*. Routledge.
- Christian Hardmeier and Marcello Federico. 2010. [Modelling pronominal anaphora in statistical machine translation](#). In *Proceedings of the 7th International Workshop on Spoken Language Translation: Papers*, pages 283–289, Paris, France.
- Amr Hendy, Mohamed Abdelrehim, Amr Sharaf, Vikas Raunak, Mohamed Gabr, Hitokazu Matsushita, Young Jin Kim, Mohamed Afify, and Hany Hassan Awadalla. 2023. [How good are GPT models at machine translation? A comprehensive evaluation](#). *CoRR*, abs/2302.09210.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *CoRR*, abs/2310.06825.
- Yuchen Eleanor Jiang, Tianyu Liu, Shuming Ma, Dongdong Zhang, Jian Yang, Haoyang Huang, Rico Senrich, Ryan Cotterell, Mrinmaya Sachan, and Ming Zhou. 2021. [Blonde: An automatic evaluation metric for document-level machine translation](#). *arXiv preprint arXiv:2103.11878*.
- Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, Shuming Shi, and Zhaopeng Tu. 2023. [Is chatgpt a good translator? yes with gpt-4 as the engine](#). *arXiv preprint arXiv:2301.08745*.
- Prathyusha Jwalapuram, Shafiq Joty, Irina Temnikova, and Preslav Nakov. 2019. [Evaluating pronominal anaphora in machine translation: An evaluation measure and a test suite](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language*

- Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2964–2975, Hong Kong, China. Association for Computational Linguistics.
- Tom Kocmi, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Thamme Gowda, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Rebecca Knowles, Philipp Koehn, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Michal Novák, Martin Popel, and Maja Popović. 2022. [Findings of the 2022 conference on machine translation \(WMT22\)](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 1–45, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Yachao Li, Junhui Li, Jing Jiang, and Min Zhang. 2024. Enhancing document-level translation of large language model via translation mixed-instructions. *arXiv preprint arXiv:2401.08088*.
- 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.
- Sameen Maruf, Fahimeh Saleh, and Gholamreza Haffari. 2021. A survey on document-level neural machine translation: Methods and evaluation. *ACM Computing Surveys (CSUR)*, 54(2):1–36.
- 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 on Association for Computational Linguistics*, ACL '02, page 311–318, USA. Association for Computational Linguistics.
- Keqin Peng, Liang Ding, Qihuang Zhong, Li Shen, Xuebo Liu, Min Zhang, Yuanxin Ouyang, and Dacheng Tao. 2023. [Towards making the most of chatgpt for machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 5622–5633. Association for Computational Linguistics.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners](#). OpenAI blog.
- 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 2020, Online, November 16-20, 2020*, pages 2685–2702. Association for Computational Linguistics.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 223–231.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.
- Longyue Wang, Chenyang Lyu, Tianbo Ji, Zhirui Zhang, Dian Yu, Shuming Shi, and Zhaopeng Tu. 2023. [Document-level machine translation with large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 16646–16661, Singapore. Association for Computational Linguistics.
- Minghao Wu, Yufei Wang, George Foster, Lizhen Qu, and Gholamreza Haffari. 2024. [Importance-aware data augmentation for document-level neural machine translation](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 740–752, St. Julian’s, Malta. Association for Computational Linguistics.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 46595–46623. Curran Associates, Inc.
- Dawei Zhu, Pinzhen Chen, Miaoran Zhang, Barry Haddow, Xiaoyu Shen, and Dietrich Klakow. 2024a. Fine-tuning large language models to translate: Will a touch of noisy data in misaligned languages suffice? *arXiv preprint arXiv:2404.14122*.

Dawei Zhu, Sony Trenous, Xiaoyu Shen, Dietrich Klakow, Bill Byrne, and Eva Hasler. 2024b. [A preference-driven paradigm for enhanced translation with large language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3385–3403, Mexico City, Mexico. Association for Computational Linguistics.

A Evaluation Metrics Shortcoming Analysis

While COMET has been shown to provide more reliable evaluations than BLEU in many cases, it is primarily trained on sentence-level translations and, as such, is not well-suited for docMT. Given that COMET lacks specific training to capture the complexities of inter-sentence dependencies and discourse-level phenomena, it is not an ideal metric for evaluating the true capabilities of LLMs in docMT tasks. Therefore, in this work, we opted to explore more appropriate evaluation methods tailored to document-level translation challenges.

Similarly, metrics like ChrF (Popović, 2015), ChrF2, and TER have made incremental progress by incorporating word-level matching mechanisms that extend beyond simple token overlap, but they still fundamentally rely on surface-level statistics. Like BLEU, these metrics do not adequately account for deeper discourse relationships, cohesion, and the broader context required for accurate docMT assessment. As a result, their limitations become more apparent when evaluating LLMs on longer texts, where capturing the overall document structure is essential.

While metrics such as CTT and AZPT are designed to address specific issues like terminology consistency and zero pronoun accuracy, they remain grounded in automatic identified lexical alignment. These metrics operate under the assumption that the presence of specific terminology or pronouns directly correlates with translation quality. However, in practice, meaning can be conveyed in multiple ways without strictly adhering to these surface-level features. This makes CTT and AZPT limited in scope, as they are unable to fully assess translation quality when alternative phrasing or omitted pronouns still preserve meaning accurately.

Blonde represents a more sophisticated approach by categorizing and analyzing discourse coherence using linguistic features such as verb tense (e.g., VBD for past tense verbs). While this is a step to-

ward capturing discourse-level phenomena, Blonde is still constrained by symbolic statistical methods. Its reliance on predefined linguistic categories means that it struggles to account for the full range of discourse phenomena that can arise in real-world documents. As a result, these metrics, despite their improvements, remain insufficient for capturing the nuances of document-level translation in its entirety.

To address these limitations, we propose leveraging LLM-as-a-judge for evaluating docMT. By employing GPT-4 with specifically designed judging prompts, we can define and assess discourse phenomena in a more abstract and flexible manner, similar to how human evaluators would approach the task. This method avoids the need to predefine all possible linguistic cases and allows for a more holistic evaluation of translation quality, ensuring that complex discourse relationships and contextual dependencies are properly recognized. In doing so, we provide more reliable and interpretable metrics and prompts for evaluating document-level translations, moving beyond the restrictive frameworks of traditional metrics.

B d-BLEU Performance

We observe that the trend in Table 5 remains consistent with Table 2, and the BLEU score shows an even stronger preference for translations that are processed separately and concatenated. It is worth to Notice that the red data point in Table 5 is influenced by the sensitivity of BLEU, where a certain generated translation contains a long-repeated incorrect token toward the end, thus lowering the overall score. When calculating the BLEU score for this sample, we find that the document receives a score near zero, despite the fact that the earlier part of the translation is mostly accurate. This sensitivity is one of the reasons why BLEU should not be used in docMT.

C GPT4-as-a-judge Evaluation Prompts

C.1 Fluency

Fluency refers to the naturalness and smoothness of a text in the target language, without awkward or unnatural phrasing. In machine translation evaluation, fluency is crucial for assessing the readability and linguistic quality of the output, which is often not fully captured by traditional metrics like BLEU. While BLEU focuses on n-gram overlap between the translation and reference text, it does not di-

Model	Eval Type	Translation Direction			
		zh-en	en-zh	de-en	en-de
Vicuna-7B	ST1	18.75	32.43	30.00	21.96
	ST2	19.99	33.52	30.87	23.35
	ST3	20.52	33.92	30.68	23.96
	DOC	19.93	32.40	30.27	22.90
Vicuna-7B-16K	ST1	19.54	28.45	29.75	21.49
	ST2	20.38	32.52	30.60	24.27
	ST3	20.43	33.15	30.56	23.95
	DOC	19.50	30.87	16.58	21.34
Vicuna-13B	ST1	21.33	37.62	31.98	26.24
	ST2	21.26	37.70	32.16	27.19
	ST3	21.89	37.97	32.15	26.94
	DOC	21.63	35.87	31.22	26.31
Vicuna-13B-16K	ST1	21.22	37.29	31.48	26.21
	ST2	21.99	37.93	31.78	26.53
	ST3	22.61	37.83	32.02	26.98
	DOC	21.84	35.01	31.60	26.03
Mistral-7B	ST1	18.69	25.75	29.50	22.37
	ST2	19.29	26.83	30.02	21.98
	ST3	18.82	26.81	30.11	22.60
	DOC	13.70	17.54	27.50	21.98

Table 5: d-BLEU score with different translation paradigms. More explanations about not using d-BLEU and about the red data point in the Table are stated in Appendix B

rectly evaluate how natural the translation sounds or whether it adheres to syntactic rules. Fluency, in contrast, provides a more nuanced evaluation of the model’s ability to produce human-like text.

In this task, we assess fluency on a scale of 1 to 5, with higher scores indicating more fluent translations. Evaluators are instructed to analyze the text and assign a score based solely on the naturalness and grammatical correctness of the model’s output.

Importantly, the fluency evaluation is conducted in isolation, decoupled from cohesion, with only inference text input, to ensure a clear focus on the text’s immediate readability. Cohesion, which refers to the grammatical and lexical connectivity between text units (Halliday and Hasan, 2014), is considered separately to avoid confounding the two metrics, as fluency and cohesion could be correlated, as it is common sense that if a text is cohesive, its flow is naturally better. See the correlation heatmaps like Figure 1 which show that our prompt design successfully decouples these two metrics.

The evaluation is supported by specific examples and justifications for the assigned score. Below is

the prompt used to guide the evaluation:

Please evaluate the fluency of the following text in the target language (English, Chinese, or German).

Instructions:

- **Task:** Evaluate the fluency of the text.
- **Scoring:** Provide a score from 1 to 5, where:
 - **5:** The text is highly fluent, with no grammatical errors, unnatural wording, or stiff syntax.
 - **4:** The text is mostly fluent, with minor errors that do not impede understanding.
 - **3:** The text is moderately fluent, with noticeable errors that may slightly affect comprehension.
 - **2:** The text has low fluency, with frequent errors that hinder understanding.

- **1:** The text is not fluent, with severe errors that make it difficult to understand.
- **Explanation:** Support your score with specific examples to justify your evaluation.

Output Format:

Provide your evaluation in the following JSON format:

```
{ "Fluency": { "Score":
"<the score>", "Explanation":
"<your explanation on how
you made the decision>" } }
```

Text to Evaluate:

"inference text"

C.2 Content Errors

Unlike fluency, which assesses the naturalness and grammatical correctness of the output, accuracy focuses on the semantic alignment between the translated text and the original reference. The evaluator’s task is to identify and categorize errors that affect the translation’s fidelity, such as mistranslations, omissions, or additions.

Rather than relying on simple n-gram matching, the evaluation emphasizes meaning preservation. The evaluator compares the translation with the reference text, identifying instances where the translation deviates in meaning. However, if the translated text conveys the same information as the reference but uses different words or phrasing, it is not considered an error, since we suspect that this phenomenon could happen in LLMs in document-level translation task. This approach ensures that the model’s output is evaluated based on its ability to faithfully represent the source content, capturing specific issues like mistranslations or information loss, and ensuring semantic integrity. The accuracy evaluation prompt is structured as follows:

Please evaluate the accuracy of the following text by comparing it to the reference text provided.

Instructions:

- **Task:** Compare the text to the reference text.
- **Identify Mistakes:** List all mistakes related to accuracy.
 - **Mistake Types:**

- * **Wrong Translation:** Incorrect meaning or misinterpretation leading to wrong information.
- * **Omission:** Missing words, phrases, or information present in the reference text.
- * **Addition:** Extra words, phrases, or information not present in the reference text.
- * **Others:** Mistakes that are hard to define or categorize.

- **Note:** If the text expresses the same information as the reference text but uses different words or phrasing, it is **not** considered a mistake.
- **Provide a List:** Summarize all mistakes without repeating the exact sentences. Provide an empty list if there are no mistakes.

Output Format:

```
{ "Accuracy": { "Mistakes": [ "<list of
all mistakes in the text, provide an empty
list if there are no mistakes>" ] }
```

Text to Evaluate:

"inference text"

C.3 Cohesion Errors

Cohesion is a critical aspect of machine translation evaluation as it ensures that the various parts of the text are well-connected and that the overall flow is logical. Unlike metrics such as fluency or accuracy, cohesion specifically examines how sentences are linked together through lexical (lexical cohesion) and grammatical (grammatical cohesion) means (Maruf et al., 2021). This is particularly important in document-level translation, where the consistency of vocabulary and the logical connection of grammatical structures across a longer text are challenging for models to maintain.

In the context of translations produced using the ST3 and DOC paradigms, evaluating cohesion allows us to assess whether the model effectively leverages contextual information to maintain consistency across the text. By decoupling cohesion from fluency, our evaluation framework enables evaluators to focus specifically on identifying lexical cohesion mistakes—such as incorrect vocabulary usage, missing synonyms, or overuse of certain terms that disrupt the flow—and grammatical

cohesion mistakes—such as errors in pronouns, conjunctions, or sentence-linking structures.

The evaluator is asked to identify any mistakes related to cohesion and categorize them as either lexical or grammatical cohesion issues. The evaluation prompt is structured as follows:

Please evaluate the cohesion of the following text by comparing it to the reference text.

Instructions:

- **Task:** Evaluate the cohesion of the text.
- **Definition:** Cohesion refers to how different parts of a text are connected using language structures like grammar and vocabulary. It ensures that sentences flow smoothly and the text makes sense as a whole.
- **Identify Mistakes:** List all mistakes related to cohesion.
 - **Lexical Cohesion Mistakes:** Issues with vocabulary usage, incorrect or missing synonyms, or overuse of certain words that disrupt the flow.
 - **Grammatical Cohesion Mistakes:** Problems with pronouns, conjunctions, or grammatical structures that link sentences and clauses.
- **Provide Lists:** Provide separate lists for lexical cohesion mistakes and grammatical cohesion mistakes. Provide empty lists if there are no mistakes.

Output Format:

```
{ "Cohesion": { "Lexical Cohesion Mistakes": [ "<list of all mistakes in the text, provide an empty list if there are no mistakes>" ], "Grammatical Cohesion Mistakes": [ "<list of all mistakes in the text, provide an empty list if there are no mistakes>" ] } }
```

Text to Evaluate:

"inference text"

D GPT4-as-a-judge Evaluation Performance

Model	Eval Type	AvgBLEU↑	Fluency↑	CE↓	LE↓	GE↓
Vicuna-7B	ST3	20.25	4.27	3.06	1.46	0.96
	DOC	19.36	4.20	3.24	1.40	0.67
Vicuna-7B-16K	ST3	20.07	4.24	3.50	1.20	0.77
	DOC	21.14	4.38	3.24	1.18	0.67
Vicuna-13B	ST3	22.46	4.09	3.31	1.57	1.07
	DOC	23.46	4.34	3.04	1.04	0.59
Vicuna-13B-16K	ST3	21.80	4.21	3.27	0.90	0.65
	DOC	22.22	4.48	2.82	0.80	0.41
Mistral-7B	ST3	18.84	3.96	4.41	1.47	1.10
	DOC	19.50	4.34	3.50	1.30	0.83

Table 6: Evaluation results (zh-en) by GPT-4 for Vicuna-7B, Vicuna-13B, and Mistral-7B under ST3 and DOC evaluation types, showing metrics AvgBLEU, Fluency, Content Errors(CE), Lexical Cohesion Errors(LE), and Grammatical Cohesion Errors(GE).

Model	Eval Type	AvgBLEU↑	Fluency↑	CE↓	LE↓	GE↓
Vicuna-7B	ST3	24.53	3.23	7.31	3.73	3.00
	DOC	21.18	3.61	6.46	3.76	2.87
Vicuna-7B-16K	ST3	26.02	4.21	2.88	1.17	0.77
	DOC	26.95	4.11	2.84	0.98	0.67
Vicuna-13B	ST3	26.76	3.59	6.84	3.79	2.39
	DOC	27.32	3.90	5.23	3.34	1.96
Vicuna-13B-16K	ST3	28.15	4.32	2.67	0.78	0.42
	DOC	28.54	4.45	2.28	0.95	0.45
Mistral-7B	ST3	22.32	3.16	6.83	4.64	2.93
	DOC	21.46	3.17	6.63	4.51	2.96

Table 7: Evaluation results (en-de) by GPT-4 for Vicuna-7B, Vicuna-13B, and Mistral-7B under ST3 and DOC evaluation types, showing metrics AvgBLEU, Fluency, Content Errors(CE), Lexical Cohesion Errors(LE), and Grammatical Cohesion Errors(GE).

Model	Eval Type	AvgBLEU \uparrow	Fluency \uparrow	CE \downarrow	LE \downarrow	GE \downarrow
Vicuna-7B	ST3	26.82	4.11	4.01	1.23	1.11
	DOC	25.64	4.31	3.14	1.67	0.66
Vicuna-7B-16k	ST3	23.56	3.61	5.74	3.52	2.52
	DOC	21.71	3.54	5.81	3.47	2.21
Vicuna-13B	ST3	27.23	4.30	3.06	1.13	0.66
	DOC	28.44	4.33	3.36	1.33	0.60
Vicuna-13B-16K	ST3	26.55	4.15	5.47	2.72	1.91
	DOC	26.28	4.18	4.7	2.91	1.92
Mistral-7B	ST3	26.09	4.10	4.73	1.49	1.37
	DOC	25.68	4.33	4.89	1.26	0.80

Table 8: Evaluation results (de-en) by GPT-4 for Vicuna-7B, Vicuna-13B, and Mistral-7B under ST3 and DOC evaluation types, showing metrics AvgBLEU, Fluency, Content Errors(CE), Lexical Cohesion Errors(LE), and Grammatical Cohesion Errors(GE).

E Correlation Visualizations

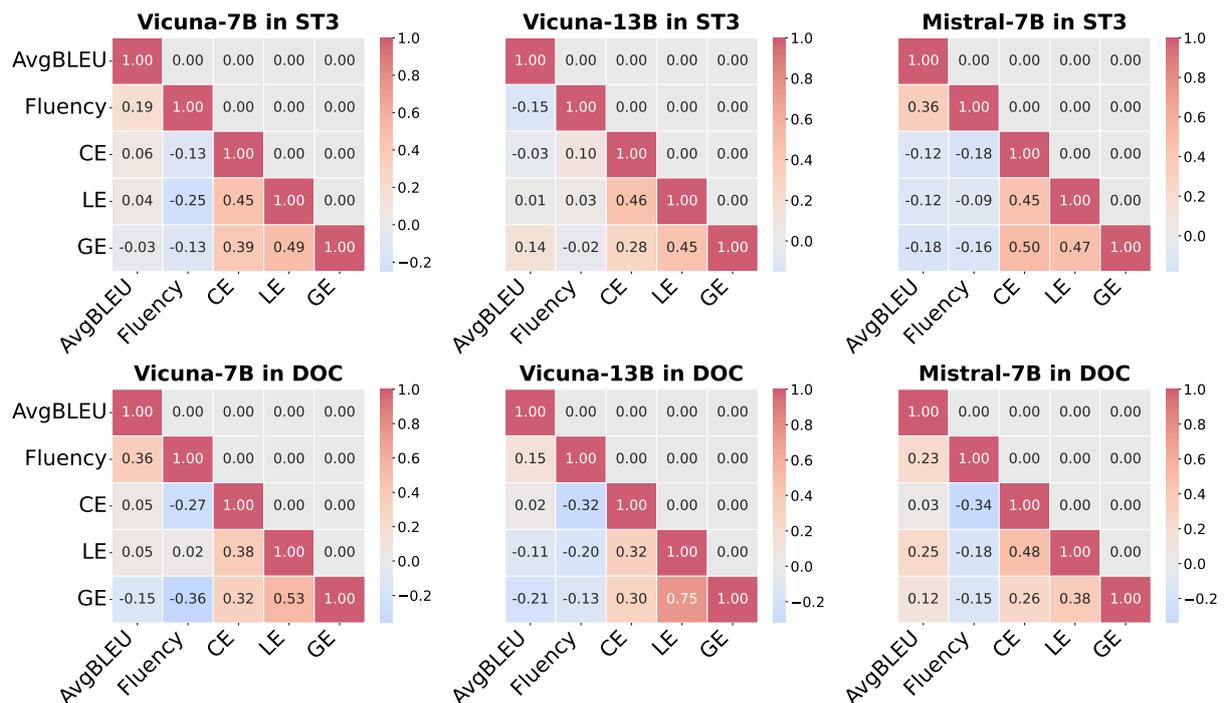


Figure 3: PCC Heatmaps among AvgBLEU, Fluency, CE, LE, GE for Vicuna-7B, Vicuna-13B, and Mistral-7B under ST3 and DOC evaluation types in translation direction of en-zh.

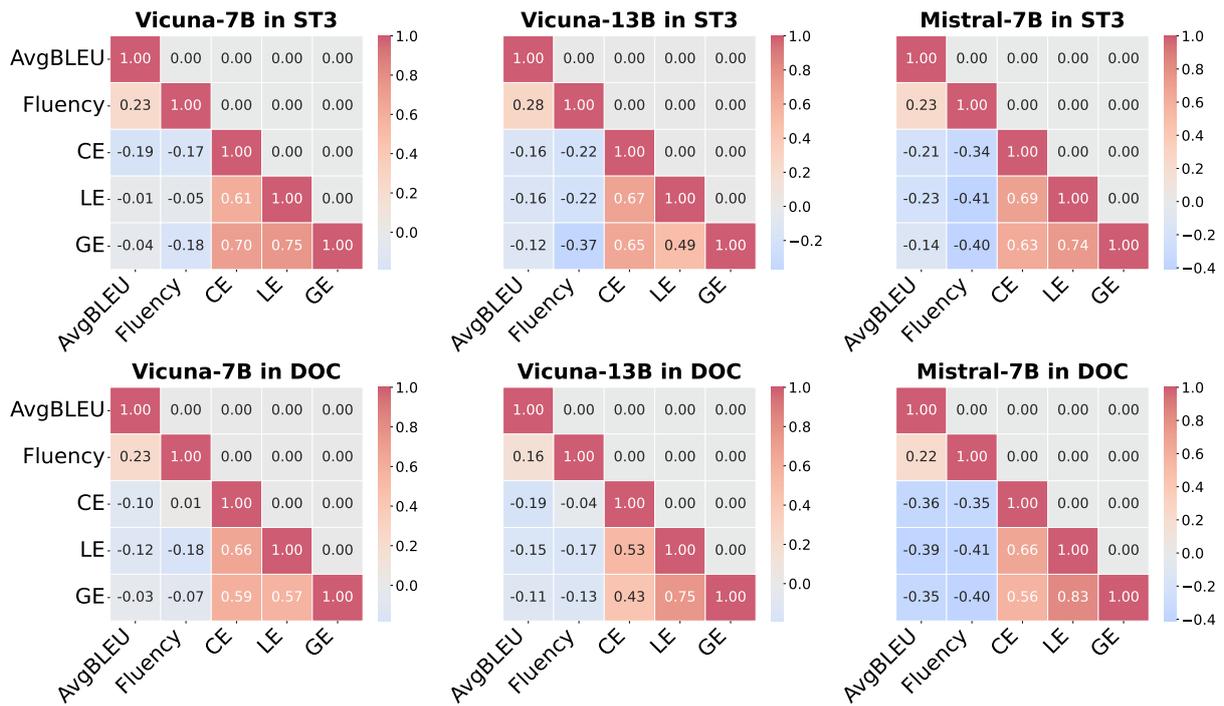


Figure 4: PCC Heatmaps among AvgBLEU, Fluency, CE(Content Errors), LE(Lexical Cohesion errors), GE(Grammatical Cohesion Errors) for Vicuna-7B, Vicuna-13B, and Mistral-7B under ST3 and DOC evaluation types in translation direction of zh-en.

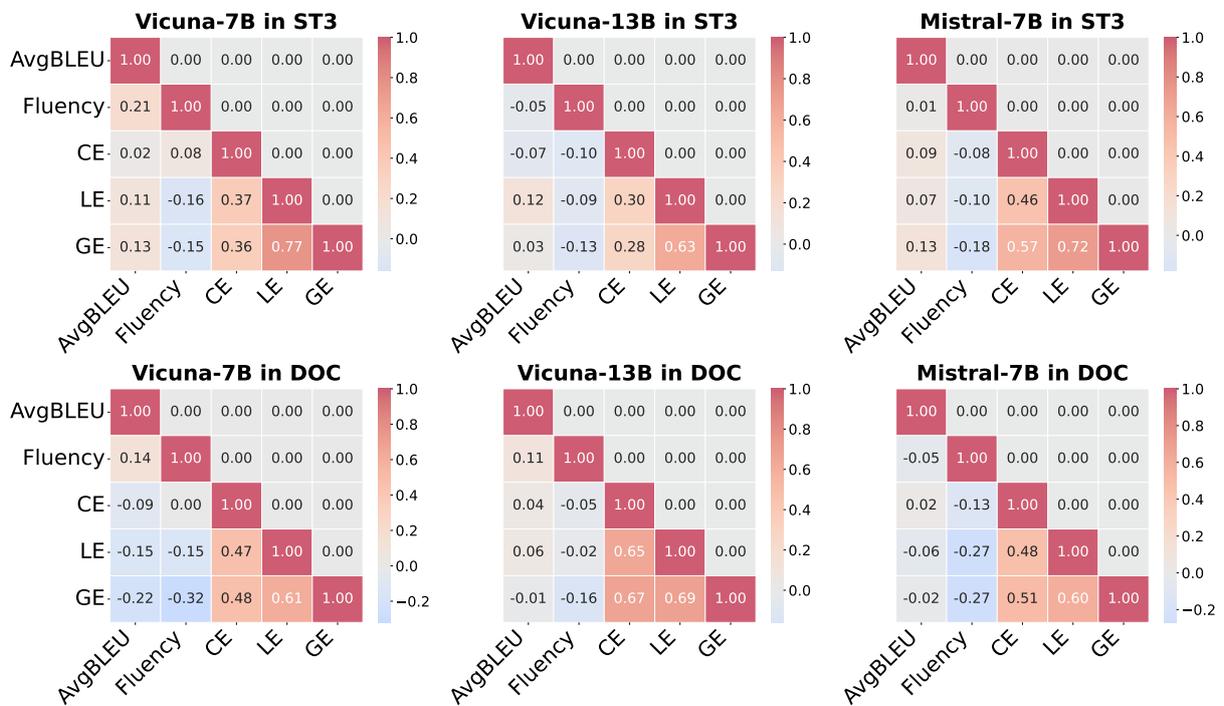


Figure 5: PCC Heatmaps among AvgBLEU, Fluency, CE(Content Errors), LE(Lexical Cohesion errors), GE(Grammatical Cohesion Errors) for Vicuna-7B, Vicuna-13B, and Mistral-7B under ST3 and DOC evaluation types in translation direction of de-en.

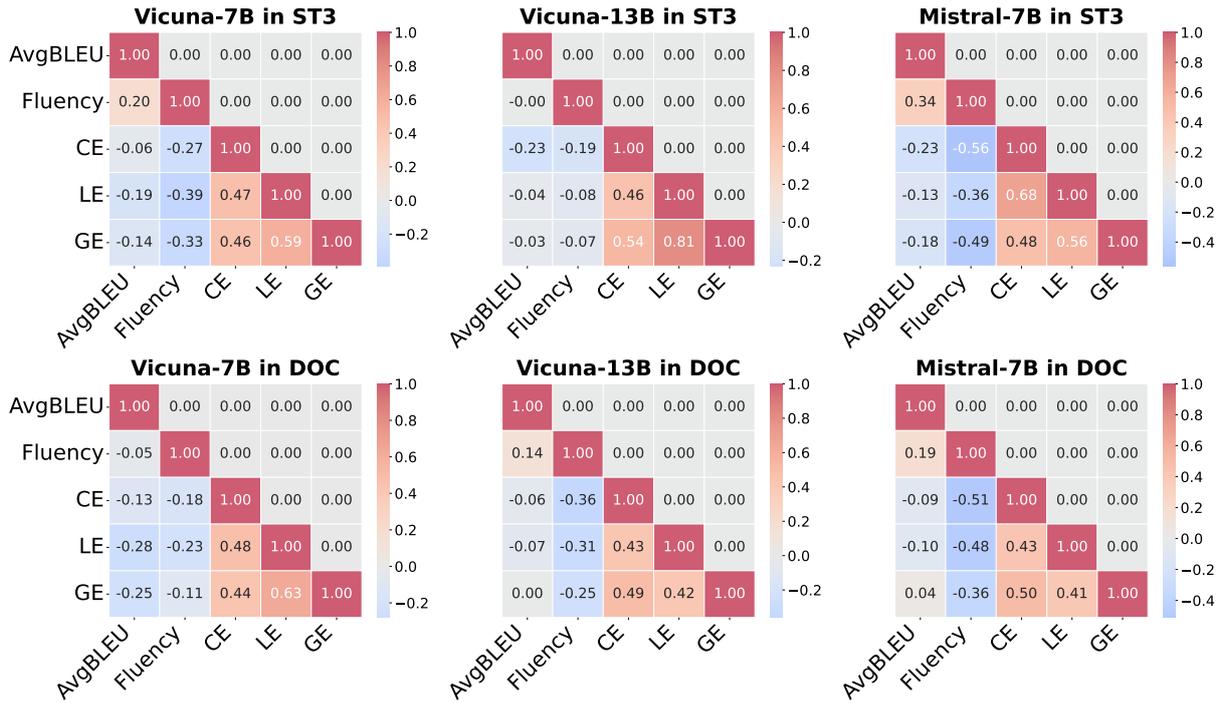


Figure 6: PCC Heatmaps among AvgBLEU, Fluency, CE(Content Errors), LE(Lexical Cohesion errors), GE(Grammatical Cohesion Errors) for Vicuna-7B, Vicuna-13B, and Mistral-7B under ST3 and DOC evaluation types in translation direction of en-de.

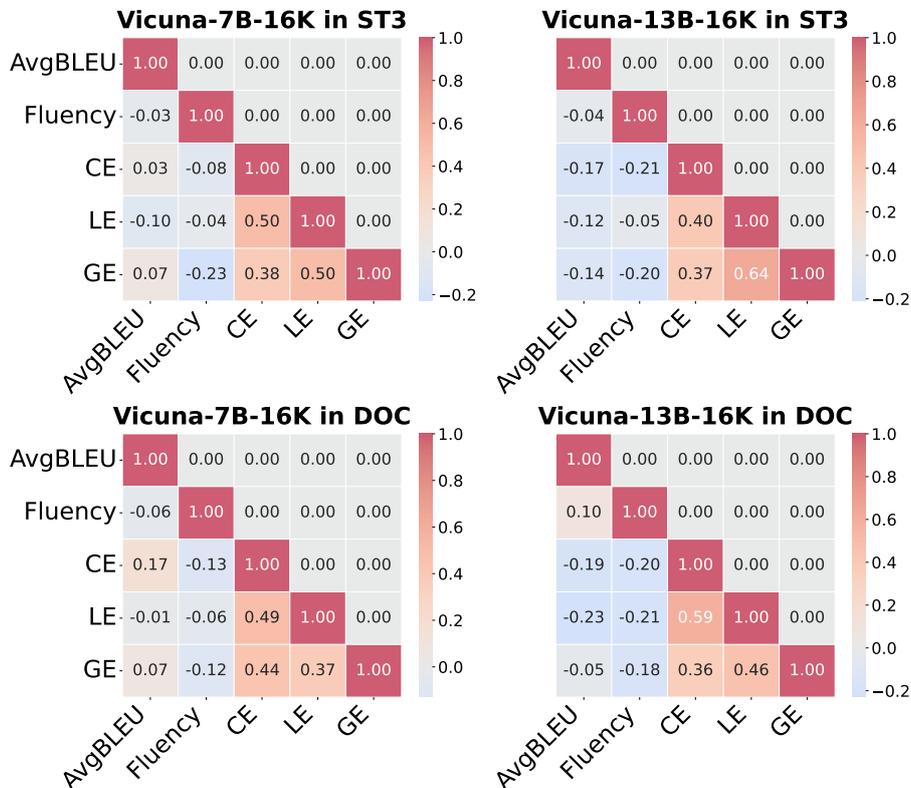


Figure 7: PCC Heatmaps among AvgBLEU, Fluency, CE, LE, GE for Vicuna-7B-16K and Vicuna-13B-16K under ST3 and DOC evaluation types in translation direction of en-zh.

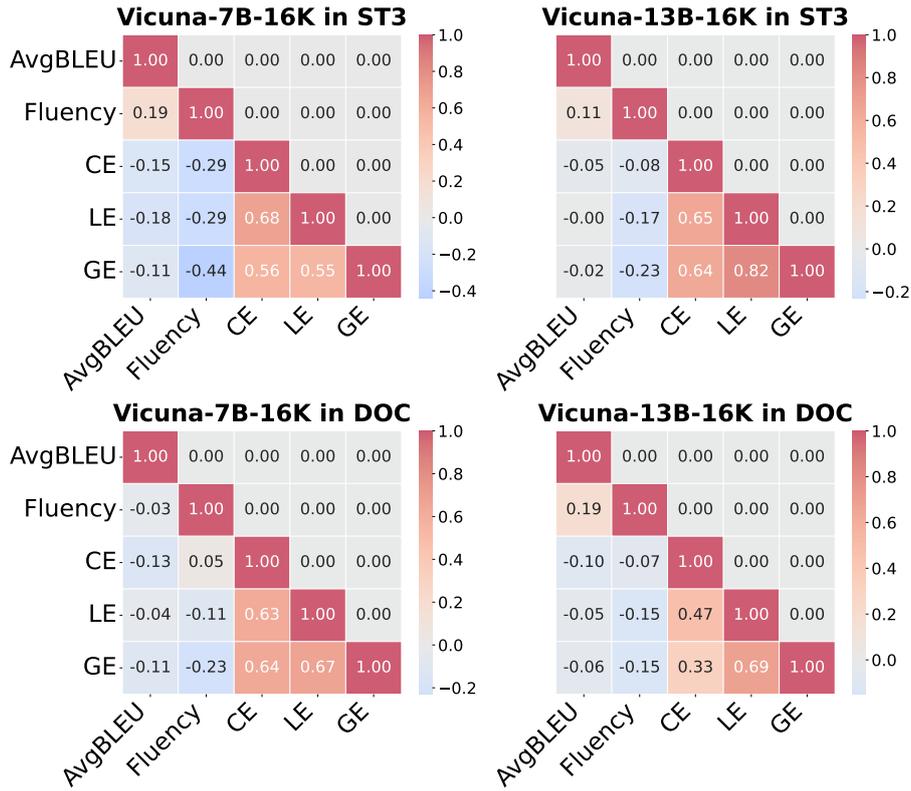


Figure 8: PCC Heatmaps among AvgBLEU, Fluency, CE(Content Errors), LE(Lexical Cohesion errors), GE(Grammatical Cohesion Errors) for Vicuna-7B-16K and Vicuna-13B-16K under ST3 and DOC evaluation types in translation direction of zh-en.

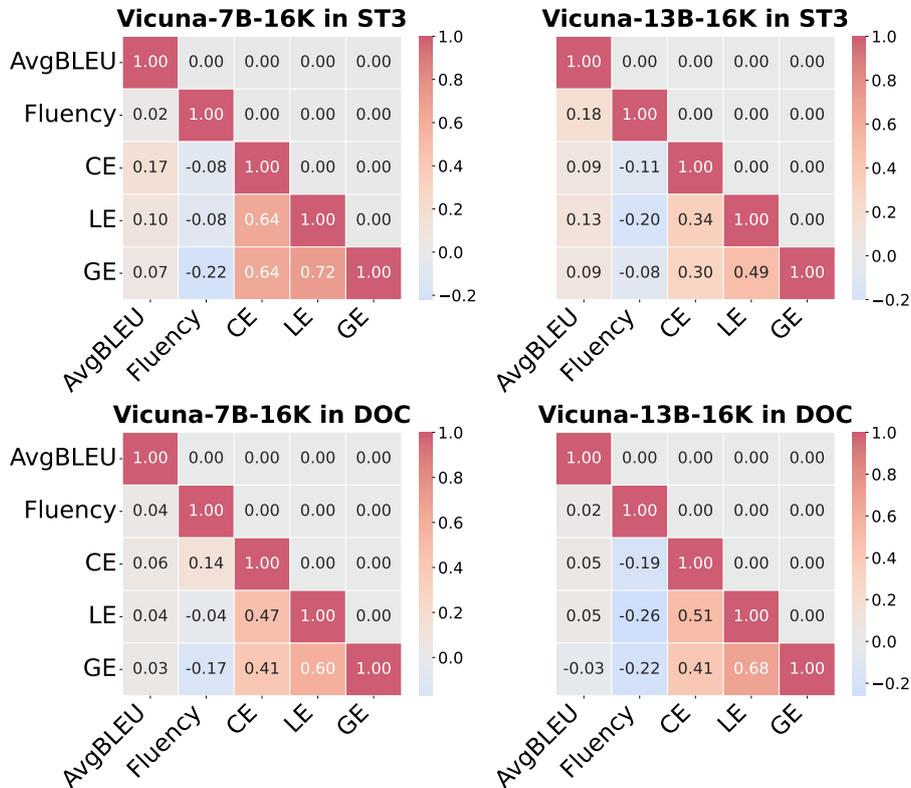


Figure 9: PCC Heatmaps among AvgBLEU, Fluency, CE(Content Errors), LE(Lexical Cohesion errors), GE(Grammatical Cohesion Errors) for Vicuna-7B-16K and Vicuna-13B-16K under ST3 and DOC evaluation types in translation direction of de-en.

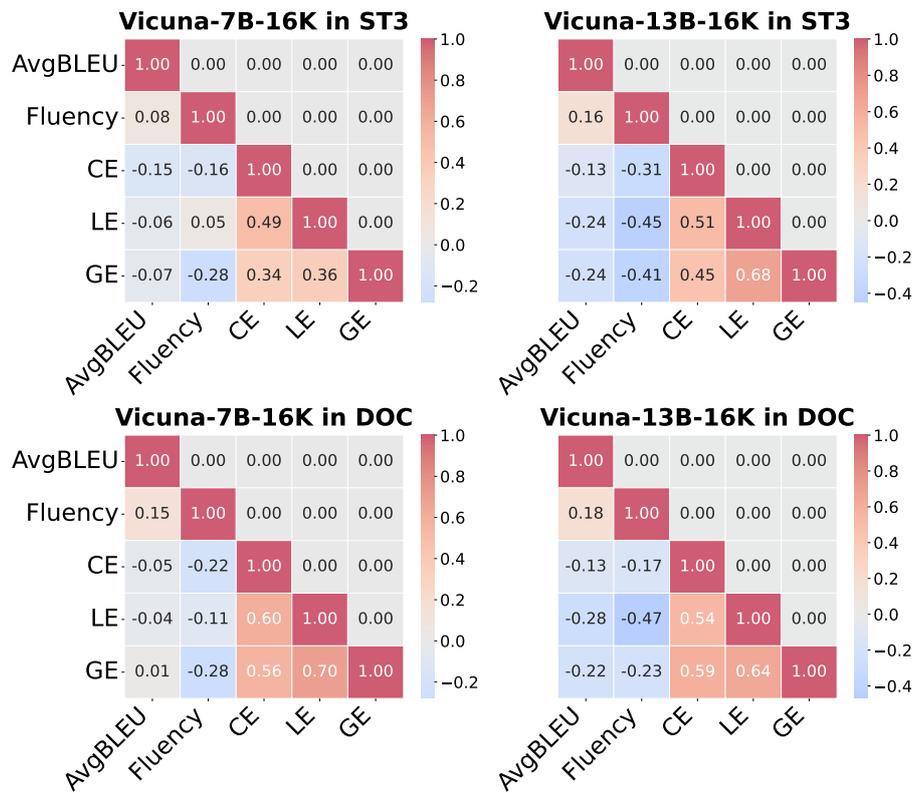


Figure 10: PCC Heatmaps among AvgBLEU, Fluency, CE(Content Errors), LE(Lexical Cohesion errors), GE(Grammatical Cohesion Errors) for Vicuna-7B-16K and Vicuna-13B-16K under ST3 and DOC evaluation types in translation direction of en-de.