

Augmenting Large Language Model Translators via Translation Memories

Yongyu Mu^{1*}, Abudurexiti Reheman^{1*}, Zhiquan Cao¹, Yuchun Fan¹,
Bei Li¹, Yinqiao Li¹, Tong Xiao^{1,2†}, Chunliang Zhang^{1,2} and Jingbo Zhu^{1,2}

¹NLP Lab, School of Computer Science and Engineering,
Northeastern University, Shenyang, China

²NiuTrans Research, Shenyang, China

lixiaoyumu9@gmail.com rexiti_neu@outlook.com

{xiaotong, zhujingbo}@mail.neu.edu.cn

Abstract

Using translation memories (TMs) as prompts is a promising approach to in-context learning of machine translation models. In this work, we take a step towards prompting large language models (LLMs) with TMs and making them better translators. We find that the ability of LLMs to “understand” prompts is indeed helpful for making better use of TMs. Experiments show that the results of a pre-trained LLM translator can be greatly improved by using high-quality TM-based prompts. These results are even comparable to those of the state-of-the-art NMT systems which have access to large-scale in-domain bilingual data and are well tuned on the downstream tasks.

1 Introduction

Marrying the world of translation memory (TM) and the world of neural machine translation (NMT) is a challenging but interesting problem in natural language processing (NLP). Previous work along this line of research either requires architecture changes of NMT models and/or additional training (Gu et al., 2018; Bulté and Tezcan, 2019; Xu et al., 2020; Hossain et al., 2020; He et al., 2021) or constructing translation knowledge base from TM (Zhang et al., 2018; Khandelwal et al., 2021; Meng et al., 2022).

More recently, researchers have been aware of the strength of prompting techniques for one-shot/few-shot machine translation (Vilar et al., 2022; Agrawal et al., 2022; Zhang et al., 2023). In particular, Reheman et al. (2023) investigated one-shot learning methods for NMT by simply viewing TMs as prompts. The result of their work is a stronger NMT system that works in the same way as usual but can be prompted when TMs are available. Interestingly, they found that the ability of NMT models to “understand” prompts plays an

*Equal contribution.

†Corresponding author.

Method	w/o-arch-change	w/o-base	few-shot
Zhang et al. (2018)	yes		
He et al. (2021)		yes	
Khandelwal et al. (2021)	yes		
Reheman et al. (2023)	yes	yes	one-shot
TMPLM (our)	yes	yes	yes

Table 1: Methods of using TM for better MT. w/o-arch-change = without architecture changes or training, w/o-base = without constructing translation knowledge base from TM, and few-shot = few-shot learning.

important role in this type of system. Prompts are still difficult to use if NMT systems are weak.

In this work, we take a step forward. We treat large language models (LLMs) as machine translation systems and prompt them with TMs (see Table 1 for a comparison of different methods). This is in part motivated by recent developments of LLMs: one of the most powerful properties of LLMs is their ability to understand and respond to complex instructions and questions (Ouyang et al., 2022; Thoppilan et al., 2022). We show that this ability is crucial for in-context learning of TM-based prompts, and LLM-based translation systems can be greatly improved by using simple instruction-like prompts. To this end, we propose **Translation Memory Prompting for large Language Models**, namely **TMPLM** - a simple but effective approach to injecting TMs into LLM translators.

We experiment with our method on a GPT-based LLM (text-davinci-003*). On translation tasks ranging over multiple languages and domains, TM-based prompting improves the LLM-based translation system by 20 to 30 BLEU points, showing better performance than a well-tuned, large-scale, in-domain NMT system on most of the tasks. We also compare different kinds of prompt templates and discuss some interesting issues, such as the role of prompting in treating LLMs as translators.

*We will refer to it as davinci-003 later in the paper.

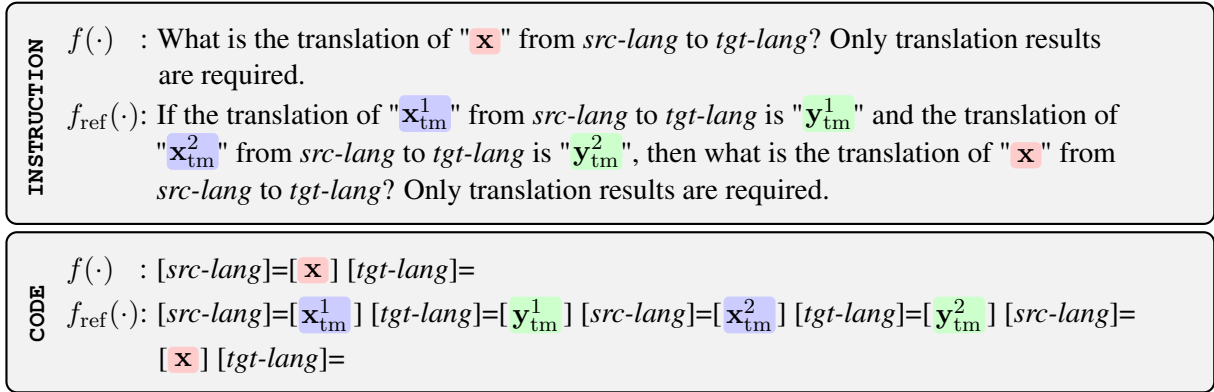


Figure 1: Two styles of template. $f(\cdot)$ denotes a template by which we represent the input sentence as the input of the translation model (such as LLM in this figure). $f_{\text{ref}}(\cdot)$ is a new template involving outputs of a TM ($k = 2$ in this example). \mathbf{x} in red stands for the sentence that needs to be translated. \mathbf{x}_{tm} in blue and \mathbf{y}_{tm} in green stand for the source and target sentence found in the TM, respectively. Both *src-lang* and *tgt-lang* need to be replaced by the names of the source and target language.

2 Prompting Methods

TM is a database that contains the bilingual translation history of professional translators. It is usually used to help the translation of the test sentence by providing similar sentence pairs, which may have translation hints, such as similar sentence patterns, phrases, lexicons, terminologies, or other translation knowledge. Either an NMT model or an LLM need to *dig out* those hints and ignore the irrelevant content. This motivates us to investigate prompting LLMs with TMs benefiting from their dazzling ability of “understand” prompts.

Suppose we have a TM database that retains a collection of pairs of sentences. Given a source-language sentence \mathbf{x} , the database returns k most similar sentences $\mathbf{X}_{\text{tm}} = \{\mathbf{x}_{\text{tm}}^1, \dots, \mathbf{x}_{\text{tm}}^k\}$ along with their corresponding translations $\mathbf{Y}_{\text{tm}} = \{\mathbf{y}_{\text{tm}}^1, \dots, \mathbf{y}_{\text{tm}}^k\}$. Now suppose we have a pre-trained translation model (either an NMT model or an LLM) that takes \mathbf{x} in some way and outputs a translation \mathbf{y} , written as

$$\mathbf{y} = \text{Trans}(f(\mathbf{x})) \quad (1)$$

where $\text{Trans}(\cdot)$ denotes the translation model, and $f(\cdot)$ denotes a template by which we represent \mathbf{x} as the input of $\text{Trans}(\cdot)$. For example, if $\text{Trans}(\cdot)$ is an NMT model, $f(\mathbf{x}) = \mathbf{x}$; if $\text{Trans}(\cdot)$ is a generative LLM, $f(\mathbf{x})$ could be an instruction involving \mathbf{x} .

We then wish to use this model to generate a new translation \mathbf{y}' by considering $(\mathbf{X}_{\text{tm}}, \mathbf{Y}_{\text{tm}})$ as instances for reference. This can be written as

$$\mathbf{y}' = \text{Trans}(f_{\text{ref}}(\mathbf{x}, \mathbf{X}_{\text{tm}}, \mathbf{Y}_{\text{tm}})) \quad (2)$$

Here $f_{\text{ref}}(\mathbf{x}, \mathbf{X}_{\text{tm}}, \mathbf{Y}_{\text{tm}})$ is a new template involving $(\mathbf{X}_{\text{tm}}, \mathbf{Y}_{\text{tm}})$.

In this work, we focus on the case in which a powerful generative LLM (such as ChatGPT) is used to perform translation. The input of $\text{Trans}(\cdot)$ could be an instruction or question-like text, and so we can design $f_{\text{ref}}(\cdot)$ in many different ways. In Figure 1, we present two types of templates: the instruction-style template and the code-style template. These designs come from a consideration of the human instruction tuning and the code training used in developing *davinci-003*. For a more extensive discussion of template design, see Appendix B.2.

It is worth emphasizing that, while we restrict ourselves to TM-based prompts in experiments, we can apply this general approach to deal with other knowledge about translation. As a simple example, we can extend $(\mathbf{X}_{\text{tm}}, \mathbf{Y}_{\text{tm}})$ to term or phrase translations. Also, when some MT systems are available, we can make use of automatic translations from other systems to define prompts.

3 Experiments

3.1 Data and LLM Setup

We tested our method (denoted by TMPLM) on three widely-used datasets of TM: DGT-TM (Steinberger et al., 2012), JRC-Acquis (JRC-A) (Steinberger et al., 2006) and the multi-domain dataset described in (Aharoni and Goldberg, 2020). To ensure a fair comparison, we adopted the same preprocessing steps as outlined in Reheman et al. (2023) for data cleanup and training/testing data split.

Data		WMT19 200M		WMT21 4B		davinci-003 175B		
		NMT	NMT+TM	NMT	NMT+TM	LLM (zero-shot)	LLM+TM (one-shot)	LLM+TM (few-shot)
DGT-TM	de → en	45.40	54.03(+8.63)	51.62	69.39(+17.77)	38.89	66.90(+28.01)	69.99 (+31.10)
	en → de	39.03	44.77(+5.74)	42.48	60.09(+17.61)	29.00	57.39(+28.39)	62.02 (+33.02)
JRC-A	de → en	45.90	50.95(+5.05)	51.72	62.99(+11.27)	40.75	62.23(+21.48)	65.55 (+24.80)
	en → de	40.10	43.41(+3.31)	41.71	56.21(+14.50)	29.83	55.01(+25.18)	57.30 (+27.47)

Table 2: BLEU scores of NMT models and LLMs on the DGT-TM and JRC-A dataset. WMT19 200M indicates WMT19 champion models (Ng et al., 2019), containing 200 million parameters. WMT21 4B indicates WMT21 champion models (Tran et al., 2021) trained by multi language-pairs data containing 4 billion parameters. One-shot and few-shot represent the results of TMPLM with $k = 1$ and $k = 5$, respectively. The BLEU improvements are reported in subscripts. See Table 6 for the COMET-22 version.

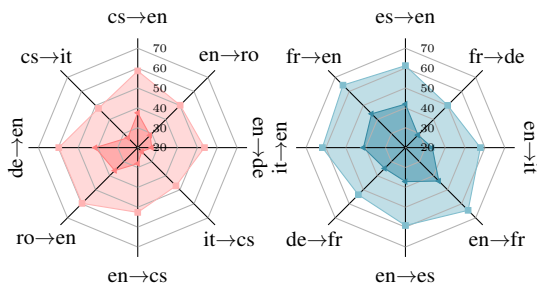


Figure 2: Comparison of LLM w/o and w/ TMs (one-shot) on 8 language-pairs from JRC-A. Points in deep and light color stand for the BLEU scores of LLM w/o and w/ TM, respectively.

For LLMs, we chose the `davinci-003` model developed by OpenAI because it is currently one of the state-of-the-art generative LLMs. The model was configured with default values of all parameters, except that the sampling temperature was set to 0. In the experiments, we used the code-style template and set k to 5 by default. The quality of translations was mainly evaluated using `multi-bleu.perl` from Moses[†]. In addition, following the recommend of using neural network-based metrics in machine translation evaluation (Freitag et al., 2022), we also used COMET-22[‡] (`wmt22-COMET-da`) (Rei et al., 2022) to make a complementary evaluation. See more details about data processing in Appendixes A.3 and A.4.

3.2 Baselines

We re-implemented Reheman et al. (2023)’s method which augments NMT systems via TM-based one-shot learning. For NMT systems, we chose two champion models in WMT: Facebook’s WMT19 en ↔ de models (Ng et al., 2019) and WMT21 multilingual models (Tran et al., 2021). These WMT models were all trained on large-scale

[†]<http://www.statmt.org/moses/>

[‡]<https://github.com/Unbabel/COMET>




Figure 3: Experiments on two impacts including different LLMs and different template styles.

bilingual data and are improved by using a series of techniques, such as back-translation and fine-tuning. As a second baseline, we chose the k NN-MT model (Khandelwal et al., 2021) because it is a very strong model for TM and NMT combination.

3.3 Translation Quality

Main Results. Table 2 shows BLEU scores on the DGT-TM and JRC-A datasets. We see, first of all, that TMPLM achieves the best result among all the systems. When TMs are not involved, the performance of LLMs is 10 BLEU points lower than that of the NMT baselines. But, when armed with TMs, LLMs obtain very large BLEU improvements. The few-shot learning+LLM system even outperforms the strong NMT+TM baseline on all of the test sets. Also, by comparing the results of WMT19 200M models and WMT21 4B models, we see that larger models help more for making use of TM (see Section 3.4 for more discussions). Besides, one-shot learning can give satisfactory results for TMPLM indicating that the most similar TM provides the most helpful translation hints. In Appendix B.4 we will see that few-shot learning yields BLEU gains in a long-tail manner.

Multi-language Experiments. We test TMPLM on more languages and run our system on data of 7 extra language pairs (i.e., 14 directions) from

Domain	k NN-MT	WMT19 200M		WMT21 4B		davinci-003 175B		
		NMT	NMT+TM	NMT	NMT+TM	LLM (zero-shot)	LLM+TM (one-shot)	LLM+TM (few-shot)
IT	45.82	38.09	40.63 _(+2.54)	38.41	46.61 _(+8.20)	20.53	47.46 _(+26.93)	51.03 _(+30.50)
Medical	54.35	41.14	45.78 _(+4.64)	47.94	55.36 _(+7.42)	37.37	58.54 _(+21.17)	60.40 _(+23.03)
Koran	19.45	17.11	17.53 _(+0.42)	23.33	19.27 _(-4.06)	17.59	18.80 _(+1.21)	20.55 _(+2.96)
Law	61.78	45.92	48.97 _(+3.05)	51.60	59.97 _(+8.37)	41.04	61.85 _(+20.81)	64.92 _(+23.88)

Table 3: Comparison of the NMT models and the k NN-MT model on the multi-domain dataset by BLEU. The COMET-22 version can be found in Table 7.




Figure 4: BLEU scores of different prompting strategies on the DGT-TM dataset. In-domain and out-domain represent demonstrations randomly selected from the TM database of the DGT-TM dataset and *newstest2017*, respectively. TM represents top- k similar translation memories (i.e., demonstrations) retrieved from the TM database of the DGT-TM dataset.

JRC-Acquis. From Figure 2, we see consistent improvements over all the language pairs. Even for non-English tasks, TMPLM can still achieve significant BLEU improvements. See Table 8 in Appendix B.3 for complete experimental results.

Multi-domain Experiments. Table 3 shows BLEU results on the multi-domain dataset. Again, the TMPLM system is robust to the domain shift. It performs best on three of the four domains.

3.4 Language Understanding Matters Most

We then investigate an interesting issue: *what kind of ability do large models have to make better use of TM-based prompts?* There are possibly three reasons, including the abilities of *translating*, *logically reasoning* and *language understanding*. However, as seen from Table 2, the baseline LLMs are not strong translation systems and their BLEU scores are generally 10 points lower than the NMT systems. The translation ability of LLMs does not turn out to be important in TM-based prompting. Note that *davinci-003* is a successor of GPT3 and is trained on additional large-scale code data. It

has been pointed out that training LLMs on code data can lead to a strong ability of logical reasoning (Liang et al., 2022). As seen in Figure 3 (a), however, no big difference between *davinci-003* and GPT3 in BLEU performance. On the other hand, *davinci-003* has a significant ability to deal with instructions because it is tuned by using feedback to human instructions. Such a property makes *davinci-003* a better text processor, and thus a stronger translation system that works with various prompts. Therefore, it is the ability of language understanding that boosts LLMs’ translation performance when prompted with TMs.

3.5 Template Styles

In Figure 3 (b), we compare the performance between the code-style and instruction-style templates on the DGT-TM en-de and de-en tasks. For systems without TMs, the instruction-style template shows similar performance as the code-style template. However, when TMs are used, the code-style template is better in most cases. In Appendix B.2, we test more templates and see a similar phenomenon that simpler templates work better.

3.6 Prompting with randomly selected demonstrations

We also compare the performance of TMPLM with the conventional few-shot, i.e., prompting LLM translators with randomly selected high quality demonstrations (Vilar et al., 2022; Agrawal et al., 2022; Zhang et al., 2023; Moslem et al., 2023; Hendy et al., 2023). We conduct experiments on the DGT-TM dataset, with demonstrations selected from the TM database of the DGT-TM dataset (in-domain) and *newstest2017* (out-domain), respectively. In Figure 4, we see, TMPLM exceeds the conventional few-shot by about 30 BLEU points indicating that LLM can benefit from TMs much more than the conventional few-shot itself. It also demonstrates the valid information hinted by TMs as explained in Section 2.




Figure 5: BLEU scores as functions of thresholds of using similar sentences in TMs on the DGT-TM and IT domain data. The left y -axis represents the BLEU scores of prompting LLMs with the translation results from NMT systems, and the x -axis represents the similarity (i.e., the FMS in Appendix A.1) thresholds by which we have a trade-off between using TMs and NMT results as prompts (1 means that we only use TMs as prompts, and 0 means that we only use NMT outputs as prompts). Deep and light red curves represent the performance of the LLMs when working with the WMT19 200M and WMT21 4B systems. Blue curves represent the proportion of the use of TMs (see the right y -axis).

3.7 Combining TMs and NMT results

To examine the impact of high-quality translations on prompting LLMs, we replace the retrieved TM with the translation result of the WMT19 and WMT21 NMT systems when the TM’s similarity is not high enough. We conducted experiments on the DGT-TM $de \rightarrow en$ data and the IT data in the multi-domain dataset because the sentence similarity distributes differently on them (see Appendix A.1). In Figure 5, we can see that the performance declines as more NMT translation results replace the TM results in prompting. This demonstrates that the quality of translations plays an important role in prompting LLMs. We also see that the performance on DGT-TM declines faster than that on IT domain. We attribute this to the better translation quality of the NMT models on the DGT-TM dataset.

There is an interesting finding that the method of prompting LLMs with the NMT results cannot surpass the NMT system itself, while the BLEU scores of prompting LLMs with TMs are always better than those of the TMs. It indicates that LLMs indeed process the prompting texts rather than simply outputting the prompting texts.

4 Conclusion

We have proposed TMPLM, an in-context learning method to prompt TMs for LLMs. By incorporat-

ing TMs into tailored templates, LLMs with TM-PLM outperforms the state-of-the-art NMT models with TM prompting. We have also demonstrated that the ability of language understanding plays an important role in prompting LLMs with TMs.

Limitations

The similarity of TMs is an important factor influencing the translations of TMPLM. However, high-similarity TMs are not always available in practical applications. It is worth studying methods to make use of relatively low-similarity translations in LLM-based translation systems.

Acknowledgements

This work was supported in part by the National Science Foundation of China (No. 62276056), the National Key R&D Program of China, the China HTRD Center Project (No. 2020AAA0107904), the Natural Science Foundation of Liaoning Province of China (2022-KF-16-01), the Yunnan Provincial Major Science and Technology Special Plan Projects (No. 202103AA080015), the Fundamental Research Funds for the Central Universities (Nos. N2216016, N2216001, and N2216002), and the Program of Introducing Talents of Discipline to Universities, Plan 111 (No. B16009).

References

- Sweta Agrawal, Chunting Zhou, Mike Lewis, Luke Zettlemoyer, and Marjan Ghazvininejad. 2022. [In-context examples selection for machine translation](#). *CoRR*, abs/2212.02437.
- Roei Aharoni and Yoav Goldberg. 2020. [Unsupervised domain clusters in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7747–7763. Association for Computational Linguistics.
- Andrzej Bialecki, Robert Muir, and Grant Ingersoll. 2012. Apache lucene 4. In *Proceedings of the SIGIR 2012 Workshop on Open Source Information Retrieval, OSIR@SIGIR 2012, Portland, Oregon, USA, 16th August 2012*, pages 17–24. University of Otago, Dunedin, New Zealand.
- Bram Bulté and Arda Tezcan. 2019. [Neural fuzzy repair: Integrating fuzzy matches into neural machine translation](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 1800–1809. Association for Computational Linguistics.

- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George F. 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 2022, Abu Dhabi, United Arab Emirates (Hybrid), December 7-8, 2022*, pages 46–68. Association for Computational Linguistics.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor O. K. Li. 2018. [Search engine guided neural machine translation](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 5133–5140. AAAI Press.
- Qiuxiang He, Guoping Huang, Qu Cui, Li Li, and Lemao Liu. 2021. [Fast and accurate neural machine translation with translation memory](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 3170–3180. Association for Computational Linguistics.
- 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.
- Nabil Hossain, Marjan Ghazvininejad, and Luke Zettlemoyer. 2020. [Simple and effective retrieve-edit-rerank text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 2532–2538. Association for Computational Linguistics.
- Urvashi Khandelwal, Angela Fan, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2021. [Nearest neighbor machine translation](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2022. [Holistic evaluation of language models](#). *arXiv preprint arXiv:2211.09110*.
- Yuxian Meng, Xiaoya Li, Xiayu Zheng, Fei Wu, Xiaofei Sun, Tianwei Zhang, and Jiwei Li. 2022. [Fast nearest neighbor machine translation](#). In *Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 555–565. Association for Computational Linguistics.
- Yasmin Moslem, Rejwanul Haque, and Andy Way. 2023. [Adaptive machine translation with large language models](#). *CoRR*, abs/2301.13294.
- Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, and Sergey Edunov. 2019. [Facebook fair’s WMT19 news translation task submission](#). In *Proceedings of the Fourth Conference on Machine Translation, WMT 2019, Florence, Italy, August 1-2, 2019 - Volume 2: Shared Task Papers, Day 1*, pages 314–319. Association for Computational Linguistics.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. [Training language models to follow instructions with human feedback](#). *arXiv preprint arXiv:2203.02155*.
- Abudurexiti Reheman, Tao Zhou, Yingfeng Luo, Di Yang, Tong Xiao, and Jingbo Zhu. 2023. [Prompting neural machine translation with translation memories](#). *arXiv preprint arXiv:2301.05380*.
- Ricardo Rei, José G. C. de Souza, Duarte M. Alves, Chrysoula Zerva, Ana C. Farinha, Taisiya Glushkova, Alon Lavie, Luísa Coheur, and André F. T. Martins. 2022. [COMET-22: unbabel-ist 2022 submission for the metrics shared task](#). In *Proceedings of the Seventh Conference on Machine Translation, WMT 2022, Abu Dhabi, United Arab Emirates (Hybrid), December 7-8, 2022*, pages 578–585. Association for Computational Linguistics.
- Ralf Steinberger, Andreas Eisele, Szymon Kłoczek, Spyridon Pilos, and Patrick Schläuter. 2012. [DGT-TM: A freely available translation memory in 22 languages](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012, Istanbul, Turkey, May 23-25, 2012*, pages 454–459. European Language Resources Association (ELRA).
- Ralf Steinberger, Bruno Pouliquen, Anna Widiger, Camelia Ignat, Tomaz Erjavec, Dan Tufis, and Dániel Varga. 2006. [The jrc-acquis: A multilingual aligned parallel corpus with 20+ languages](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation, LREC 2006, Genoa, Italy, May 22-28, 2006*, pages 2142–2147. European Language Resources Association (ELRA).
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. [Lamda: Language models for dialog applications](#). *arXiv preprint arXiv:2201.08239*.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012, Istanbul, Turkey, May 23-25, 2012*, pages 2214–2218. European Language Resources Association (ELRA).

Chau Tran, Shruti Bhosale, James Cross, Philipp Koehn, Sergey Edunov, and Angela Fan. 2021. [Facebook ai’s WMT21 news translation task submission](#). In *Proceedings of the Sixth Conference on Machine Translation, WMT@EMNLP 2021, Online Event, November 10-11, 2021*, pages 205–215. Association for Computational Linguistics.

David Vilar, Markus Freitag, Colin Cherry, Jiaming Luo, Viresh Ratnakar, and George Foster. 2022. Prompting palm for translation: Assessing strategies and performance. *arXiv preprint arXiv:2211.09102*.

Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.

Tong Xiao, Jingbo Zhu, Hao Zhang, and Qiang Li. 2012. [NiuTrans: An open source toolkit for phrase-based and syntax-based machine translation](#). In *The 50th Annual Meeting of the Association for Computational Linguistics, Proceedings of the System Demonstrations, July 10, 2012, Jeju Island, Korea*, pages 19–24. The Association for Computer Linguistics.

Jitao Xu, Josep Maria Crego, and Jean Senellart. 2020. [Boosting neural machine translation with similar translations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1580–1590. Association for Computational Linguistics.

Biao Zhang, Barry Haddow, and Alexandra Birch. 2023. Prompting large language model for machine translation: A case study. *arXiv preprint arXiv:2301.07069*.

Jingyi Zhang, Masao Utiyama, Eiichiro Sumita, Graham Neubig, and Satoshi Nakamura. 2018. [Guiding neural machine translation with retrieved translation pieces](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 1325–1335. Association for Computational Linguistics.

A Detailed Experimental Setup

A.1 Retrieval of Similar Sentences

Following [Reheman et al., 2023](#), we adopt a word-level fuzzy matching strategy, with the numbers and punctuation marks removed. Specifically, we first use the search engine Apache Lucene ([Bialecki et al., 2012](#)) to acquire the Top500 similar TMs from TM database, then rerank the most similar TM by using the length normalized Levenshtein Distance, given by

$$\text{FMS}(X, S) = 1 - \frac{\text{LD}(X, S)}{\max(|X|, |S|)} \quad (3)$$

where $\text{FMS}(\cdot, \cdot)$ denotes the Fuzzy Match Score, $\text{LD}(\cdot, \cdot)$ denotes the word level Levenshtein Distance, and $|\cdot|$ denotes the length of a sentence.

A.2 Details of Datasets

Datasets and their language directions used in our experiments are listed here.

- The DGT-TM dataset([Tiedemann, 2012](#)), which is bidirectional in English-German;
- The JRC-Acquis (JRC-A) dataset([Steinberger et al., 2006](#)), which includes 8 language pairs and 16 directions: English-German, English-French, German-French, English-Italian, English-Romanian, English-Spanish, English-Czech, and Czech-Italian;
- The multi-domain dataset ([Aharoni and Goldberg, 2020](#)), which includes 4 domains in the German to English direction: Medical, Law, IT, and Koran.

The statistics of these TM and the corresponding similarity ratios of retrieved sentences in the FMS metric are shown in Table 4.

A.3 Data Pre-processing

For the DGT-TM, JRC-A and multi-domain datasets, we clean the data using the scripts provided by [Reheman et al. \(2023\)](#)’s work. To construct the test set and TM database for the DGT-TM and JRC-A datasets, we process each language direction separately. Specifically, we randomly extract 3,000 sentence pairs from each dataset as the test set, and use the remaining sentence pairs as the TM database. For the multi-domain dataset, we use its original test set as our test set and its original training set as the TM database. We use the FMS algorithm on the split data to obtain the TM corresponding to the test set. In particular, for the few-shot experiments, we retrieved the k most similar sentence pairs from the TM database for each test sentence.

Finally, we replace the escaped characters in the dataset and use Moses[§] decoder detokenizer to recover the tokenized data before feeding it to the `davinci-003` system.

A.4 Data Post-processing

`davinci-003` always generates redundant symbols at the beginning and end of sentences, including: ‘”’, ‘\n’, ‘[’, ‘]’, and other escaped

[§]<http://www.statmt.org/ Moses/>

Dataset	Lang	Domain	TM scale	FMS				
				[0, 0.2)	[0.2, 0.4)	[0.4, 0.6)	[0.6, 0.8)	[0.8, 1.0)
DGT-TM	En-De	-	3.1M	2%	23%	16%	17%	42%
	De-En	-	3.1M	4%	26%	17%	17%	36%
JRC-A	En-De	-	423K	6%	33%	18%	13%	30%
	De-En	-	423K	6%	33%	18%	15%	28%
	En-Fr	-	424K	3%	34%	19%	14%	30%
	Fr-En	-	424K	3%	33%	19%	15%	30%
	De-Fr	-	846K	9%	34%	16%	12%	29%
	Fr-De	-	846K	8%	34%	16%	12%	30%
	En-It	-	433K	7%	32%	18%	14%	29%
	It-En	-	433K	7%	32%	17%	16%	28%
	En-Ro	-	273K	7%	39%	21%	14%	19%
	Ro-En	-	273K	6%	37%	22%	15%	20%
	En-Es	-	432K	2%	34%	20%	16%	28%
	Es-En	-	432K	2%	34%	20%	16%	28%
	En-Cs	-	681K	12%	33%	17%	12%	26%
	Cs-En	-	681K	13%	32%	15%	13%	27%
Cs-It	-	714K	11%	31%	17%	14%	27%	
It-En	-	714K	12%	32%	16%	13%	27%	
multi-domain	De-En	IT	223K	14%	18%	28%	26%	14%
	De-En	Koran	18K	2%	26%	33%	28%	11%
	De-En	Law	467K	8%	31%	18%	14%	28%
	De-En	Medical	248K	7%	23%	20%	17%	33%
WMT14	En-De	-	4.5M	18%	68%	12%	1%	1%
WMT19	De-En	-	30M	13%	65%	19%	2%	1%

Table 4: TMs and proportions of the retrieved sentences in different ranges of FMS.

characters. The occurrences of these characters is regular and can be removed uniformly by scripts. Consequently, before scoring, we use NiuTrans (Xiao et al., 2012) word segmentation tool for Chinese and Moses decoder’s `tokenizer.perl` for all other languages. Finally we use `multi-bleu.perl` for scoring.

A.5 More Prompt Templates

We try a large number of prompt templates, as shown in Table 5. Without special specification, the instruction-style template with TM is the #1, and without TM is the #2, and the code-style template with TM is the #17, and without TM is the #18. In particular, in the multi-language experiment, we use the instruction-style template. The template for all of the few-shot experiments is obtained by increasing the number of TMs in #17.

Punctuation has a significant impact on the generation results. For example, using template #13, if the source sentence ends with ‘:’, it will lead the model to continue generating words but not stop in an appropriate number of decoding steps. Meanwhile, although many templates have a similar form, their performance still differs. We believe that adding a strong boundary signal to the templates helps the model to know where to end.

B More Experimental Results

B.1 Evaluation by COMET-22

Except for the BLEU scores, we also provide the COMET-22 scores as seen in Table 6 and Table 7. We can see that despite LLM’s poor performance on zero-shot, prompting LLM with a few TMs can achieve significant improvement. On the other hand, the few-shot learning+LLM system can still outperform the strong NMT+TM baseline in most cases.

B.2 Performance of Different Prompt Templates

In order to explore the effect of using different prompt templates on the performance of `davinci-003`, we use 20 prompt templates in the `de → en` direction of the DGT-TM dataset for experiments. Seen from table 5, the code-style template is better than the instruction-style template in most cases.

B.3 Experiments on More languages

We perform multi-lingual experiments on the JRC-A dataset, and in these experiments, we use the instruction-style template as shown in Figure 1. Table 8 shows the complete experiment results for the

No.	Prompt Template	With TM	Sample	BLEU
1	If the translation of " x_{tm} " from <i>src-lang</i> to <i>tgt-lang</i> is " y_{tm} ", then what is the translation of " x " from <i>src-lang</i> to <i>tgt-lang</i> ? Only translation results are required.	Yes	If the translation of "I have an apple." from English to German is "Ich habe einen Apfel." then what is the translation of "I have an orange." from English to German? Only translation results are required.	63.97
2	What is the translation of " x " from <i>src-lang</i> to <i>tgt-lang</i> ? Only translation results are required.	No	What is the translation of "I have an apple." from English to German? Only translation results are required.	38.38
3	If " x_{tm} " translated into <i>tgt-lang</i> is " y_{tm} ", then what is the translation of " x " should be if translated into <i>tgt-lang</i> ? Only translation results are required.	Yes	If "I have an apple." translated into German is "Ich habe einen Apfel." then what is the translation of "I have an orange." should be if translated into German? Only translation results are required.	61.9
4	What is the translation of " x " should be if translated into <i>tgt-lang</i> ? Only translation results are required.	No	What is the translation of "I have an apple." should be if translated into German? Only translation results are required.	37.93
5	If $[x_{tm}]$ translated into <i>tgt-lang</i> is $[y_{tm}]$, then what is the translation of $[x]$ should be if translated into <i>tgt-lang</i> ? Only translation results are required.	Yes	If [I have an apple.] translated into German is [Ich habe einen Apfel.] then what is the translation of [I have an orange.] should be if translated into German? Only translation results are required.	61.78
6	Translate <i>src-lang</i> to <i>tgt-lang</i> . [<i>src-lang</i>]: $[x_{tm}]$ [<i>tgt-lang</i>]: $[y_{tm}]$ [<i>src-lang</i>]: $[x]$ [<i>tgt-lang</i>]:	Yes	Translate English to German. [English]: [I have an apple.] [German]: [Ich habe einen Apfel.] [English]: [I have an orange.] [German]:	65.25
7	Translate <i>src-lang</i> to <i>tgt-lang</i> . [<i>src-lang</i>]= $[x_{tm}]$ [<i>tgt-lang</i>]= $[y_{tm}]$ [<i>src-lang</i>]= $[x]$ [<i>tgt-lang</i>]=	Yes	Translate English to German. [English]=[I have an apple.] [German]=[Ich habe einen Apfel.] [English]=[I have an orange.] [German]=	66.02
8	Translate <i>src-lang</i> to <i>tgt-lang</i> . [<i>src-lang</i>]= $[x_{tm}]$ [<i>tgt-lang</i>]= $[y_{tm}]$ [<i>src-lang</i>]= $[x]$ [<i>tgt-lang</i>]=	Yes	Translate English to German. [English]=[I have an apple.] [German]=[Ich habe einen Apfel.] [English]=[I have an orange.] [German]=	66.08
9	Translate <i>src-lang</i> to <i>tgt-lang</i> . [<i>src-lang</i>]= $[x_{tm}]$ [<i>tgt-lang</i>]= $[y_{tm}]$ [<i>src-lang</i>]= $[x]$ [<i>tgt-lang</i>]=	Yes	Translate English to German. [English]=[I have an apple.] [German]=[Ich habe einen Apfel.] [English]=[I have an orange.] [German]=	66.20
10	if <i>src-lang</i> = $[x_{tm}]$ then <i>tgt-lang</i> = $[y_{tm}]$; if <i>src-lang</i> = $[x]$ then <i>tgt-lang</i> =	Yes	if English = [I have an apple.] then German = [Ich habe einen Apfel.]; if English = [I have an orange.] then German =	66.75
11	<i>src-lang</i> =" x_{tm} " <i>tgt-lang</i> =" y_{tm} " <i>src-lang</i> =" x " <i>tgt-lang</i> =	Yes	English="I have an apple." German="Ich habe einen Apfel." English="I have an orange." German=	66.28
12	<i>src-lang</i> = $[x_{tm}]$ <i>tgt-lang</i> = $[y_{tm}]$ <i>src-lang</i> = $[x]$ <i>tgt-lang</i> =	Yes	English=[I have an apple.] German=[Ich habe einen Apfel.] English=[I have an orange.] German=	65.37
13	[<i>src-lang</i>] x_{tm} [<i>tgt-lang</i>] y_{tm} [<i>src-lang</i>] x [<i>tgt-lang</i>]	Yes	[English] I have an apple. [German] Ich habe einen Apfel. [English] I have an orange. [German]	58.47
14	[<i>src-lang</i>]: $[x_{tm}]$ [<i>tgt-lang</i>]: $[y_{tm}]$ [<i>src-lang</i>]: $[x]$ [<i>tgt-lang</i>]:	Yes	[English]: [I have an apple.] [German]: [Ich habe einen Apfel.] [English]: [I have an orange.] [German]:	65.54
15	[<i>src-lang</i>]: $[x]$ [<i>tgt-lang</i>]:	No	[English]: [I have an orange.] [German]:	39.83
16	[<i>src-lang</i>] = $[x_{tm}]$ [<i>tgt-lang</i>] = $[y_{tm}]$ [<i>src-lang</i>] = $[x]$ [<i>tgt-lang</i>] =	Yes	[English] = [I have an apple.] [German] = [Ich habe einen Apfel.] [English] = [I have an orange.] [German] =	66.45
17	[<i>src-lang</i>]= $[x_{tm}]$ [<i>tgt-lang</i>]= $[y_{tm}]$ [<i>src-lang</i>]= $[x]$ [<i>tgt-lang</i>]=	Yes	[English]=[I have an apple.] [German]=[Ich habe einen Apfel.] [English]=[I have an orange.] [German]=	66.90
18	[<i>src-lang</i>]= $[x]$ [<i>tgt-lang</i>]=	No	[English]=[I have an orange.] [German]=	38.89
19	{ <i>src-lang</i> }= $\{x_{tm}\}$ { <i>tgt-lang</i> }= $\{y_{tm}\}$ { <i>src-lang</i> }= $\{x\}$ { <i>tgt-lang</i> }=	Yes	{English}={I have an apple.} {German}={Ich habe einen Apfel.} {English}={I have an orange.} {German}={}	65.48
20	{[<i>src-lang</i>]}= $\{x_{tm}\}$ {[<i>tgt-lang</i>]}= $\{y_{tm}\}$ {[<i>src-lang</i>]}= $\{x\}$ {[<i>tgt-lang</i>]}=	Yes	{[English]}={I have an apple.} {[German]}={Ich habe einen Apfel.} {[English]}={I have an orange.} {[German]}={}	63.32

Table 5: Comparison of prompt templates in one-shot TM (i.e., $k = 1$). Abbreviations are same as Figure 1.

multi-language experiment. Great BLEU improvements are obtained on these datasets.

B.4 Impact of k

To explore the effect of k on the performance of davinci-003 in the few-shot experiments, we

conduct experiments with k from 1 to 9 in both directions of the DGT-TM dataset. Figure 6 shows a long-tail performance gain as k increases.

Data		WMT19 200M		WMT21 4B		davinci-003 175B		
		NMT	NMT+TM	NMT	NMT+TM	LLM (zero-shot)	LLM+TM (one-shot)	LLM+TM (few-shot)
DGT-TM	de → en	85.99	87.28 _(+1.29)	87.10	89.28 _(+2.18)	83.86	88.74 _(+4.88)	89.47 _(+5.61)
	en → de	85.52	86.91 _(+1.39)	86.89	89.01 _(+2.12)	82.24	88.52 _(+6.28)	89.44 _(+7.20)
JRC-A	de → en	85.85	85.80 _(-0.05)	86.68	87.70 _(+1.02)	84.15	87.79 _(+3.64)	88.46 _(+4.31)
	en → de	86.53	86.25 _(-0.28)	87.39	88.88 _(+1.49)	84.03	88.20 _(+4.17)	88.84 _(+4.81)

Table 6: COMET-22 scores of NMT models and LLMs on the DGT-TM and JRC-A dataset.

Domain	WMT19 200M		WMT21 4B		davinci-003 175B		
	NMT	NMT+TM	NMT	NMT+TM	LLM (zero-shot)	LLM+TM (one-shot)	LLM+TM (few-shot)
IT	83.04	83.87 _(+0.83)	83.54	85.09 _(+1.55)	72.44	86.05 _(+13.61)	87.27 _(+14.83)
Medical	83.30	83.61 _(+0.31)	84.92	84.97 _(+0.05)	80.76	84.97 _(+4.21)	86.63 _(+5.87)
Koran	72.42	72.00 _(-0.42)	75.09	72.23 _(-2.86)	73.35	73.65 _(+0.30)	74.34 _(+0.99)
Law	85.80	85.53 _(-0.27)	86.75	87.23 _(+0.48)	83.30	87.49 _(+4.19)	88.47 _(+5.17)

Table 7: COMET-22 scores of NMT models and LLMs on the multi-domain dataset.

Lang Direction	Cs-En		Cs-It		De-En		De-Fr	
	→	←	→	←	→	←	→	←
w/o TM	58.58	52.58	47.93	47.22	59.74	53.56	53.52	50.08
w/ TM	37.50	28.02	27.02	22.85	36.34	28.74	34.78	28.93
Δ	+21.08	+24.56	+20.91	+24.37	+23.40	+24.82	+18.74	+21.15
Lang Direction	Es-En		Fr-En		It-En		Ro-En	
	→	←	→	←	→	←	→	←
w/o TM	61.25	59.18	64.45	64.60	61.80	57.71	59.66	50.18
w/ TM	41.89	37.10	44.06	43.78	41.13	33.53	41.37	27.06
Δ	+19.36	+22.08	+20.39	+20.78	+20.67	+24.18	+18.29	+23.12

Table 8: Experiment results on 8 language-pairs from JRC-A.

Data		NMT System	FMS										
			0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
DGT-TM	en → de	WMT19	57.39	57.53	57.70	58.02	57.71	57.08	56.07	54.73	53.05	48.67	37.06
	de → en		66.90	66.90	66.88	66.34	65.63	64.30	62.70	61.04	58.45	54.15	44.21
	de → en		WMT21	66.90	66.90	67.06	67.02	66.74	65.84	64.58	63.08	60.59	57.33
IT	de → en	WMT19	47.46	44.48	43.85	43.30	40.78	36.92	35.78	32.55	29.27	28.17	26.65
	de → en	WMT21	47.46	44.01	43.19	42.31	37.53	34.35	33.31	30.17	27.23	26.25	24.90
Law	de → en	WMT19	61.85	61.84	61.67	60.89	59.61	58.04	56.75	55.04	53.00	50.10	44.12
	de → en	WMT21	61.85	61.83	62.00	61.99	61.39	60.34	59.32	57.98	56.31	53.93	47.56
Medical	de → en	WMT19	58.54	58.45	58.32	58.05	57.25	55.34	54.06	52.62	50.03	47.01	38.85
	de → en	WMT21	58.54	58.45	58.32	58.11	57.30	55.44	54.14	52.74	20.15	47.14	38.97

Table 9: Performance of replacing the low-matching part of TMs at different thresholds of FMS with the translation results from NMT. For example, FMS 0.2 in first row means that TMs with FMS less than 0.2 are replaced by NMT translation results.

B.5 Impact of Orders of TM results

To observe the effect of constructing the prompt template with different TMs similarity orders on the performance in the few-shot experiments, we

constructed two types of prompt templates in the DGT-TM dataset with a few-shot sample size of 5. One is arranged in descending order of TMs similarity, and the TM adjacent to the sentence to




Figure 6: BLEU scores of different k on the DGT dataset

Lang Direction	Descending	Ascending
de \rightarrow en	69.99	70.01
en \rightarrow de	62.02	62.30

Table 10: The performance comparison of different templates which is constructed based on the similarity of TM when the number of few-shot samples is 5.

Model	TM	WMT14 En2De	WMT19 De2En
Transformer-base	w/o TM	27.59	39.67
	w/ TM	21.86	40.22
text-davinci-003	w/o TM	29.58	40.85
	w/ TM	28.11	36.63

Table 11: Comparison of performance on WMT dataset.

be translated is the lowest similarity. The other one is arranged in ascending order of TMs similarity, and the TM adjacent to the sentence to be translated is the highest similarity. The results are shown in Table 10.

B.6 Performance on the WMT Datasets

We conduct experiments on WMT14 en \rightarrow de and WMT19 de \rightarrow en directions. We use the same method as that used on the multi-domain dataset to process these two benchmarks. It is worth noting that the data obtained on these two benchmarks have a low similarity of TMs, as shown in Table 11. Table 11 shows the performance of the LLM and baseline models on the WMT14 en \rightarrow de and WMT19 de \rightarrow en datasets.

Model	BLEU
text-davinci-003	66.90
davinci(GPT3)	65.48
text-curie-001	42.30
text-babbage-001	37.72
text-ada-001	14.65

Table 12: Comparison of performance with different size models on DGT-TM de \rightarrow en.

B.7 Performance of Different Sized Models

Moreover, we conduct experiments using “small” models such as text-curie-001 and text-babbage-001. But their performance is far away behind davinci-003 whose outputs contain null in lines sometimes. We attribute this to the lack of emergent abilities of big models (Wei et al., 2022). The results are shown in Table 12.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitation In Page 5
- A2. Did you discuss any potential risks of your work?
Limitation In Page 5
- A3. Do the abstract and introduction summarize the paper's main claims?
Section1, in Page 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
No response.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
No response.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
No response.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
No response.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
No response.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
No response.

C Did you run computational experiments?

Section 3

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 3.2

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

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

Section 3.1

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

We draw conclusions on the basis of experiments conducted on a wide range of datasets and language pairs.

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

Appendix A

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

Left blank.

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

No response.

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

No response.

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

No response.

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

No response.

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

No response.