

Narrowing the Gap between Zero- and Few-shot Machine Translation by Matching Styles

Weiting Tan, Haoran Xu, Lingfeng Shen, Shuyue Stella Li
Kenton Murray, Philipp Koehn, Benjamin Van Durme, Yunmo Chen

Johns Hopkins University
{wtan12, hxu64, lshen30, sli136, kenton, phi, vandurme, yunmo}@jhu.edu

Abstract

Large language models trained primarily in a monolingual setting have demonstrated their ability to generalize to machine translation using zero- and few-shot examples with in-context learning. However, even though zero-shot translations are relatively good, there remains a discernible gap comparing their performance with the few-shot setting. In this paper, we investigate the factors contributing to this gap and find that this gap can largely be closed (for about 70%) by matching the writing styles of the target corpus. Additionally, we explore potential approaches to enhance zero-shot baselines without the need for parallel demonstration examples, providing valuable insights into how these methods contribute to improving translation metrics.

1 Introduction

Recent advancements in large language models (LLMs) have revolutionized Natural Language Processing field as such models (OpenAI, 2023; Ouyang et al., 2022; Chowdhery et al., 2022; Touvron et al., 2023, *inter alia*) can easily adapt to a new task through prompt (in-context) learning where the task instruction and demonstrations (examples to guide LLMs on the task) are provided to the model. Such capability opens up new opportunities for machine translation, which is traditionally trained/fine-tuned on large amounts of parallel corpus (Brown et al., 1993; Bahdanau et al., 2016; Vaswani et al., 2017; Team et al., 2022, *inter alia*). Recent work (Vilar et al., 2022; Zhang et al., 2023; Jiao et al., 2023; Hendy et al., 2023) has found that prompt-based methods perform well on language models trained primarily on monolingual data, rivaling state-of-the-art systems trained specifically for machine translation tasks on benchmark datasets.

Comparing zero-and few-shot outputs (§2), we observe a huge gap in terms of BLEU (Papineni

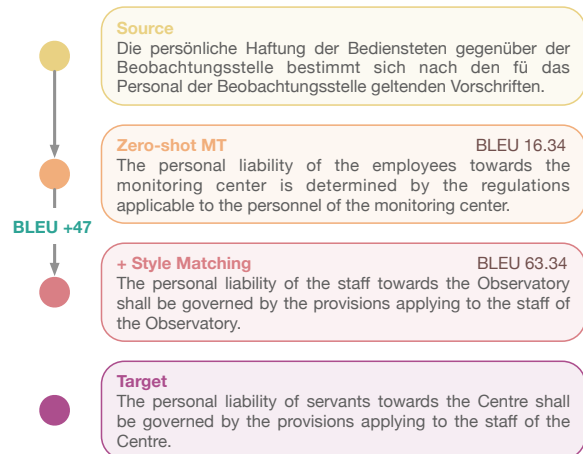


Figure 1: An example of translations under different settings. The zero-shot translation conveys the most semantic meaning of its source but lacks the ability to match with the style of its target sentence, resulting in a low BLEU score. With style-matching prompting, the LLM is able to generate a much better translation without accessing additional parallel examples.

et al., 2002) score, despite the acceptable quality of zero-shot translations. As shown in Figure 1, the zero-shot translation already conveys the meaning but the few-shot translation is better matched with the target (thus obtaining a better BLEU score).

We present qualitative and quantitative analysis on zero- and few-shot translation to understand their quality difference. Our evaluation (§2) reveals that the performance gap stems mostly from writing styles rather than semantics. This finding motivates us to quantify the degree of style match between translations and references (§3) and develop a data-efficient style-learning prompting strategy (§4).¹ Compared to few-shot translation, our prompting strategy relies solely on retrievals from in-domain target corpora and demonstrates effectiveness in closing approximately 70% of the gap between zero-and few-shot translation. We provide an

¹ The style-learning prompting only has access to samples from target corpora. We denote “few-shot” as the experiment setup that requires paired parallel demonstrations.

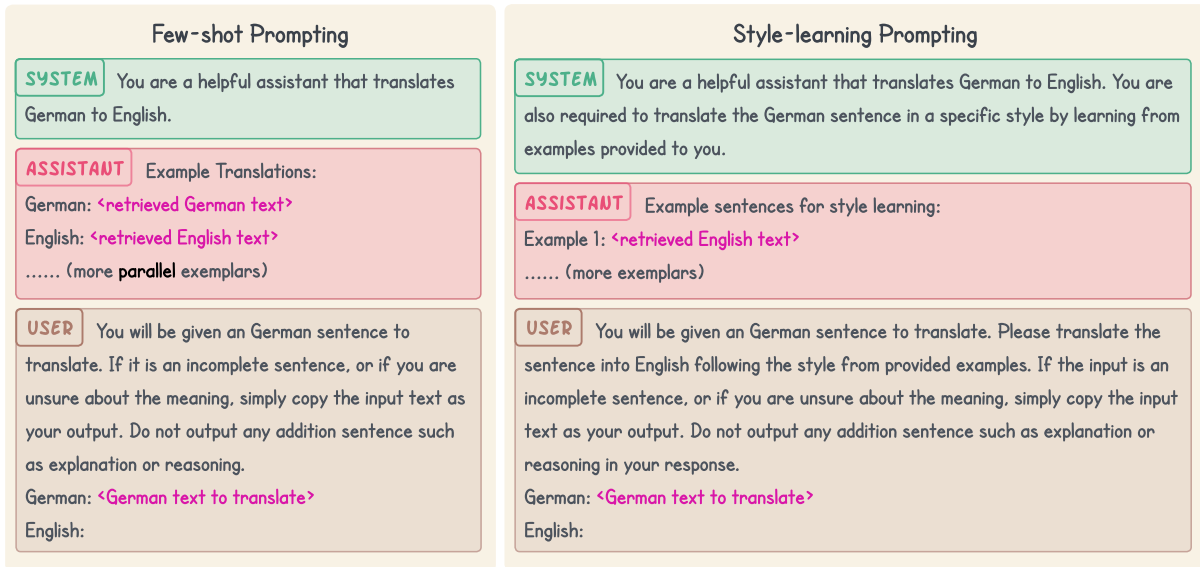


Figure 2: Prompt templates for instructing models to perform machine translation. *Left* is the prompt template used for zero- and few-shot translation; in the few-shot scenario, source and target language pairs are required; in the zero-shot setting, example translations are not provided (*i.e.* no assistant message is sent). *Right* is the prompt template that incorporates style-learning instruction which only requires retrieving monolingual sentences from target corpora. In fact, style-learning prompting can also be considered as a kind of few-shot, but it only has access to target corpora. We denote “few-shot” as the experiment setup that requires parallel paired demonstrations.

example in Figure 1 to illustrate the effectiveness of the proposed style-matching approach.

2 The Gap between Zero- and Few-shot

For our experiments, we use the German-English data splits in Aharoni and Goldberg (2020)² that comprises IT, Subtitle, Law, Medical, and Koran domains. Due to the data quality issue, we keep Law, Medical, and Koran domains as they are relatively cleaner and representative of specialized domains other than generic daily language. For the details of setups and discussion about the dataset, we defer readers to Appendix A.

Based on prior work on in-context learning (Brown et al., 2020; Zhang et al., 2023; Vilar et al., 2022; Jiao et al., 2023; Hendy et al., 2023; Sia and Duh, 2023, *inter alia*), we design prompts to instruct LLMs to perform zero- and few-shot translations. Specifically, for few-shot translation, we employ a retrieval-based prompting method (Hendy et al., 2023; Moslem et al., 2023; Vilar et al., 2022) that retrieves source and target pairs from the training corpus as in-context learning (ICL) demonstrations.

To facilitate this process, we create a prompt template illustrated in Figure 2 *Left*. In this

² The dataset is originally collected by Koehn and Knowles (2017).

Method	BLEU	COMET
Vanilla Transformer	41.4	80.4
GPT3.5 0-shot	31.4	82.1
GPT3.5 5-shot	47.6	84.6

Table 1: We compare average BLEU scores for prompt-based methods (using gpt-3.5-turbo-0301 endpoint to translate from German to English) and vanilla transformer (a 6-layer encoder-decoder structure as introduced by Vaswani et al. (2017)) trained on the in-domain dataset. We found that with few-shot translation, LLM could generate translation that achieves much better performance than 0-shot as well as domain-specific models.

template, the retrieved examples are embedded within the assistant message to aid in the few-shot translation. To retrieve relevant examples, we utilize the BM25 (Robertson et al., 1994) retriever, with source sentences serving as queries. By searching the source language corpus, we identify the top matches and extract aligned parallel sentence pairs as demonstration examples.

We evaluate translation outputs using BLEU (Papineni et al., 2002) and COMET (Rei et al., 2020), and the results are presented in Table 1. In comparison to the few-shot performance (5-shot), zero-shot translation exhibits a noticeable gap of 16.2 points on the BLEU score. However, the difference is less significant when considering COMET, as the




Figure 3: *Left*: trigram overlap analysis between *translations* and *retrieved examples*; the gaps between zero- and few-shot indicate that choices of lexicons in few-shot have been greatly impacted by retrieved demonstrations. *Mid*: averages of Tree Edit Distance (TED) between dependency parse trees from references and zero- or few-shot translations. Few-shot has a lower TED, showing that its grammar is more conformed with references. *Right*: averages of TED between constituency parses from references and zero- or few-shot translations. Again, the few-shot results with lower TED are more conformed in syntactical structure with references.

few-shot approach only outperforms the zero-shot counterpart by 2.5 points. To provide a better understanding of the COMET difference, we present a few examples in Figure 7 to illustrate how the score difference is reflected in translations.³

Examining the example with a COMET difference of 2.7 points, we observe that such a scale of difference, when mapped to lexicons, corresponds to only a few word changes, while the underlying semantics remain largely consistent with the reference. This minor difference in COMET scores indicates that zero-shot translations have already conveyed the semantic meaning of the source sentence, albeit with some variations in lexical choices and sentence structure.

3 How Much Does Style Contribute?

We have observed that the lexical difference between zero-shot and few-shot translations is more noticeable than their semantic difference. Following the traditional linguistic definition of language style (McDonald and Pustejovsky, 1985; Jin et al., 2022, *inter alia*), we can interpret this phenomenon as a distinction in writing style, where semantics remains consistent while aspects such as word choices, syntax, tones, etc., are altered.

To examine the potential impact of writing styles being learned through few-shot demonstrations, we investigate various aspects associated with writing styles to determine whether few-shot translation exhibits more consistent styles with target domains. We follow common techniques used in prior work (Calvo et al., 2014; Jankowska, 2017; Stamatatos et al., 2000) to analyze style matching through three aspects: lexical overlaps, syntactical structures, and

³ These examples serve as reference points to help readers qualitatively grasp the disparity in COMET scores.

word dependencies. These aspects provide valuable insights into the output style of different systems.

3.1 Choice of Lexicons

To understand how few-shot demonstrations affect lexicon usage in translations, we perform an analysis that identifies common n-gram presence between *retrieved samples* and *translations*. The match rate difference between zero- and few-shot indicates how much these examples change the choices of lexicons in translations. We compute average trigram match rates between top-5 retrieved samples and zero- and few-shot translations (shown in Figure 3 *Left*).⁴ Across all domains, the n-gram match rates are much higher for few-shot translations, indicating that few-shot translations are being steered to be lexical-wise more similar to demonstrations. Given that few-shot translations obtain better scores on BLEU in general, we attribute a non-trivial portion of performance gain to lexicons learned from its demonstrations, *therefore better examples matched with target domains can potentially provide better guidance for few-shot translation*. A qualitative example is provided in Figure 6.

3.2 Organization of Sentence Structures

Word Semantic Relations The degree of style matching between translations and the references can be measured by comparing the roles of different words under dependency parse trees (shown in Figure 3 *Mid*).⁵ We quantify the difference between parses using Tree Edit-Distance (TED),

⁴ We explored a range of n-grams from unigram to 10-gram while observing a similar gap, thus presenting only the trigram results for simplicity.

⁵ Dependency parsing and constituency parsing were performed using the spaCy toolkit (<https://spacy.io/>)




Figure 4: Results of few-shot, zero-shot, and zero-shot with style-learning prompting evaluated using BLEU and COMET. The x-axis is the number of demonstrations used in the prompt; in few-shot, they are parallel sentence pairs, whereas in zero-shot with style learning, they are just target samples. The y-axis is the metric score. Both few-shot and style-learning prompting improve from the zero-shot baseline substantially even with only one example.

where smaller mean closer (or more similar) between trees.⁶ Few-shot translations exhibit a smaller overall TED, indicating closer alignments in word semantic relations with references.

Syntactical Structure Similarity We also apply TED on constituent trees to understand similarities of syntactical structures between translations and references.⁷ Our results show that few-shot translations also demonstrates higher similarities in syntactical structures with target domains.

4 Narrowing the Gap via Matching Styles

From §3, we demonstrated that few-shot conforms to the target style better than zero-shot translation. Motivated by this observation, we investigate approaches to mitigate the style difference so that zero-shot translation can be improved without access to additional parallel examples.

We propose a novel style-learning prompting approach that leverages target corpora, as illustrated in Figure 2 Right. Our hypothesis is that style-relevant information primarily resides in the in-domain target corpus (in our case, English). Therefore, by retrieving samples from the corresponding monolingual corpus, the style-learning instruction would enhance style consistency in translation. This idea is in the same spirit as prior work in test time domain adaptation (Singh and Ortega, 2022; Zheng et al., 2022; Zhang et al., 2022; Hu et al., 2019, *inter alia*), where target domain information is used to guide model inference. In our case, we allow language models

(LLMs) to access examples from the target corpora, enabling them to adjust their language styles.

Our approach has two steps: first, we employ zero-shot translation; then we take the translation output as a query to retrieve examples from the target language’s corpus and *add them to the original translation prompt with a style-learning instruction*. By revising prompts with target examples, the approach allows us to align styles in translations without source-side demonstrations.

The performance comparison⁸ among few-shot, zero-shot, and zero-shot with style matching is shown in Figure 4, where the number of demonstrations ranges from 1 to 50. From the results, few-shot prompting exhibits the best performance by leveraging both source and target information. Our style-learning prompts largely improve the zero-shot performance, reducing the gap by approximately 70%. To further validate these findings, we conducted experiments using the Llama2-7b model (Touvron et al., 2023), fixing the number of demonstrations at 5 for both German-English and English-German translation directions. As depicted in Table 2⁹, the trend remains consistent, demonstrating that style-learning prompts effectively bridge the gap between zero-shot and few-shot translations. Interestingly, Llama2-7b shows less improvement with style-learning prompts and has a notably lower zero-shot baseline compared to GPT3.5, suggesting that the impact of style-learning varies among models, dependent on their initial translation capabilities.

Overall, few-shot prompting is still more effective and we attribute such advantages of few-

⁶ TED is computed using the Zhang-shasha algorithm (Zhang and Shasha, 1989), a classical algorithm for computing similarity between two parses.

⁷ We removed nodes at the bottom level so that TED would not reflect lexical differences.

⁸ The breakdowns for domains are provided in Figure 8.

⁹ Some numbers for GPT3.5 are slightly different from Figure 4 as they are rerun experiments over a 500 test subset

Direction	Method	Law		Med		Koran		Average	
		GPT3.5	Llama2	GPT3.5	Llama2	GPT3.5	Llama2	GPT3.5	Llama2
DE → EN	Zero-shot	37.2	32.5	44.5	39.2	17.9	14.4	33.2	28.7
	Style	48.9	46.1	57.8	44.6	19.8	15.5	42.2(+9.0)	35.4(+6.7)
	Few-shot	59.3	58.9	64.1	58.9	20.2	16.6	47.9(+14.7)	44.8(+16.1)
EN → DE	Zero-shot	32.1	22.6	38.8	28.7	12.8	8.3	27.9	19.9
	Style	41.1	28.0	47.8	32.9	23.5	14.9	37.4(+9.5)	25.2(+5.3)
	Few-shot	49.6	48.2	54.5	51.3	20.3	24.8	41.4(+13.5)	41.5(+21.6)

Table 2: GPT3.5 and Llama2-7b’s performance (BLEU score) across different domains with zero-shot, few-shot, and style learning prompting. The style-learning prompting strategy improves both GPT3.5 and Llama2’s performance with only monolingual information.

Prompt Style	Fewshot				Style			
	1-shot	5-shot	10-shot	Avg	1-shot	5-shot	10-shot	Avg
Tier 0	43.4	47.9	48.1	44.9	38.5	42.2	42.8	41.2
Tier 1	36.7	38.8	39.7	38.4 (-14%)	34.6	36.1	36.6	35.7 (-13%)
Tier 2	36.2	38.1	39.2	37.8 (-16%)	34.4	35.8	36.4	35.5 (-14%)
Tier 3	36.2	38.1	39.2	37.8 (-16%)	34.2	35.7	36.0	35.3 (-14%)

Table 3: Prompt performance (BLEU scores averaged across domains; more details in Table 5) with retrieved samples of varying quality. The evaluation results show that using a single demonstration from Tier 0 is more effective than using 10 demonstrations from Tier 1.

shot over style-learning prompting to:

- The few-shot setting has access to additional paired demonstrations, which can provide better guidance, such as language alignment, compared to monolingual examples employed in style-learning prompting.
- The retriever queries the target corpus with the potentially noisy translation generated from zero-shot translation for style-learning prompting.

Despite the remaining performance gap, our style-learning prompting method demonstrates increased data efficiency and adaptability of the language model to new domains where only monolingual resources are available.

Ablation on Retrieval Quality For both few-shot and our style-learning prompting, the quality of retrieved examples is important. We retrieved the 100 most similar samples with BM25 and separated them into four groups (tier 0 to 3, with 0 being the most similar group). As shown in Table 3, for both few-shot and style-learning, they benefit more from just a single example from tier 0 than from the 10 samplers from tier 1, indicating the importance of retrieval quality. For the full results, please refer to Table 5. In Figure 5, we also show that a dense retriever (based on sentence-transformer) obtains better demonstrations for

the style-learning prompt. Though improving retrievers is orthogonal to our contribution of prompting strategy in this paper, we have shown that retriever is indeed essential in our pipeline for demonstration curation. We envision that a retriever that specifically attends to fine-grained style might be developed to further improve the style-learning prompting performance.

5 Conclusion

In this paper, we investigate the prompt-based zero- and few-shot translation performance on different domains. We attribute the performance gap to the style difference, and through our analysis, we verified that few-shot translation benefits from its retrieved samples by learning the target domain’s style. We then design a style-learning prompting strategy that only requires monolingual information and effectively bridges the gap between zero- and few-shot translation.

Limitations

Model: This paper’s experiment is currently based on the GPT3.5-0301 snapshot and LLaMa2-7b model, and we will include ICL performance from other LLMs in the future. However, we anticipate that the observed trends and findings would likely extend to other LLMs.

Translation Dataset: The translation datasets used in this study exclusively involve German-to-English translation, which represents a high-resource language pair. Consequently, the results may not be directly applicable to low-resource translation tasks or languages in which LLMs do not perform well. We leave the investigation of low-resource languages to future work.

References

- 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*, pages 7747–7763, Online. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2016. [Neural machine translation by jointly learning to align and translate](#).
- Peter F. Brown, Stephen A. Della Pietra, Vincent J. Della Pietra, and Robert L. Mercer. 1993. [The mathematics of statistical machine translation: Parameter estimation](#). *Computational Linguistics*, 19(2):263–311.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Hiram Calvo, Andrea Segura-Olivares, and Alejandro Garcia. 2014. [Dependency vs. constituent based syntactic n-grams in text similarity measures for paraphrase recognition](#). *Computación y Sistemas*, 18.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling language modeling with pathways](#).
- Thomas Cover and Peter Hart. 1967. [Nearest neighbor pattern classification](#). *IEEE Transactions on Information Theory*, 13(1):21–27.
- 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](#).
- Junjie Hu, Mengzhou Xia, Graham Neubig, and Jaime Carbonell. 2019. [Domain adaptation of neural machine translation by lexicon induction](#).
- Magdalena Jankowska. 2017. [Author style analysis in text documents based on character and word n-grams](#).
- Wenxiang Jiao, Wenxuan Wang, Jen tse Huang, Xing Wang, and Zhaopeng Tu. 2023. [Is chatgpt a good translator? yes with gpt-4 as the engine](#).
- Di Jin, Zhijing Jin, Zhiting Hu, Olga Vechtomova, and Rada Mihalcea. 2022. [Deep Learning for Text Style Transfer: A Survey](#). *Computational Linguistics*, 48(1):155–205.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- David D. McDonald and James D. Pustejovsky. 1985. [A computational theory of prose style for natural language generation](#). In *Second Conference of the European Chapter of the Association for Computational Linguistics*, Geneva, Switzerland. Association for Computational Linguistics.
- Yasmin Moslem, Rejwanul Haque, John D. Kelleher, and Andy Way. 2023. [Adaptive machine translation with large language models](#).
- OpenAI. 2023. [Gpt-4 technical report](#).
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#).
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association*

- for *Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Stephen Robertson, Steve Walker, Michel Beaulieu, and Michael Gatford. 1994. Okapi at TREC-3. In *TREC 3: Proceedings of the Third Text Retrieval Conference*, pages 109–126. National Institute of Standards and Technology (NIST).
- Suzanna Sia and Kevin Duh. 2023. [In-context learning as maintaining coherency: A study of on-the-fly machine translation using large language models](#). In *Proceedings of Machine Translation Summit XIX, Vol. 1: Research Track*, pages 173–185, Macau SAR, China. Asia-Pacific Association for Machine Translation.
- Ayush Singh and John E. Ortega. 2022. [Addressing distribution shift at test time in pre-trained language models](#).
- Efstathios Stamatatos, Nikos Fakotakis, and George K. Kokkinakis. 2000. Automatic text categorization in terms of genre and author. *Computational Linguistics*, 26:471–495.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#).
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- David Vilar, Markus Freitag, Colin Cherry, Jiaming Luo, Viresh Ratnakar, and George Foster. 2022. [Prompting palm for translation: Assessing strategies and performance](#).
- Biao Zhang, Barry Haddow, and Alexandra Birch. 2023. [Prompting large language model for machine translation: A case study](#).
- Hongxiao Zhang, Hui Huang, Jiale Gao, Yufeng Chen, Jinan Xu, and Jian Liu. 2022. [Iterative constrained back-translation for unsupervised domain adaptation of machine translation](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5054–5065, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Kaizhong Zhang and Dennis Shasha. 1989. [Simple fast algorithms for the editing distance between trees and related problems](#). *SIAM J. Comput.*, 18:1245–1262.
- Xin Zheng, Zhirui Zhang, Shujian Huang, Boxing Chen, Jun Xie, Weihua Luo, and Jiajun Chen. 2022. [Non-parametric unsupervised domain adaptation for neural machine translation](#).

A Experiment Setup

Model: For all our experiments, we use the publicly available ChatGPT *gpt-3.5-turbo-0301*¹⁰ model. For hyper-parameters, we use *temperature* 0.3 throughout our experiments, and we set all other hyper-parameters using their default values.

Evaluation Metrics: We employ BLEU (Papineni et al., 2002) and COMET (Rei et al., 2020) to compare the reference and hypothesis.

Datasets For the dataset, we rely on the domain data splits in Aharoni and Goldberg (2020), originally collected by Koehn and Knowles (2017). This corpus includes five domain-specific (Law, Medical, Koran, IT, and Subtitles) datasets for German-English translation. We did not use IT and Subtitles domains because the sentences in these two domains are noisy and short (with only 9 words/sentence for IT and 8.2 words/sentence for *Subtitles*). From our preliminary study, we also found that lots of noisy sentences are not filtered out in this dataset. Therefore we used ChatGPT to automatically filter the noisy pairs from the test set (we are not exposed to the dataset and for the details, please refer to appendix A). In this paper, we report evaluation results based on the cleaned dataset¹¹.

Domain	#Sentence	#Words/Sentence
Law	1907	28.6
Medical	1665	16.9
Koran	1629	20.3

Table 4: Test data statistics after applying GPT3.5 filter following appendix A

Data Cleaning with GPT-3.5 To clean the test set without being exposed to the sample, we first perform zero-shot translation on the devset following Figure 2. Then we rank the generated outputs by their BLEU score and take the 20 worst German-English samples (most of which are low-quality because of noise in the sentence pair, as you can find in Table 8). Note that we take the original source-target sentence pairs instead of the zero-shot output because we only use zero-shot’s performance as a scorer to obtain the worst samples. Then we prompt GPT-3.5 to evaluate the quality of these ill-formed pairs and get the evaluation results. The criteria section of the prompt shown in Table 6, Table 7, Table 8 comes from GPT-3.5’s evaluation results. Lastly, we use these prompts to ask GPT-3.5 whether each test sentence pair is a good translation and we filter out those pairs that GPT-3.5 predicts "No".

B Ablation on ICL Demonstrations’ Quality

To further understand the impact of demonstrations’ quality, we ablate on the quality of samples retrieved with the lexical retriever – BM25. We use BM25 retriever to obtain 100 most similar samples from the training corpus for few-shot translation. From the most similar 100 samples, we craft four different chunks, each of 25 samples. We rank these chunks of retrieved data as tier 0, tier 1, tier 2, and tier 3 where tier 0 consists of 25 samples that have the highest BM25 scores, tier 1 has the next 25 samples with the highest similarity score, and so on. From Table 5, we see that the performance of prompt (averaged across domain) degrades by 14% from tier 0 to tier 1, showing that the quality of high-similarity samples that are essential for the high performance. Our finding provides a different view compared to prior work (Hendy et al., 2023; Vilar et al., 2022) which finds random and searched samples (with embedding-based k-nearest-neighbor (Cover and Hart, 1967) methods) have no effect on the performance. We hypothesize that such difference comes from the nature of the task and dataset since our domain-specific dataset has more conformed styles while prior works tested on general benchmark datasets from WMT, which makes style-matching harder with few examples.

¹⁰ <https://platform.openai.com/docs/guides/chat>

¹¹ The cleaned dataset will be released at: anonymized.com

Prompt Style:		Fewshot				Style-Transfer			
Domain	Tier	1-shot	5-shot	10-shot	Avg	1-shot	5-shot	10-shot	Avg
Average	Tier 0	43.4	47.9	48.1	44.9	38.5	42.2	42.8	41.2
	Tier 1	36.7	38.8	39.7	38.4 (-14%)	34.6	36.1	36.6	35.7 (-13%)
	Tier 2	36.2	38.1	39.2	37.8 (-16%)	34.4	35.8	36.4	35.5 (-14%)
	Tier 3	36.2	38.1	39.2	37.8 (-16%)	34.2	35.7	36.0	35.3 (-14%)
Law	Tier 0	54.9	59.3	60.0	58.1	45.5	48.9	48.9	47.8
	Tier 1	43.8	47.1	48.2	46.4 (-20%)	40.7	42.3	42.8	41.9 (-12%)
	Tier 2	42.7	45.4	47.4	45.1 (-22%)	40.3	42.1	42.5	41.7 (-13%)
	Tier 3	42.8	46.5	47.6	45.6 (-21%)	40.0	42.0	42.8	41.6 (-13%)
Medical	Tier 0	57.7	64.1	63.6	57.5	50.2	57.8	59.4	55.8
	Tier 1	46.4	48.5	49.3	48.1 (-16%)	44.3	45.9	47.0	45.7 (-18%)
	Tier 2	46.5	48.0	48.7	47.7 (-17%)	44.3	45.7	46.4	45.5 (-18%)
	Tier 3	46.1	47.5	48.4	47.3 (-17%)	43.9	45.7	45.4	49.0 (-19%)
Koran	Tier 0	17.6	20.2	20.5	19.0	19.7	19.8	20.2	19.9
	Tier 1	19.8	20.6	21.7	20.7 (+9%)	18.7	20.0	20.1	19.6 (-1%)
	Tier 2	19.2	20.9	21.7	20.6 (+8%)	18.5	19.6	20.2	19.4 (-2%)
	Tier 3	19.4	20.6	21.1	20.4 (+7%)	18.5	19.5	19.9	19.3 (-3%)

Table 5: Prompt performance with retrieved samples of varying quality

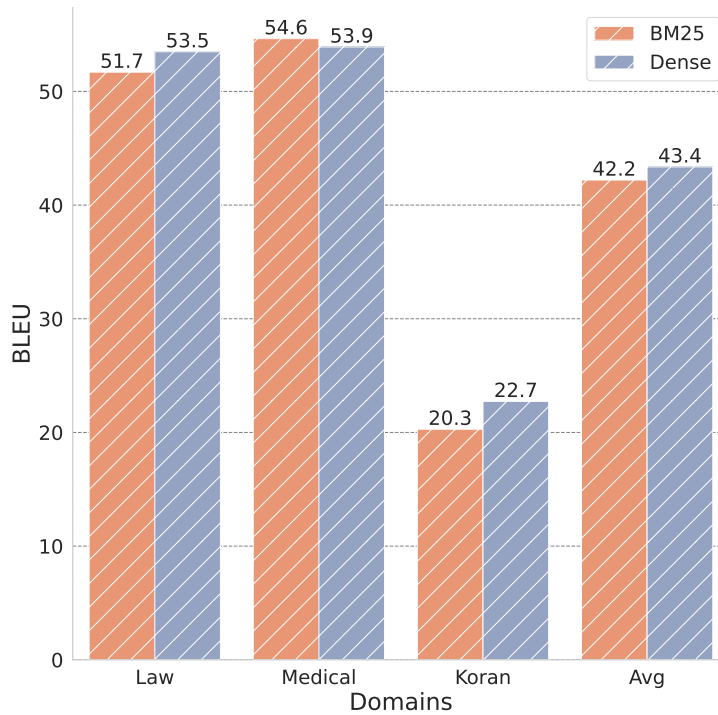


Figure 5: BM25 vs Dense retriever’s effect on style-learning prompting’s translation quality (with k=5 examples)

[Source]: Die Informationen sind zur Erfuellung der Beaufsichtigungsaufgabe nach Unterabsatz 1 bestimmt.
 [Target]: the information shall be for the purpose of performing the task of overseeing referred to in the first subparagraph.
 [0-shot]: The information is intended to fulfill the supervisory task according to subparagraph 1.

Few-shot Demo:

[German]: Die Informationen sind zur Erfuellung der Aufgabe nach Unterabsatz 1 bestimmt.

[English]: the information shall be for the purpose of performing the task referred to in the first subparagraph

[1-shot translation]: The information is intended for the purpose of fulfilling the supervisory task referred to in subparagraph 1.

Style Demo:

Such exchange of information must be intended for the performance of the supervisory task of the authorities or bodies mentioned

[1-shot style transfer]: The information is intended for the performance of the supervisory task under paragraph 1.

Figure 6: The example shown here is an illustration of the improvement from few-shot and style-learning translation. The text in orange is the text that GPT3.5 learns from the demonstration, which results in better lexical overlap and higher syntactic similarity with the reference text. In this example, the Few-shot prompt’s demo has better quality as it overlaps more with the target sentence (in fact, the whole sentence can be found in the target sentence, as shown with the underline). The retriever for the few-shot prompts has better quality as it is able to utilize both the source and target language’s corpus.

7.5	[reference]: Whereas, under Article 34 (1bis), the Convention is open to accession by regional economic integration organizations [few-shot]: <u>Under Article 34 (1 bis), the Convention is open to accession</u> by regional economic integration organizations [style-learning]: <u>According to Article 34, paragraph 1, any regional organization for economic integration</u> <u>can become a party to the agreement</u>
5.1	[reference]: It is considered that the imposition of measures will restore fair competition on the market. [few-shot]: It is <u>considered</u> that the <u>imposition</u> of measures will restore fair competition to the market. [style-learning]: It is <u>believed</u> that the <u>introduction</u> of measures will restore fair competition to the market.
2.7	[reference]: they must come from holdings, centres or organizations which are subject to regular official veterinary checks in accordance with paragraph [few-shot]: They must come from a <u>holding</u> , centre or establishment which is subject to regular official veterinary <u>checks</u> in accordance with paragraph 3. [style-learning]: They must come from an establishment, center, or <u>facility</u> subject to regular official veterinary <u>inspections</u> in accordance with paragraph 3
0.7	[reference]: The Cooperation Committee shall take decisions by common accord. [few-shot]: The Cooperation Committee <u>takes</u> its decisions unanimously. [style-learning]: The Cooperation Committee <u>adopts</u> its decisions unanimously.

Figure 7: Here we provide 4 examples of comparison of COMET scores between few-shot and style-learning prompting outputs. The number on the left is the COMET score difference and the highlighted words are the major difference between the two outputs. We see that all these samples have similar semantics and the difference is mostly from lexical choice (except for the 7.5 cases where sentence structure is also changed but the semantics is still very similar). Therefore, we believe zero-and few-shot outputs’ performance gap originates from styles rather than semantics.

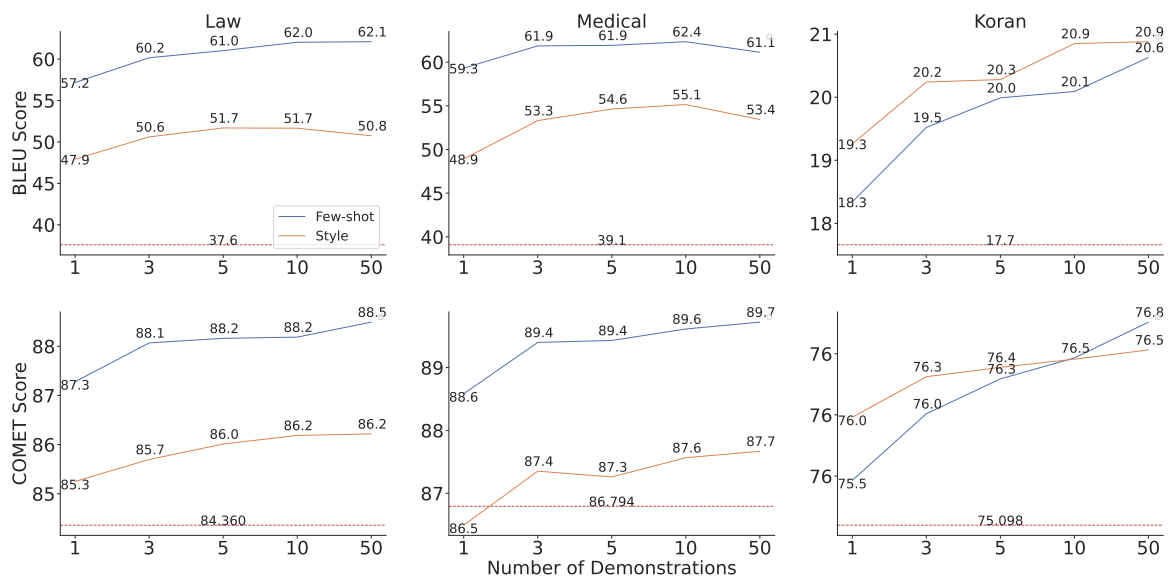


Figure 8: Full Results of few-shot and style-learning prompting across three domains, evaluated by BLEU and COMET (COMET score is multiplied by 100 for better visualization). The x-axis is the number of demonstrations used in the prompt and the y-axis is the evaluation result. The dashed red line is the performance of zero-shot translation and we see that both prompts improve the domain-specific translation even just given 1 demo.




Figure 9: *Left*: 3-gram overlap analysis between *translations* and *retrieved examples*. Compared to zero-shot, the output from style-learning prompting has a much higher match rate with retrieved samples. *Mid*: Average Tree Edit Distance (TED) between the dependency parse tree of reference and outputs from zero-shot or style-learning prompting. Style-learning prompting results in a lower TED, showing its more grammar with the reference. *Right*: Average TED between the constituency parse of reference and zero-shot/style-learning outputs. Again, the style-learning prompting achieves lower TED and has a more conformed syntactical structure with the reference.

You are a human evaluator who judges the quality of parallel data for German-English translation. Below are some examples of low-quality parallel data along with the criteria for filtering them out. [German]: DIESE HINTERBLIEBENENVERSORGUNG ENTSPRICHT :

[English]: THAT PENSION SHALL BE EQUAL TO A PERCENTAGE OF THE PENSION ACCRUING TO THE MEMBER OR FORMER MEMBER OF THE COMMISSION OR OF THE COURT UNDER ARTICLE 9 AT THE DATE OF DEATH, NAMELY:

[Criteria]: The translation is overly verbose and does not accurately convey the meaning of the German sentence. The English translation includes unnecessary repetitions and lacks clarity.

[German]: 3 2 0 0 BAR Information und Veröffentlichung BAR 200000 BAR BAR 200000 BAR

[English]: 3 2 0 0 BAR Information and publishing BAR 200000 BAR BAR 200000 BAR

[Criteria]: The translation includes non-translated elements (such as "BAR") that do not provide any meaningful information in the target language. It appears to be a result of incorrect processing or formatting.

[German]: - gelegentlich erfolgen,

[English]: - are of an occasional nature,

[Criteria]: The translation fails to capture the meaning of the German sentence accurately. It provides a more general interpretation, which does not convey the intended sense of occasional occurrences.

[German]: a) höher sein als die Einheit,

[English]: (a) be greater than 1;

[Criteria]: The translation does not accurately convey the meaning of the German sentence. It provides a different interpretation, suggesting that the value should be greater than 1, whereas the German sentence simply states "higher than the unit."

[German]: - Exportação para a Polónia.

[English]: - Exportação para a Polónia.

[Criteria]: The translation is not in English but rather includes Portuguese words. It appears to be a mistake or a mix-up between different language pairs."""

Given the criteria above, is following sentence pair a good translation? Output Yes if it is a good translation, output No if it is a bad translation.

[German]: <test source sentence> [English]: <test target sentence>

Table 6: Filtering Prompt for Law Domain

You are a human evaluator who judges the quality of parallel data for German-English translation. Below are some examples of low-quality parallel data along with the criteria for filtering them out.

[German]: Brivudin behandelt wurden.
[English]: Xeloda should also not be used in the following groups...
[Criteria]: Incompleteness or Omission: When important information is missing in the translation or not adequately conveyed.

[German]: 12/23
[English]: 12/ 22
[Criteria]: Irrelevant or Unrelated Information: When the translation includes information that is not relevant to the original sentence.

[German]: Midazolam zur Injektion).
[English]: Examples include the cholesterol-reducing agent atorvastatin, the antibiotics...
[Criteria]: Sentence Structure and Syntax: When the translation has incorrect word order or sentence structure.

[German]: Hyperthyreose
[English]: Thyroid dysfunction most often presenting as hypothyroidism or hyperthyroidism
[Criteria]: Incorrect Terminology or Inaccurate Vocabulary: When the translation uses incorrect or inaccurate terms.

[German]: 1120
[English]: 112 Further training, language courses...
[Criteria]: Mistranslation or Misinterpretation: When the translation conveys a different meaning than the original sentence.

[German]: Die Wirksamkeit von Revasc als Gerinnungshemmer wurde in vier Studien...
[English]: ho death/ re-MI at Day 30 were not statistically different...
[Criteria]: Inconsistent Terminology: When the translation uses inconsistent or contradictory terms.

[German]: 7,9 $\mu\text{mol/l}$.
[English]: Ki values in human liver microsomes were 27, 7.5 and 7.9 $\mu\text{mol/l}$, respectively.
[Criteria]: Grammatical Errors: When the translation contains grammatical mistakes or incorrect usage of language rules.

[German]: Uber diesen Minimalwert hinausgehende
[English]: The recommendation of a minimum yield of 2.0×10^6 CD34+ cells/ kg
[Criteria]: Lack of Clarity or Ambiguity: When the translation is unclear or ambiguous, making it difficult to understand the intended meaning.

Given the criteria above, is following sentence pair a good translation? Output Yes if it is a good translation, output No if it is a bad translation.

[German]: <test source sentence> [English]: <test target sentence>

Table 7: Filtering Prompt for Koran Domain

You are a human evaluator who judges the quality of parallel data for German-English translation. Below are some examples of low-quality parallel data along with the criteria for filtering them out.

[German]: Doch, mit Sicherheit!
[English]: Why not?
[Criteria]: The English translation does not capture the meaning of the German sentence, which should be translated as "Certainly, for sure!" or "Yes, definitely!" instead of "Why not?"

[German]: Dahin werdet ihr kommen müssen.
[English]: The Prophets like Eisa and Uzair who were worshipped are exempt from this, and so are Maryam, and trees and the moon etc.)
[Criteria]: The English translation is completely unrelated to the original German sentence. It introduces a different topic and provides irrelevant information. A more accurate translation would be "There you will have to go."

[German]: " Wir sind zugrunde gerichtet!
[English]: (And say:) "We have fallen into debt;
[Criteria]: The English translation does not convey the intended meaning of the German sentence. A more coherent translation would be "We are destroyed!" or "We are ruined!"

Given the criteria above, is following sentence pair a good translation? Output Yes if it is a good translation, output No if it is a bad translation.

[German]: <test source sentence> [English]: <test target sentence>

Table 8: Filtering Prompt for Medical Domain