
Human-in-the-loop Machine Translation with Large Language Model

Xinyi Yang*

nlp2ct.xinyi@gmail.com

Runzhe Zhan*

nlp2ct.runzhe@gmail.com

Derek F. Wong †

derekw@um.edu.mo

Junchao Wu

nlp2ct.junchao@gmail.com

Lidia S. Chao

lidiasc@um.edu.mo

NLP²CT Lab, Department of Computer and Information Science, University of Macau

Abstract

The large language model (LLM) has garnered significant attention due to its in-context learning mechanisms and emergent capabilities. The research community has conducted several pilot studies to apply LLMs to machine translation tasks and evaluate their performance from diverse perspectives. However, previous research has primarily focused on the LLM itself and has not explored human intervention in the inference process of LLM. The characteristics of LLM, such as in-context learning and prompt engineering, closely mirror human cognitive abilities in language tasks, offering an intuitive solution for human-in-the-loop generation. In this study, we propose a human-in-the-loop pipeline that guides LLMs to produce customized outputs with revision instructions. The pipeline initiates by prompting the LLM to produce a draft translation, followed by the utilization of automatic retrieval or human feedback as supervision signals to enhance the LLM’s translation through in-context learning. The human-machine interactions generated in this pipeline are also stored in an external database to expand the in-context retrieval database, enabling us to leverage human supervision in an offline setting. We evaluate the proposed pipeline using the GPT-3.5-turbo API on five domain-specific benchmarks for German-English translation. The results demonstrate the effectiveness of the pipeline in tailoring in-domain translations and improving translation performance compared to direct translation instructions. Additionally, we discuss the experimental results from the following perspectives: 1) the effectiveness of different in-context retrieval methods; 2) the construction of a retrieval database under low-resource scenarios; 3) the observed differences across selected domains; 4) the quantitative analysis of sentence-level and word-level statistics; and 5) the qualitative analysis of representative translation cases.

Keywords: Machine Translation, Large Language Model, Human-in-the-loop, In-context Learning, Prompt Engineering, Natural Language Processing

1 Introduction

Large language models (LLMs) have exhibited remarkable proficiency in comprehending natural language prompts (OpenAI, 2023; Touvron et al., 2023), enabling them to execute vari-

*Equal Contribution. Xinyi Yang contributes to the experiments, data curation, and analysis. Runzhe Zhan contributes to the methodology, code skeleton, and paper drafting.

† Corresponding Author.

ous controllable generation tasks based on human instructions. Furthermore, LLMs can acquire knowledge from limited demonstrations that are relevant to the input data and generate desired outputs through analogy. This paradigm, known as in-context learning (ICL) (Dong et al., 2023), represents a significant advancement in prompt engineering and offers insights into adapting LLMs to downstream tasks without the need for fine-tuning the models.

Machine translation (MT) serves as a representative sequence-to-sequence task that also requires tailoring models to produce domain-specific translations. Traditional approaches to building domain-specific MT models involve fine-tuning pre-trained models with domain data or utilizing domain adaptation techniques to transfer in-domain MT models to out-of-domain requirements. However, these methods are suited for accessible, medium-scale MT models, which may not be suitable for LLMs. Notably, certain LLMs available through application programming interfaces (API) lack accessible weight matrices. Furthermore, optimizing LLM parameters with domain data can be expensive under resource-limited scenarios. Consequently, current research on LLM-based MT predominantly focuses on ICL, including in-context selection methods (Agrawal et al., 2022), in-context prompt engineering (Zhang et al., 2023), and the systematic evaluation of LLM-based MT (Hendy et al., 2023; Jiao et al., 2023). While these lines of research present empirical studies investigating ICL in the MT task, they still exhibit a dearth of exploration in customizing LLMs for domain-specific needs. Given the representative capabilities of LLMs, which rely on generating outputs based on provided instructions, it is intuitive to leverage LLMs to refine general MT outputs for different domains.

Nevertheless, there remain challenges and issues when utilizing ICL to adapt LLMs for domain-specific needs. Firstly, ICL demonstrations for MT typically comprise source input and target reference, lacking domain features. Secondly, the ICL retrieval database is usually constructed using separate labeled data, and the retrieved demonstrations fail to capture LLMs’ translation preferences. Last but not least, the adaptation of black-box LLMs does not benefit from parameter optimization, thereby limiting adaptation methods to modifying ICL inputs alone. In response to these challenges, we propose integrating LLM-specific translation feedback into ICL inputs, enabling the model to learn from both relevant input-output pairs and domain preferences.

Specifically, the proposed pipeline consists of two essential parts: feedback collection, and in-context refinement. To collect LLM-specific feedback associated with domains, we first request the LLM to produce domain-specific translations and obtain feedback by comparing its translation with a reference translation. The feedback takes the form of a sequence of revision instructions, indicating the necessary edits to transform the LLM’s translation into the reference translation. Ideally, these revision instructions originate from human feedback sources. However, due to resource limitations, we simulate this process and generate synthetic human feedback in this study. Subsequently, these translation texts and feedback are stored together in the ICL retrieval database. For in-context refinement, when faced with new in-domain translation requests, the pipeline initially prompts the LLM to generate a draft translation. It then retrieves similar translation pairs and their revision histories as in-context demonstrations tied to the specific domain. Finally, the model refines the draft translation based on the retrieved domain-specific demonstrations. Any new human-machine interactions generated within this pipeline are incorporated to expand the in-context retrieval database. Overall, the primary concept revolves around enabling the LLM to revise its outputs by learning from relevant domain-specific revision feedback.

We conduct experiments using the proposed pipeline in the German-English translation direction across five domains, utilizing the GPT-3.5 Turbo API as a testbed for the black-box LLM. The results demonstrate that the proposed pipeline enhances domain translation performance in selected domains, as indicated by four automated evaluation metrics. To further elu-

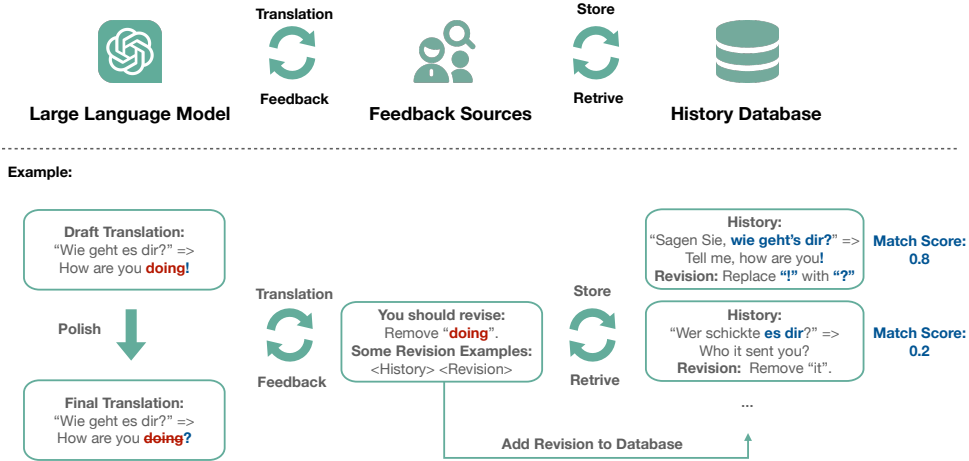


Figure 1: Illustration of the proposed human-in-the-loop translation method in the context of the large language model.

to evaluate the effectiveness of our approach, we discuss the pipeline and results through an ablation study, quantitative analysis, and qualitative analysis.

2 Methodology

We leave the discussion of related work in the Appendix A due to page limitation. The overall pipeline is illustrated in Figure 1. The feedback retrieved from a data store aims to provide domain-specific revision demonstrations to LLM. Furthermore, it can also be jointly used with external human supervision in real-world applications. In addition to the feedback that exists in the database, any newly produced feedback for current text will be also recorded into the database to be used as a candidate for ICL retrieval in the future.

2.1 Feedback Collection

To correct LLM’s bias in domain translation through the ICL paradigm, we must first construct an ICL retrieval database that reflects the gap between LLM’s translation preferences and domain preferences. To do this, we ask LLM to translate several domain texts first, and then use automated methods or human intervention to generate feedback on LLM’s translation. To simulate the process of human feedback, we use an automated evaluation method based on edit distance theory to generate feedback for the translations. Computing edit distance is a dynamic programming problem, where the cost matrix reflects what editing operations are needed to transform from the machine translation to the reference translation. Specifically, bottom-up recursion yields the minimum cost of editing the translation, so we can generate human-like feedback by back-tracing the optimal alignment of the cost matrix and converting it to natural language. There are three kinds of editing operations are considered in our case: deletion, insertion, and substitution. Given an LLM’s translation h and reference translation r , the cost matrix $D(i, j)$ indicates the edit distance between $h_{<i}$ and $r_{<j}$. Let the cost of deletion, insertion, and




Figure 2: An example of describing the optimal path of minimum edit distance.

substitution be 1, the cost matrix $D(i, j)$ can be calculated as:

$$\begin{aligned}
D(i, 0) &= i \\
D(0, j) &= j \\
D(i, j) &= \begin{cases} D(i-1, j-1) & \text{if } h_i = r_j \\ \min \begin{cases} D(i-1, j) + 1 \\ D(i, j-1) + 1 \\ D(i-1, j-1) + 1 \end{cases} & \text{if } h_i \neq r_j \end{cases} \quad (1)
\end{aligned}$$

We can obtain an optimal path of minimum edit distance by back-tracing the cost matrix D and then convert this path to the natural language instruction as shown in Figure 2. Once the feedback of a specific test instance is produced, we combine it with the source text, LLM’s translation, and reference translation to form the new ICL demonstration instance. This enables LLM to learn from its deviations from in-domain reference translations and how to polish the draft translation.

2.2 In-context refinement

We construct a database for ICL retrieval using the automatic feedback generation method mentioned in the previous section. In this section, we will describe how to retrieve and utilize these ICL demonstrations as well as feedback records for conducting a two-stage translation.

Demonstration Retrieval We retrieve the relevant ICL demonstrations by evaluating the relevance between test instances and examples stored in the database. Specifically, we follow the previous exploratory work on ICL for MT and use two retrieval metrics: the BM25 score (Robertson et al., 1995) and the BM25-rerank score (Agrawal et al., 2022). BM25 is a common retrieval metric that evaluates the relevance between the query and the documents. Here we consider each source sentence in the ICL database as a document and calculate the BM25 score. BM25-Rerank, as a post-screening method for BM25, first selects the N samples with the highest BM25 scores and re-scores these samples based on the n-gram recall score R to select the top- K samples as ICL demonstrations. Let the input in source language be s , and a demonstration candidate stored in the database be c , then the recall score can be calculated as:

$$R = \exp \frac{1}{n} \sum_{i=1}^n \log \frac{\text{Count}(i\text{-gram} \in s \cap c)}{\text{Count}(i\text{-gram} \in s)} \quad (2)$$

Two-stage Translation As mentioned in the previous section, we first ask the LLM to generate a draft translation, and then polish the draft translation by providing the ICL demonstrations,

resulting in a two-stage translation process. In practice, we implement the two-stage translation through the multi-turn dialog feature of the GPT-3.5 API. In addition, we empirically found that the LLM may incorrectly modify the draft translation after observing the ICL demonstrations. We attribute this phenomenon to the fact that the scale limitation of the retrieval database, which may lead to the inclusion of some irrelevant samples. Therefore, we asked the LLM further compare the polished translations with the draft translations at the second stage, and finally select the higher quality one. It is worth noting that the LLM did not observe the reference translation of the test instance during the whole process but only compared the translation quality by means of self-reflection.

3 Experiments

3.1 Data and Evaluation

We verify the effectiveness of the proposed pipeline on a multi-domain German-English translation benchmark (Aharoni and Goldberg, 2020). The test data keeps the same setting as the same benchmark, which involves five domains including IT, Koran, Law, Medical, and Subtitles. We randomly sample 2,000 samples from the training set of each domain as the source of constructing the ICL retrieval database. To complete this process, we use GPT-3.5 API to generate the translation and employ the edit-distance-based method mentioned in Section 2.1 to produce human-like feedback. We use several automated metrics to evaluate the translation quality, including BLEU (Papineni et al., 2002), TER (Snover et al., 2006), BERTScore (Zhang et al., 2020), and COMET (Rei et al., 2020). BLEU and TER are the traditional metrics that evaluate the text overlap whereas the others can evaluate the semantic overlap based on neural networks. We also found that APIs sometimes produce hallucinations or refuse to translate some sentences. To make a fair comparison, we manually check the translation results and remove the invalid results, and only evaluate the performance of the sentences that are successfully translated by all the methods.

3.2 Settings

The experiments were conducted using GPT-3.5 API. The decoding temperature and the top_p parameters are set to 1 by default. When evaluating the relevance of demonstrations, we first retrieve the top $K=200$ demonstrations with the highest BM25 scores and then select top- N demonstrations with the highest 4-gram re-rank scores as the finalized ICL demonstrations.

3.3 Main Results

Table 1 presents the automated evaluation results of the baseline methods. The experimental findings demonstrate that the proposed approach effectively enhances the performance of GPT-3.5-Turbo baseline. Importantly, we observe that the impact of the proposed HIL method varies across different domains, a topic that will be thoroughly examined in Section 4.2. In addition, we were limited to varying the number of ICL demonstrations from 1 to 3 due to constraints on request tokens. Nonetheless, the results strongly indicate that providing more ICL demonstrations leads to improved performance. Moreover, when evaluating the performance with neural metrics, the differences in scores are not substantial compared to the traditional metrics. We postulate that the process of refining the draft translation may not deviate significantly from the original semantic content but rather brings it closer to specific translation preferences in certain domains. These findings will be elucidated through a detailed case study in the subsequent section.

| | IT | | | | Koran | | | |
|----------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | BLEU | TER | BERT-F | COMET | BLEU | TER | BERT-F | COMET |
| GPT-3.5-Turbo | 34.4 | 62.3 | 93.2 | 82.5 | 16.2 | 74.1 | 90.6 | 73.4 |
| +1-shot HIL | 29.0 | 77.3 | 92.5 | 81.1 | 15.7 | 80.8 | 90.5 | 72.8 |
| +2-shot HIL | 33.9 | 64.7 | 92.9 | 82.3 | 15.8 | 77.7 | 90.6 | 73.2 |
| +3-shot HIL | 32.6 | 68.8 | 93.0 | 82.2 | 16.5 | 76.0 | 90.7 | 73.6 |
| +Compare HIL | 35.2 | 61.3 | 93.3 | 82.8 | 16.6 | 74.6 | 90.7 | 73.8 |
| | Law | | | | Medical | | | |
| GPT-3.5-Turbo | 37.6 | 54.7 | 93.7 | 83.8 | 40.0 | 59.4 | 93.9 | 83.4 |
| +1-shot HIL | 36.2 | 59.6 | 93.5 | 83.1 | 36.6 | 67.9 | 93.3 | 82.1 |
| +2-shot HIL | 37.0 | 56.6 | 93.6 | 83.5 | 39.2 | 63.0 | 93.6 | 83.0 |
| +3-shot HIL | 36.7 | 56.5 | 93.6 | 83.7 | 38.4 | 63.0 | 93.7 | 83.2 |
| +Compare HIL | 37.7 | 54.5 | 93.8 | 83.9 | 40.9 | 60.1 | 93.9 | 83.6 |
| | Subtitles | | | | | | | |
| GPT-3.5-Turbo | 27.9 | 64.8 | 93.0 | 80.0 | | | | |
| +1-shot HIL | 26.3 | 69.4 | 92.5 | 78.8 | | | | |
| +2-shot HIL | 26.4 | 67.1 | 92.6 | 79.4 | | | | |
| +3-shot HIL | 27.4 | 64.6 | 93.0 | 79.8 | | | | |
| +Compare HIL | 28.0 | 64.1 | 93.1 | 80.1 | | | | |

Table 1: Automated evaluation results of different translation strategies. “ K -shot HIL” indicates the proposed HIL method with K demonstrations used. “Compare HIL” indicates using the comparison strategy to finalize the two-stage translation.

4 Analysis

4.1 Effects of ICL Retrieval Methods

To explore the potential impact of different demonstration retrieval methods on our proposed HIL translation workflow, we conducted experiments using two strategies: BM25 and BM25 Re-rank, in a 3-shot translation scenario. The comparative results are summarized in Table 2. Based on the automated metrics, both BM25 and BM25 Re-Rank methods exhibited strengths and weaknesses in various domains. Specifically, BM25 Re-Rank slightly outperformed BM25 in terms of the BLEU metric. The advantage of BM25 Re-Rank was particularly evident in the IT domain, as it achieved higher scores than BM25 across all metrics. However, this conclusion was reversed in the Law domain. The advantage of the BM25 Re-Rank strategy lies in its ability to filter out repetitive context with identical queries. Consequently, this approach selects more relevant examples related to the IT domain, leading to an improvement in the quality of the translation output within this domain. The BM25 method for demonstration retrieval focuses on document frequency and keyword matching, making it more effective in ensuring proper usage of legal terminology. This observation underscores the necessity of adopting different strategies for demonstration retrieval across different domains to ensure the selection of contextually relevant and domain-specific examples for the target sentences.

4.2 Domain Differences

In general, the HIL approach exhibits superior performance compared to the GPT-3.5 API baseline across all five domains, with particularly notable advantages in the IT and Medical domains. However, the differences in performance are relatively smaller in Law and Subtitles domains. These variations can be attributed to the distinct sentence styles and structures prevalent in each

| Method | BM25 | | | | BM25 Re-Rank | | | |
|-----------|-------------|-------------|-------------|-------------|--------------|-------------|-------------|-------------|
| | BLEU | TER | BERT-F | COMET | BLEU | TER | BERT-F | COMET |
| IT | 34.9 | 62.3 | 93.1 | 82.5 | 35.2 | 61.3 | 93.3 | 82.8 |
| Koran | 16.2 | 74.3 | 90.7 | 73.7 | 16.6 | 74.6 | 90.7 | 73.8 |
| Law | 38.0 | 54.2 | 93.8 | 84.2 | 37.7 | 54.5 | 93.8 | 83.9 |
| Medical | 40.6 | 58.8 | 94.0 | 83.8 | 40.9 | 60.1 | 94.0 | 83.6 |
| Subtitles | 28.2 | 64.5 | 93.0 | 80.2 | 28.0 | 64.1 | 93.1 | 80.1 |

Table 2: Automated evaluation results for 3-shot HIL with different ICL retrieval strategies.

domain. Upon analyzing the translation results, it becomes evident that HIL excels in IT and Medical domains by effectively aligning terminology with the reference translations. For example, consider the phrase “Returns a character string” in the IT domain, HIL correctly recognizes the need to use the third-person singular form of the word “returns” and avoids translating “character string” simply as “string”. While these words may not be technical terms, their specific usage preferences in the IT domain are crucial, and such nuances cannot be captured by the GPT-3.5 API baseline. Conversely, in domains such as Law and Subtitles, where HIL’s performance is comparatively lower, the sentences tend to adhere to specific legal clauses or follow a more colloquial and concise style. As GPT-3.5 is a multi-domain language model, it already possesses substantial knowledge related to these domains, leading to satisfactory draft translations in the initial output, thereby reducing the necessity for extensive corrections through demonstrations. Furthermore, it is worth noting that the quality of the data within the demonstration pool may not be very high, as they were randomly sampled from the training set. This behavior could also have a negative impact on the final results.

4.3 Quantitative Analysis

We conducted quantitative analysis on the translation results with the help of `compare-mt`¹ (Neubig et al., 2019) toolkit.

Part of Speech (POS) We applied Stanford’s POS tagging toolkit (Toutanova et al., 2003) to label the target-side text and subsequently examined the translation outcomes of both the baselines and HIL method across different POS categories. The results are presented in Figure 3. Overall, HIL outperforms the baselines in all noun categories, including singular or mass nouns (NN), plural nouns (NNS), and singular proper nouns (NNP). Additionally, among the three verb types, HIL exhibited superior performance to the baseline in base form verbs (VB) and third person singular present verbs (VBZ), with the most noticeable advantage observed in base form verbs (VB). Furthermore, the advantage of HIL is particularly evident in POS categories where the baseline exhibits lower translation accuracy. These findings underscore the effectiveness of the HIL approach in handling various parts of speech and its ability to deliver improved translation outputs, especially in challenging linguistic contexts.

Sentence Length Figure 4 presents the translation performance of both the baseline and HIL models for sentences of varying lengths. HIL exhibits a clear advantage over the baseline for sentences that are shorter than 10 words as well as those longer than 60 words. Unfortunately, HIL performs less effectively for medium-length sentences compared to the baseline’s initial draft. This observation indicates that there might be challenges specific to this sentence length range that warrant further investigation and potential refinement of the HIL approach.

¹<https://github.com/neulab/compare-mt>




Figure 3: The translation comparison in terms of different POS tags.

In conclusion, our analysis of translation performance across different POS categories and sentence lengths reveals that HIL exhibits exceptional proficiency in translating both nouns and verbs, particularly excelling in extreme-short and extreme-long sentences.




Figure 4: The translation comparison across various sentence lengths.

4.4 Case Study

Below, we present an illustrative example from the IT domain to showcase the advantages of HIL in terms of terminology translation. Additionally, we compare the proposed HIL method with the ordinary ICL method without revision feedback. The results are shown in Table 3. In the given example, while the German word “Handbuch” was translated as “manual” in both the draft and ICL translations, the reference and retrieved demonstration suggest that “Handbook” is a more accurate translation. Regarding the terminology usage of the IT domain, “Handbook” precisely describes a professional document type, whereas “manual” might be more generic and provide less specific information. Among the three translation methods, the HIL translation successfully captures this crucial information based on the provided demonstrations. On the

| IT Domain | |
|----------------|--|
| Source | Das Handbuch zu & ksnapshot; |
| Reference | The & ksnapshot; Handbook |
| Demonstrations | <ol style="list-style-type: none"> 1. <input>Das Handbuch zu & kontakt; <hypothesis>The manual for & kontakt; <reference>& kontakt; Handbook <revision>“the” should be deleted. <u>“manual” should be deleted.</u> “for” should be deleted.”, “handbook” should be inserted after “kontakt;” 2. <input>Das Handbuch zu & kanagram; <hypothesis>The man-ual for & kanagram; <reference>& kanagram; Handbook <revision>“the” should be deleted. <u>“manual” should be deleted.</u> “for” should be deleted. “handbook” should be inserted after “kanagram;” 3. ... |
| Draft | The manual for & ksnapshot; |
| ICL | The manual for &ksnapshot; |
| HIL | The & ksnapshot; Handbook |

Table 3: An example result of three different translation strategies. “Draft” represents the preliminary translation results obtained at the initial turn in our HIL pipeline. “ICL” presents the translation results achieved using ordinary ICL demonstrations without revision feedback.

other hand, HIL also learns from revisions in the demonstration and opts to translate the original sentence as “The &ksnapshot; Handbook”, aligning more accurately with the word order of the original text. This demonstrates how HIL can effectively incorporate valuable revision feedback to produce more contextually appropriate and accurate translations.

5 Conclusions and Future Work

In this paper, we present an empirical study focused on enhancing the translation capabilities of the LLM by integrating concrete feedback within the translation process. Our objective is to establish a human-in-the-loop machine translation pipeline, where human feedback plays a pivotal role. To simulate this concept, we utilize an automated feedback method, leveraging the GPT-3.5 API as our testbed, which yields effective results. In the future, our plan is to collect human feedback to create a novel dataset and conduct experiments on the proposed pipeline using this dataset. This approach will enable us to further validate and implement our human-in-the-loop machine translation system in the context of LLM, enhancing its practical applicability and performance.

Acknowledgement

This work was supported in part by the Science and Technology Development Fund, Macau SAR (Grant Nos. FDCT/0070/2022/AMJ, FDCT/060/2022/AFJ), the Multi-year Research Grant from the University of Macau (Grant No. MYRG2020-00054-FST), and the Research Program of Guangdong Province (Grant No. 2220004002576). This work was performed in part at SICCC which is supported by SKL-IOTSC, and HPCC supported by ICTO of the University of Macau.

References

- Agrawal, S., Zhou, C., Lewis, M., Zettlemoyer, L., and Ghazvininejad, M. (2022). In-context examples selection for machine translation. *ArXiv preprint*, abs/2212.02437.
- Aharoni, R. and Goldberg, Y. (2020). Unsupervised domain clusters in pretrained language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7747–7763. Association for Computational Linguistics.
- Dong, Q., Li, L., Dai, D., Zheng, C., Wu, Z., Chang, B., Sun, X., Xu, J., and Sui, Z. (2023). A survey for in-context learning. *ArXiv preprint*, abs/2301.00234.
- Hendy, A., Abdelrehim, M. G., Sharaf, A., Raunak, V., Gabr, M., Matsushita, H., Kim, Y. J., Afify, M., and Awadalla, H. H. (2023). How good are gpt models at machine translation? a comprehensive evaluation. *ArXiv preprint*, abs/2302.09210.
- Jiao, W., Wang, W., tse Huang, J., Wang, X., and Tu, Z. (2023). Is chatgpt a good translator? yes with gpt-4 as the engine. volume abs/2301.08745.
- Neubig, G., Dou, Z.-Y., Hu, J., Michel, P., Pruthi, D., and Wang, X. (2019). compare-mt: A tool for holistic comparison of language generation systems. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 35–41. Association for Computational Linguistics.
- OpenAI (2023). Gpt-4 technical report. *ArXiv preprint*, abs/2303.08774.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (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. Association for Computational Linguistics.
- Rei, R., Stewart, C., Farinha, A. C., and Lavie, A. (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. Association for Computational Linguistics.
- Robertson, S. E., Walker, S., Jones, S., Hancock-Beaulieu, M. M., Gatford, M., et al. (1995). Okapi at trec-3. *Nist Special Publication Sp*, 109:109.
- Snover, M., Dorr, B., Schwartz, R., Micciulla, L., and Makhoul, J. (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. Association for Machine Translation in the Americas.
- Toutanova, K., Klein, D., Manning, C. D., and Singer, Y. (2003). Feature-rich part-of-speech tagging with a cyclic dependency network. In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 252–259.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., and Lample, G. (2023). Llama: Open and efficient foundation language models. *ArXiv preprint*, abs/2302.13971.
- Wang, D., Wei, H., Zhang, Z., Huang, S., Xie, J., and Chen, J. (2022). Non-parametric online learning from human feedback for neural machine translation. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 11431–11439. AAAI Press.

Wu, X., Xiao, L., Sun, Y., Zhang, J., Ma, T., and He, L. (2022). A survey of human-in-the-loop for machine learning. *Future Generation Computer Systems*, 135:364–381.

Zhang, B., Haddow, B., and Birch, A. (2023). Prompting large language model for machine translation: A case study. *ArXiv preprint*, abs/2301.07069.

Zhang, T., Kishore, V., Wu, F., Weinberger, K. Q., and Artzi, Y. (2020). Bertscore: Evaluating text generation with BERT. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

A Appendix: Related Work

A.1 ICL-based MT

Agrawal et al. (2022) proposed an ICL example selection method for machine translation, aiming to explore the impact of ICL examples on translation output quality. To address the issues in existing BM25 retrieval methods, the proposed approach re-ranks the top 100 candidate sentences selected by the BM25 score and introduces a re-ranking score, thereby maximizing the coverage of input words. Zhang et al. (2023) aims to explore different prompts in the context of LLM-based machine translation. In comparison to previous related research, the main innovation of this study lies in exploring how to design prompts for LLMs to enhance their translation capabilities from three different perspectives. Specifically, the research investigates different prompting strategies, the utilization of unlabeled data, and the flexibility of transfer learning. These research achievements demonstrate the significant value of ICL in improving the quality of machine translation systems.

A.2 Human-in-the-Loop MT

The concept of Human-in-the-Loop (HIL) (Wu et al., 2022) aims to leverage user feedback for optimizing the model. Building on this idea, Wang et al. (2022) proposes a novel non-parametric online learning method called kNN-over-kNN (KoK) that does not alter the model structure. KoK is a plug-and-play non-parametric approach that learns online based on human feedback, reducing the number of user interactions and improving machine translation model performance. The online learning process of KoK involves three steps: decoding, correcting, and adapting. In the decoding phase, the MT system translates the source sentence, and the output is obtained by weighting the KoK method and kNN-MT. In the correcting phase, users provide corrections of the machine-translated text, resulting in the post-edited translation. Finally, in the adapting phase, post-edited and source sentences are jointly used to expand the data repositories of token-kNN and policy-kNN, thereby optimizing the model. Through these three steps, the KoK framework can promptly influence the kNN translation model’s decision-making. It is worth noting that, as of now, the HIL method has not been applied to LLM translation. Our paper focuses on exploring the potential integration of HIL into LLM translation to further enhance its performance and capabilities.