

Beyond Human-Only: Evaluating Human-Machine Collaboration for Collecting High-Quality Translation Data

Zhongtao Liu, Parker Riley, Daniel Deutsch, Alison Lui,
Mengmeng Niu, Apu Shah and Markus Freitag

Google

{zhongtao,prkriley,danddeutsch,alisonlui,mniu,apurva,freitag}@google.com

Abstract

Collecting high-quality translations is crucial for the development and evaluation of machine translation systems. However, traditional human-only approaches are costly and slow. This study presents a comprehensive investigation of 11 approaches for acquiring translation data, including human-only, machine-only, and hybrid approaches. Our findings demonstrate that human-machine collaboration can match or even exceed the quality of human-only translations, while being more cost-efficient. Error analysis reveals the complementary strengths between human and machine contributions, highlighting the effectiveness of collaborative methods. Cost analysis further demonstrates the economic benefits of human-machine collaboration methods, with some approaches achieving top-tier quality at around 60% of the cost of traditional methods. We release a publicly available dataset¹ containing nearly 18,000 segments of varying translation quality with corresponding human ratings to facilitate future research.

1 Introduction

Collecting high-quality translations efficiently presents significant challenges. Traditional approaches rely heavily on different tiers of human translators, ranging from professional linguists to junior bilingual speakers (Zouhar and Bojar, 2024). While these approaches can produce high-quality translations, they are often expensive, time-consuming, and challenging to scale for large datasets.

Recent advancements in machine translation with large language models (OpenAI, 2024; Gemini, 2024) have demonstrated models' impressive abilities to generate human-like translations. However, recent research (Yan et al., 2024) tends to

¹The dataset can be found at <https://github.com/google-research/google-research/tree/master/collaborative-tr-collection>.




Figure 1: Our 11 translation systems, organized by initial translation type (human or machine) and post-editing type (none, human, or machine). Detailed system descriptions are provided in Section 2.2.

position human translators and machine translation systems as competitors rather than potential collaborators, which could result in efficient alternatives for addressing the limitations of traditional translation data collection methods.

In this paper, we aim to fill the gap by comprehensively investigating the potential of human-machine collaboration to efficiently collect high-quality translation data. We hypothesize that combining the strengths of humans and machines could lead to higher quality, cost-efficient translation collection methods. To verify the hypothesis, we explore 11 different methods for acquiring translation data, including human-only, machine-only, and various hybrid methods.

Our research seeks to answer the following key questions:

- Can human-machine collaborative approaches produce translations of comparable or higher quality than traditional human-only or machine-only methods?
- How do different collaborative methods impact translation quality, and where do the improvements primarily originate?
- What are the cost implications of these various approaches, and can human-machine collabo-

ration offer a more cost-efficient solution for high-quality translation collection?

Our findings demonstrate that human-machine collaboration can match or even exceed human-only translation quality while being more cost-efficient. We present detailed error analyses to reveal the complementary strengths of the collaborative methods and conduct a thorough cost analysis to illustrate the economic benefits of collaborative approaches.

To support future research, we also release a publicly available dataset containing nearly 18,000 segments of varying translation quality with corresponding human ratings.

2 Collecting Translations

Translating text from one language to another can either be done by bilingual annotators or machine translation systems. However, both cases are prone to producing errors in their translations, including well-trained expert translators (Freitag et al., 2023). As such, translations can be post-edited, a process of correcting a translation, either manually or with a model, that often yields higher-quality translations.

Both steps of this process — the initial translation collection and the post-editing — can either be done with humans or with models, each with their own advantages and disadvantages in terms of speed, quality, cost, and scalability. In this work, we explore how combinations of human and machines for both steps of this pipeline can combine to produce high-quality translations.

2.1 Data Sources

We use the test sets provided by the WMT23 General MT Shared Task (Kocmi et al., 2023) and collect new translations using several methods. These data sets comprise 460 English-German (EnDe) paragraph-level segments and 1175 Chinese-English (ZhEn) sentence-level segments with human rating annotations.

2.2 Data Collection Systems

Figure 1 illustrates the combinations from the two dimensions: initial translation and post-editing methods from either human annotators or machines. Machines may be either large language models (LLMs) or machine translation (MT) systems. This results in the 11 systems in the figure, named ac-

cording to the source of the initial translation with a suffix representing the post-editing approach.

In this work, we use several different sources for the initial translation:

- **OrigHumanRef** and **HumanRef** are human translations collected by professional translators. We refer to the original reference provided by the WMT23 General MT Shared Task (Kocmi et al., 2023) as **ORIGHUMANREF**. We collected a new from-scratch professional translation **HUMANREF** following the standard annotation steps.
- **BestWMT** is the top-ranked MT system picked from the official results of WMT23 General Translation Task: GPT4-5shot (OpenAI, 2024) for EnDe and Lan-BridgeMT (Wu and Hu, 2023) for ZhEn, representing the state-of-the-art MT capability we can access.
- **MidWMT** is a middle-ranked MT system from the official results of WMT23 General Translation Task: ONLINE-G for both EnDe and ZhEn, representing the conventional MT quality we can use.

We additionally explore the following different methods for post-editing translations:

- **HumanPE** refers to the post-edit service provided by a separate batch of linguists. **HumanPEx2** means the translation going through two independent rounds of post-edits from professional translators. There is no translator overlap between the two batches.
- **LLMRefine** (Xu et al., 2024) is one of the state-of-the-art post-edit approaches leveraging error feedback for pin-pointing corrections. Here we reproduced its error-feedback process and leverage Gemini-1.0 Ultra (Gemini, 2024) with the reported prompts to generate post-edited text.

2.3 Evaluation

In this paper, we use Multidimensional Quality Metrics (MQM; Lommel et al., 2014; Freitag et al., 2021) to evaluate translation quality. MQM is the state-of-the-art human evaluation framework for MT. In MQM, expert raters identify error spans within translations, which are automatically converted to numeric scores. Lower scores indicate fewer errors and thus higher quality.




Figure 2: Cross-BLEU scores for different EnDe translation collection approaches.

3 Data Quality Analysis

In this section, we seek to understand how our collected translations differ from each other (§3.1), how those differences correspond to changes in quality (§3.2), and what those results indicate about the value of human-machine collaboration in terms of quality (§3.3).

3.1 Lexical Overlap and Similarity

Figure 2 presents a cross-BLEU (Freitag et al., 2022) similarity matrix for English-German translations, which measures lexical similarity between pairs of translations. See Figure 8 in Appendix A.1 for Chinese-English results. Higher scores indicate greater similarity.

One prominent pattern in these results is that systems based on the same initial translation retain high similarity even after post-editing. This indicates that **post-editing still preserves some characteristics of the original translation**. Also, translation systems based on MT (BESTWMT, MIDWMT, and their post-edited versions) are more similar to each other than to translations based on an initial human translation.

Table 1 presents a subset of the information in Figures 2 and 8, to emphasize the interaction between using human- vs. model-based approaches for the initial translation and post-edit. This illustrates the trend that **humans and machines tend to make more changes to translations from the other group**.

Source	+ HUMANPE	+ LLMREFINE
EN-DE		
HUMANREF	95	72
BESTWMT	81	83
MIDWMT	70	77
ZH-EN		
HUMANREF	88	79
BESTWMT	84	89
MIDWMT	68	71

Table 1: Cross-BLEU score comparison between different post-edited versions of the same translation. Lower numbers indicate less similarity and more changes from the initial translation.

Reference	MQM per segment
HUMANREF+LLMREFINE	2.76
ORIGHUMANREF	2.80
BESTWMT+HUMANPE	2.98
BESTWMT	3.30
BESTWMT+LLMREFINE	3.36
HUMANREF+HUMANPE	3.53
HUMANREF+HUMANPEX2	3.70
HUMANREF	3.79
MIDWMT+HUMANPE	3.83
MIDWMT+LLMREFINE	4.02
MIDWMT	6.45

Table 2: English-German MQM human evaluation results. Lower scores represent higher translation quality.

3.2 MQM Quality Evaluation

Tables 2 and 3 present the MQM human evaluation results. The solid lines denote *significance clusters*, where every system in a cluster is statistically significantly better than every system below that cluster, based on random permutation tests with 10,000 trials, where a p -value of less than $\alpha = 0.05$ is considered significant.

The results reveal that **HumanRef+LLMRefine** and **BestWMT+HumanPE** are the overall winners, with each appearing in the best significance cluster in both language pairs.

Tables 2 and 3 show that post-edits, both HUMANPE and LLMREFINE, **demonstrate a positive impact on initial translations**. These methods consistently either elevate the translation quality to a higher level of significance or preserve the existing quality.

Reference	MQM per segment
HUMANREF+LLMREFINE	1.82
HUMANREF+HUMANPEX2	1.82
BESTWMT+HUMANPE	1.87
HUMANREF+HUMANPE	1.91
BESTWMT+LLMREFINE	1.94
HUMANREF	2.05
BESTWMT	2.22
MIDWMT+HUMANPE	2.23
MIDWMT+LLMREFINE	2.45
MIDWMT	3.98
ORIGHUMANREF	5.63

Table 3: Chinese-English MQM human evaluation results. Lower scores represent higher translation quality.

3.3 Human-Machine Collaboration

Table 4 and Figure 3 present a detailed analysis that highlights the quality benefits of human-machine collaboration.

Figure 3 shows the gains in quality that can be provided by our post-edit approaches. The gains are more pronounced when starting with a lower-quality translation (MIDMT), but even high-quality translations (HUMANREF, BESTWMT) can be improved. The quality differences between initial translations are greatly reduced after post-editing, but not eliminated.

Table 4 provides a finer-grain analysis of the effect of each post-edit approach on different initial translations. In both language pairs, an LLM-based method provides the most benefit when starting with human translation, while human post-editing provides the most benefit for machine translation. Recall that in Table 1 we showed that model-based methods and humans make more changes to translations from the other group; here we see that these changes are also net-positive. This indicates that **human-machine collaboration is an effective way to achieve high-quality translations.**

4 Error Analysis

Here we present more detailed error analysis of both initial translation and post-edit stages to investigate where the quality improvements originate. Using English-German as an example, we first present analysis of the initial translations in Section 4.1 to understand the initial error distributions. Then, we further investigate the error-correction dynamics during post-editing in Section 4.2 to un-

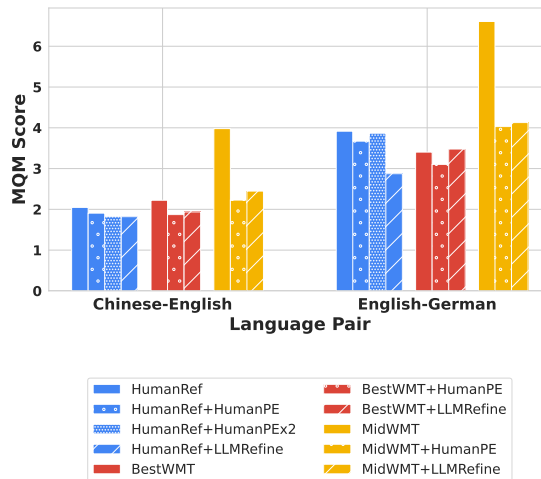


Figure 3: MQM Scores for different translation systems across two language pairs: Chinese-English and English-German. Bars represents the average MQM scores for each translation system. The systems are grouped and colored by initial translation and further categorized by post-editing method with different fill patterns. Lower MQM scores indicate better quality.

derstand why human-machine collaboration stands out from other approaches.

4.1 Error distribution from Initial Translation

We present the error type and severity distribution of the English-German initial translations in Table 5 and that of Chinese-English in Appendix A.2. It shows that accuracy-related errors are the primary source of major errors and that the distribution of minor errors is more evenly spread across different categories. Importantly, these error distributions are similar for both human and machine-based initial translations. This consistency provides a solid foundation for comparing different post-editing techniques in our downstream analysis.

4.2 Error Correction from Post-Editing

To understand the error-correction dynamics of different post-editing methods and find the origin of the improvements of human-machine collaboration, we explored three key questions:

- Do different post-editing methods agree on which segments to modify?
- How do different post-editing methods affect the total number of major and minor errors?
- How do different post-editing methods affect

Source	Init. Translation	+ HumanPE		+ HumanPEx2		+ LLMRefine		
		Score ↓	Δ	Score ↓	Δ	Score ↓	Δ	
EN-DE								
HUMANREF	3.79	3.53	-0.26	3.70	-0.09	2.76	-1.03	
BESTWMT	3.30	2.98	-0.32	-	-	3.36	+0.06	
MIDWMT	6.45	3.83	-2.62	-	-	4.02	-2.43	
ZH-EN								
HUMANREF	2.05	1.91	-0.14	1.82	-0.23	1.82	-0.23	
BESTWMT	2.22	1.87	-0.35	-	-	1.94	-0.28	
MIDWMT	3.98	2.23	-1.75	-	-	2.45	-1.53	

Table 4: MQM human evaluation comparison of each post-edit approach on different initial translations. Lower MQM scores indicates better quality.

Error Type	OrigHumanRef	HumanRef	BestWMT	MidWMT
No-error	200	175	144	116
Major	186	263	211	470
Fluency	27 (15%)	29 (11%)	25 (12%)	88 (19%)
Accuracy	114 (61%)	149 (57%)	108 (51%)	275 (59%)
Style	22 (12%)	54 (21%)	45 (21%)	60 (13%)
Minor	659	745	891	1084
Fluency	257 (39%)	242 (32%)	439 (49%)	538 (50%)
Accuracy	181 (27%)	211 (28%)	186 (21%)	226 (21%)
Style	141 (21%)	190 (26%)	178 (20%)	213 (20%)

Table 5: Error type and severity distributions of English-German MQM human evaluation results.

the total number of high- and low-quality segments?

We first examine how often human post-editors (HUMANPE) and machine post-editing methods (LLMREFINE) agree on which segments need correction. Figure 4 shows that both methods identify more segments for editing in lower-quality initial translations as evidenced by the shrinking "No Change" (yellow) section. Notably, agreement between HUMANPE and LLMREFINE increased from 23.9% for high-quality HUMANREF translations to 67.4% for lower-quality MidWMT translations as observed by the expanded "HumanPE & LLMRefine" (purple) section. This suggests more consensus on obvious errors in lower-quality texts while greater divergence for higher-quality translations in editing approaches. A detailed numerical breakdown is shown in Table 8 in Appendix A.3.

Another interesting observation from Figure 4 is that HUMANPE (purple + red) identifies a larger proportion of segments needing correction in BESTWMT (48.9% + 19.3% = 68.2%) compared to HUMANREF (23.9% + 11.5% = 35.4%), despite

²Tabular statistics are provided in Tables 8 and 9 in Appendix A.3.




Figure 4: Agreement between HumanPE and LLM-Refine in identifying segments requiring post-edit on English-German data. Each pie chart² represents a different initial translation source.

the superior quality of BESTWMT over HUMANREF as shown in Table 2. This suggests that human post-editors might overlook certain errors in human translations due to their familiar patterns. Conversely, the unfamiliar patterns in machine-generated text may make errors more salient to human editors. This interpretation is consistent with the pattern depicted in Figure 2, where human post-editors make fewer changes to human translations than to machine translations. This observation provides one plausible explanation for the necessity of human-machine collaboration in achieving high-quality translations. Detailed statistics for English-German are shown in Table 8 and similar trends are also observed in Chinese-English in Table 9 in Appendix A.3.

We wish to investigate how many errors are corrected during post-editing, but because it is difficult to automatically determine whether an individual error was corrected, we instead examine how the total number of errors changes after post-editing, also considering severity. Figure 5 shows that both human and machine post-editing reduce overall error counts across different initial translation qual-





Figure 5: Error changes percentages by different post-editing approaches on English-German data. The percentages present the changes in error counts for each post-editing method compared to its initial translation. A negative indicates a decrease in errors, while positive value indicates an increase in the error type.


ities. LLMREFINE outperforms HUMANPE on HUMANREF initial translation in reducing major errors (-27.8% vs. -2.7%), while HUMANPE is superior for BESTWMT, decreasing major-error segments compared to the increase for LLMREFINE (-4.3% vs. +8.5%). On MIDWMT, both methods show substantial improvements, with HUMANPE moderately ahead of LLMREFINE (-41.3% vs. -37.0% decrease in major-error segments). These findings highlight the complementary strengths of human and machine post-editing methods, indicating that a hybrid method is likely the most effective strategy for reducing errors, regardless of the initial translation’s origin. Similar trends are also observed in Chinese-English in Figure 9 in Appendix A.3.

To understand the error correction dynamics for each segment, we analyzed how MQM scores change before and after post-editing. Ideally, post-editing would fix existing errors while minimizing the introduction of new ones. However, as Figure 6 illustrates, post-editing is not guaranteed to improve every segment: while some segments are improved, others are worsened.

Figure 6 compares HUMANPE and LLMREFINE on HUMANREF initial translations for English-German data. Both methods reduce the number of high-scoring (low-quality) segments ($MQM \geq 5$). Notably, LLMREFINE outperforms HUMANPE by showing fewer quality-degrading corrections and more quality-improving ones. LLMREFINE minimizes low-to-high-scoring



(a) From HUMANREF to HUMANREF+HUMANPE




(b) From HUMANREF to HUMANREF+LLMREFINE


Figure 6: Segment-level quality shift through HUMANPE and LLMREFINE from English-German HUMANREF. Each segment is categorized into one of three groups based on its MQM score: 1) high-scoring segments with $MQM \geq 5$; 2) low-scoring segments with $0 < MQM < 5$; 3) error-free segments with $MQM=0$. Higher MQM scores indicate more numerous/severe errors and accordingly lower translation quality.

transitions with a narrower flow from low-scoring segments ($0 < MQM < 5$) to high-scoring segments compared to HUMANPE. Moreover, LLMREFINE achieves a significant reduction in high-scoring segments by 6.5% (from 29.1% to 22.6%) compared to HUMANPE’s 1.3% (from 29.1% to 27.8%), suggesting that it is more effective at achieving post-editing gains while preserving originally good translations. A similar trend is observed for Chinese-English with HUMANREF initial translation in Figure 10 in Appendix A.3.

To demonstrate that the quality improvement is not solely due to the capabilities of LLMREFINE, we conducted further experiments as shown in Figure 7. This figure compares HUMANPE and LLMREFINE with BESTWMT initial translations. Interestingly, the results are reversed: HUMANPE



(a) From BESTWMT to BESTWMT+HUMANPE



(b) From BESTWMT to BESTWMT+LLMREFINE

Figure 7: Segment-level quality shift through HUMANPE and LLMREFINE from English-German BESTWMT.

outperforms LLMREFINE in this scenario, showing fewer quality-degrading corrections and more quality-improving ones. HUMANPE demonstrates a significantly wider flow from high-scoring segments to low-scoring ones. It achieves a notable reduction in high-scoring segments by 3.2% (from 26.5% to 23.3%), while LLMREFINE sees an increase of 0.7% (from 26.5% to 27.2%). Furthermore, HUMANPE significantly increases the No-Error segments (MQM=0) by 12.2% (from 33.5% to 45.7%) compared to LLMREFINE’s 2.2% increase.

The distinct performance differences of LLMREFINE and HUMANPE in these two sets of experiments highlight that **the quality improvement stems primarily from the complementary strengths of human and machine collaboration**, rather than the superior capability of either LLMREFINE or HUMANPE alone. This underscores the importance of leveraging both human and machine strengths in achieving optimal translation quality.

Systems	Quality Rank		Costs
	EnDe	ZhEn	
HUMANREF	3	2	1X
HUMANREF+HUMANPE	2	1	1.6X
HUMANREF+HUMANPEX2	2	1	2.2X
HUMANREF+LLMREFINE	1	1	1X
BESTWMT+HUMANPE	1	1	0.6X

Table 6: Quality rank and costs comparison of different data collection systems. 1st rank indicates the translation quality belongs to the highest quality significance cluster in Table 2 and 3.

5 Costs Analysis

So far we have focused on comparing quality between various translation data collection approaches. However, practical considerations make it important to consider the trade-off between quality and costs. Table 6 analyzes relative human annotation costs between various approaches, along with the rank of the significance cluster that each method appeared in. The exact costs for the human annotation conducted in this study are confidential (although all annotators were paid fair market wages), so we instead use relative costs, based on the industry standard that post-editing text of a given length takes less time (and accordingly costs less) than producing a translation of that length. We specifically assume that human post-editing costs around 60% of what human translation does. According to existing literature (Plitt and Masselot, 2010; Zouhar et al., 2021; Green et al., 2013) and internal statistics, we believe it’s a fair assumption, although the exact costs can vary upon different vendors, languages, task size, etc.

With this framework, the best combination of quality and cost appears to be human post-editing of high-quality MT (BESTWMT+HUMANPE), attaining quality in the top significance cluster in both language pairs with only 60% of the human annotation cost of collecting an initial human translation. Meanwhile, we see that one or two rounds of human post-editing of an initial human translation increases costs without a meaningful gain in quality, while just applying an LLM post-editor (HUMANREF+LLMREFINE) brings quality to the top significant cluster with no additional human annotation cost, making it a viable option when human translations are already collected. It’s worth noting that LLM inference costs are negligible (on the order of dollars per million tokens) compared to human annotation costs, further enhancing the

cost-effectiveness of LLM-based approaches. This indicates that **human-machine collaboration can be a faster, more cost-efficient alternative** to traditional collection of translations from humans, optimizing both quality and resource allocation by leveraging the strengths of both humans and machines.

6 Related Work

There have been a few studies investigating methods of acquiring high-quality translations. Recently, [Zouhar et al. \(2024\)](#) proposed collecting high-quality translations by building consensus between multiple translators. [Zouhar and Bojar \(2024\)](#) proposed collecting multiple translations from different tiers of human translators with careful budget calculations to optimize cost-efficiency.

Human Post-Edits Computer-aided translation tools are now widely used by professional translators for interactive translation and post-editing ([Alabau et al., 2014](#); [Federico et al., 2014](#); [Green et al., 2014](#); [Denkowski, 2015](#); [Sin-wai, 2014](#); [Kenny, 2012](#)). [Carl et al. \(2011\)](#) have shown that human translators work faster and make fewer mistakes when editing machine translations than when translating from scratch. [Toral et al. \(2018\)](#) supports this, demonstrating even greater improvements with neural machine translation compared to phrase-based systems. [Zouhar et al. \(2021\)](#) investigates the relationship between machine translation quality and post-editing efforts and found no straightforward relationship. On the other hand, [Popovic et al. \(2016\)](#) suggested that post-edits should be used carefully for MT evaluation due to the bias of each post-edit towards its MT system. Further, [Toral \(2019\)](#) showed that human post-edits are simpler and more normalised in language than human translations from scratch.

Automatic Refinement [Lin et al. \(2022\)](#) showed how the errors that humans make differ from those made by MT systems. They constructed a Translation Error Correction (TEC) corpus with professional translators and showed that models trained on it outperform Automatic Post-Editing (APE) models ([Knight and Chander, 1994](#)) that are trained to correct MT output. Since the emergence of LLMs, new refinement approaches based on detailed MQM annotations have appeared ([Xu et al., 2023](#); [Fernandes et al., 2023](#)). [Xu et al. \(2024\)](#) showed that these refinement method can be used

to improve the quality of human translations.

Meanwhile, machines have been extensively evaluated and utilized as an alternative to human annotators for data collection ([Zouhar et al., 2021](#); [Yan et al., 2024](#)).

In contrast to the above methods, we investigate the interaction between humans and machines in the initial translation and post-editing stages, including detailed analysis of the resulting changes in quality while also considering cost-efficiency.

7 Conclusion

We investigate various approaches for gathering translation data, including human-only, machine-only, and hybrid approaches. Our results demonstrate that human-machine collaboration can consistently generate high-quality translations at a lower cost than human-only methods. Through detailed error analysis, we uncovered the nuances of error correction dynamics and highlighted the advantages of human-machine collaborative methods. Our cost analysis also demonstrates the cost-efficiency of human-machine collaboration methods. Finally, we release to the public a dataset of roughly 18,000 translation segments of varying quality from different collection methods along with human ratings, to facilitate further research in this area.

Limitations

This study focuses on two language pairs, English-German and Chinese-English. They are chosen due to the extensive study in the WMT23 metrics shared task (Freitag et al., 2023) and the availability of data from various translation systems from the WMT23 general shared task (Kocmi et al., 2023). While our analysis provides support for the findings presented in this work and we offer a plausible explanation for the observed results, it is important to acknowledge certain variables are not accounted for in this work, including using translators or post-editors with varying quality levels, different systems for translation and post-editing, utilizing sentence or paragraph datasets from other domains, and higher or lower resource language pairs beyond the two investigated here. Therefore, we cannot guarantee the observed trends will generalize to different datasets.

We want to especially highlight the need for further exploration of the quality variance observed among human translators, such as ORIGHUMANREF and HUMANREF in the English-German translation task. The current study’s limited annotation budget and timeline restricted the depth of this investigation. Future experiments aimed at examining the impact of post-editing on annotator agreement would be particularly interesting and valuable.

Ethical Statement

The source data used for translation and post-edits is accessible to the public. We’re certain that the data annotated by human labors is free from risk or toxic content. We used an internal, proprietary tool to collect human translation, post-edits, and evaluation data. The annotators were compensated fairly and were not required to disclose any personal details during the annotation process. All the test data used in this study are publicly available and annotators were allowed to label sensitive information if necessary. The annotators are fully informed that the data they collected will be used for research purposes.

References

Vicent Alabau, Christian Buck, Michael Carl, Francisco Casacuberta, Mercedes García-Martínez, Ulrich Germann, Jesús González-Rubio, Robin L Hill, Philipp Koehn, Luis A Leiva, et al. 2014. Casmacat:

A computer-assisted translation workbench. In *Proceedings of the Demonstrations at the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 25–28.

Michael Carl, Barbara Dragsted, Jakob Elming, Daniel Hardt, and Arnt Lykke Jakobsen. 2011. The process of post-editing: A pilot study. *Copenhagen Studies in Language*, 41(1):131–142.

Michael Denkowski. 2015. Machine translation for human translators.

Marcello Federico, Nicola Bertoldi, Mauro Cettolo, Matteo Negri, Marco Turchi, Marco Trombetti, Alessandro Cattelan, Antonio Farina, Domenico Lupinetti, Andrea Martines, et al. 2014. The matecat tool. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: System Demonstrations*, pages 129–132.

Patrick Fernandes, Daniel Deutsch, Mara Finkelstein, Parker Riley, André Martins, Graham Neubig, Ankush Garg, Jonathan Clark, Markus Freitag, and Orhan Firat. 2023. [The devil is in the errors: Leveraging large language models for fine-grained machine translation evaluation](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 1066–1083, Singapore. Association for Computational Linguistics.

Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, errors, and context: A large-scale study of human evaluation for machine translation](#). *Transactions of the Association for Computational Linguistics*, 9:1460–1474.

Markus Freitag, David Grangier, Qijun Tan, and Bowen Liang. 2022. [High quality rather than high model probability: Minimum Bayes risk decoding with neural metrics](#). *Transactions of the Association for Computational Linguistics*, 10:811–825.

Markus Freitag, Nitika Mathur, Chi-kiu Lo, Eleftherios Avramidis, Ricardo Rei, Brian Thompson, Tom Kocmi, Frederic Blain, Daniel Deutsch, Craig Stewart, Chrysoula Zerva, Sheila Castilho, Alon Lavie, and George Foster. 2023. [Results of WMT23 metrics shared task: Metrics might be guilty but references are not innocent](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 578–628, Singapore. Association for Computational Linguistics.

Team Gemini. 2024. [Gemini: A family of highly capable multimodal models](#).

Spence Green, Jeffrey Heer, and Christopher D. Manning. 2013. [The efficacy of human post-editing for language translation](#). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’13, page 439–448, New York, NY, USA. Association for Computing Machinery.

- Spence Green, Sida I. Wang, Jason Chuang, Jeffrey Heer, Sebastian Schuster, and Christopher D. Manning. 2014. [Human effort and machine learnability in computer aided translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1225–1236, Doha, Qatar. Association for Computational Linguistics.
- Dorothy Kenny. 2012. *Electronic Tools and Resources for Translators*.
- Kevin Knight and Ishwar Chander. 1994. Automated postediting of documents. In *AAAI*, volume 94, pages 779–784.
- Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Philipp Koehn, Benjamin Marie, Christof Monz, Makoto Morishita, Kenton Murray, Makoto Nagata, Toshiaki Nakazawa, Martin Popel, Maja Popović, and Mariya Shmatova. 2023. [Findings of the 2023 conference on machine translation \(WMT23\): LLMs are here but not quite there yet](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 1–42, Singapore. Association for Computational Linguistics.
- Jessy Lin, Geza Kovacs, Aditya Shastry, Joern Wuebker, and John DeNero. 2022. [Automatic correction of human translations](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 494–507, Seattle, United States. Association for Computational Linguistics.
- Arle Lommel, Hans Uszkoreit, and Aljoscha Burchardt. 2014. Mqm: Un marc per declarar i descriure mètriques de qualitat de la traducció. *Tradumàtica: traducció i tecnologies de la informació i la comunicació*, (12):455–463.
- Team OpenAI. 2024. [Gpt-4 technical report](#).
- Mirko Plitt and François Masselot. 2010. [A productivity test of statistical machine translation post-editing in a typical localisation context](#). In *Prague Bulletin of Mathematical Linguistics*.
- Maja Popovic, Mihael Arčan, and Arle Lommel. 2016. [Potential and limits of using post-edits as reference translations for MT evaluation](#). In *Proceedings of the 19th Annual Conference of the European Association for Machine Translation*, pages 218–229.
- Chan Sin-wai. 2014. The development of translation technology 1967–2013. In *Routledge Encyclopedia of Translation Technology*, pages 2–31. Routledge.
- Antonio Toral. 2019. [Post-editeese: an exacerbated translationese](#). In *Proceedings of Machine Translation Summit XVII: Research Track*, pages 273–281, Dublin, Ireland. European Association for Machine Translation.
- Antonio Toral, Martijn Wieling, and Andy Way. 2018. [Post-editing effort of a novel with statistical and neural machine translation](#). *Frontiers in Digital Humanities*, 5:9.
- Yangjian Wu and Gang Hu. 2023. [Exploring prompt engineering with GPT language models for document-level machine translation: Insights and findings](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 166–169, Singapore. Association for Computational Linguistics.
- Wenda Xu, Daniel Deutsch, Mara Finkelstein, Juraj Juraska, Biao Zhang, Zhongtao Liu, William Yang Wang, Lei Li, and Markus Freitag. 2024. [Llmrefine: Pinpointing and refining large language models via fine-grained actionable feedback](#).
- Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Wang, and Lei Li. 2023. [INSTRUCTSCORE: Towards explainable text generation evaluation with automatic feedback](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5967–5994, Singapore. Association for Computational Linguistics.
- Jianhao Yan, Pingchuan Yan, Yulong Chen, Judy Li, Xianchao Zhu, and Yue Zhang. 2024. [Gpt-4 vs. human translators: A comprehensive evaluation of translation quality across languages, domains, and expertise levels](#).
- Vilém Zouhar, Martin Popel, Ondřej Bojar, and Aleš Tamchyna. 2021. [Neural machine translation quality and post-editing performance](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10204–10214, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Vilém Zouhar and Ondřej Bojar. 2024. [Quality and quantity of machine translation references for automatic metrics](#).
- Vilém Zouhar, Věra Kloudová, Martin Popel, and Ondřej Bojar. 2024. [Evaluating optimal reference translations](#).

A Appendix

A.1 Cross-BLUE scores

Figure 8 presents the cross-BLEU similarity matrix for Chinese-English translation systems.

A.2 Error Distribution of Initial Translation

Table 7 presents the error type and severity distributions of Chinese-English MQM human evaluation results.




Figure 8: Cross-BLEU scores for different Chinese-English translation collection approaches

Error Type	OrigHumanRef	HumanRef	BestWMT	MidWMT
No-error	225	453	490	304
Major				
Total	955	284	260	648
Fluency	18 (2%)	15 (5%)	5 (2%)	23 (4%)
Accuracy	851 (89%)	229 (81%)	221 (85%)	568 (88%)
Style	31 (3%)	12 (4%)	5 (2%)	18 (3%)
Minor				
Total	1902	1303	1238	1704
Fluency	625 (33%)	431 (33%)	364 (29%)	522 (31%)
Accuracy	599 (31%)	412 (32%)	388 (31%)	447 (26%)
Style	629 (33%)	402 (31%)	436 (35%)	677 (40%)

Table 7: Error type and severity distributions of Chinese-English MQM human evaluation results.

A.3 Error Correction from Post-Editing

Figures 9, 10, and 11 present Chinese-English results comparable to the English-German results presented in Section 4.2. Table 8 presents the same data as in Figure 4 for English-German, and Table 9 presents the same for Chinese-English.






Figure 9: Error changes percentages by different post-editing approaches on Chinese-English data. The percentages present the changes in error counts for each post-editing method compared to its initial translation. A negative indicates a decrease in errors, while positive value indicates an increase in the error type.



(a) From HUMANREF to HUMANREF+HUMANPE



(b) From HUMANREF to HUMANREF+LLMREFINE


Figure 10: Segment-level quality shift through HUMANPE and LLMREFINE from Chinese-English HUMANREF.

Initial Translation	Total Seg	HumanPE	LLMRefine	HumanPE & LLMRefine	Human Only	LLMRefine Only
HumanRef	460	163 (35.4%)	235 (51.1%)	110 (23.9%)	53 (11.5%)	125 (27.2%)
BestWMT	460	314(68.3%)	281 (61.1%)	225 (48.9%)	89 (19.3%)	56 (12.2%)
MidWMT	460	400 (87.0%)	324 (70.4%)	310 (67.4%)	90 (19.6%)	14 (3%)


Table 8: Numerical breakdown of the agreement between HUMANPE and LLMREFINE in identifying segments requiring post-editing in English-German

Initial Translation	Total Seg	HumanPE	LLMRefine	HumanPE & LLMRefine	Human Only	LLMRefine Only
HumanRef	1175	558 (47.5%)	225 (19.1%)	161 (13.7%)	397 (33.8%)	64 (5.4%)
BestWMT	1175	830 (70.6%)	133 (11.3%)	123 (10.5%)	707 (60.2%)	10 (0.9%)
MidWMT	1175	1006 (85.6%)	408 (34.7%)	399 (34.0%)	607 (51.7%)	9 (0.8%)

Table 9: Numerical breakdown of the agreement between HumanPE and LLMRefine in identifying segments requiring post-editing in Chinese-English



(a) From BESTWMT to BESTWMT+HUMANPE



(b) From BESTWMT to BESTWMT+LLMREFINE

Figure 11: Segment-level quality shift through HUMANPE and LLMREFINE from Chinese-English BESTWMT.