

Hindsight Quality Prediction Experiments in Multi-Candidate Human-Post-Edited Machine Translation

Malik Marmonier Benoît Sagot Rachel Bawden

Inria, Paris Center

{firstname.lastname}@inria.fr

Abstract

This paper investigates two complementary paradigms for predicting machine translation (MT) quality: source-side difficulty prediction and candidate-side quality estimation (QE). The rapid adoption of Large Language Models (LLMs) into MT workflows is reshaping the research landscape, yet its impact on established quality prediction paradigms remains underexplored. We study this issue through a series of “hindsight” experiments on a unique, multi-candidate dataset resulting from a genuine MT post-editing (MTPE) project. The dataset consists of over 6,000 English source segments with nine translation hypotheses from a diverse set of traditional neural MT systems and advanced LLMs, all evaluated against a single, final human post-edited reference. Using Kendall’s rank correlation, we assess the predictive power of source-side difficulty metrics, candidate-side QE models and position heuristics against two gold-standard scores: TER (as a proxy for post-editing effort) and COMET (as a proxy for human judgment). Our analysis yields three primary findings: (1) On the source side, the predictive power of difficulty metrics is highly contingent on the reference metric used; features that strongly correlate with COMET (e.g., segment length, neural predictors) show much weaker correlation to TER. (2) On the candidate side, we find a significant mismatch between QE model rankings and final human-adjudicated quality, and further show that modern QE metrics are significantly more aligned with the quality of traditional neural MT outputs than with those from general-purpose LLMs. (3) While we confirm a statistically significant positional bias in document-level LLMs (i.e., the tendency for translation quality to degrade for segments occurring later in a document) its practical impact on translation quality appears to be negligible. These findings highlight that the architectural shift towards LLMs alters the reliability of established quality prediction methods while simultaneously mitigating previous challenges in document-level translation.

Keywords: machine translation, translation difficulty prediction, quality estimation, position bias

1. Introduction

The ability to anticipate machine translation (MT) quality is valuable to both researchers and professional translators. It is typically addressed through two complementary paradigms, which we group under the umbrella term “quality prediction.”

On the source side, translation difficulty prediction aims to estimate translation effort a priori by analyzing the source text. This can inform decisions on resource allocation and help avoid wasted effort on segments unlikely to meet minimal quality requirements for human post-editing (Fericola et al., 2023). Conversely, this can also drive the selection of more challenging source segments for the creation of evaluation sets in MT research (Proietti et al., 2025). On the candidate side, Quality Estimation (QE; cf. Callison-Burch et al. 2012) assesses the quality of a machine-generated translation without a target reference, enabling downstream tasks such as triaging translations by degrees of post-editing required, or selecting the best hypothesis among multiple candidates.

These approaches were largely developed and benchmarked within an MT landscape that has been transformed in complex ways due to the rapid rise of large language models (LLMs) and their widespread adoption in translation use cases (Chat-

terji et al., 2025). LLMs, thanks to their general-purpose nature, have spurred massive industrial investment far exceeding that of specialized MT systems. A notable result of this scaling is a dramatic expansion of model capabilities with regard to context-window length, making document-level translation widely accessible for the first time. Yet, the practical consequences of these evolutions for the established methods of quality prediction remain underexplored.

This paper investigates this issue through a series of “hindsight” experiments conducted on a unique dataset with substantial ecological validity, as it is the byproduct of a genuine, multi-candidate post-editing project. For each of over 6,000 source segments, this dataset contains nine translation hypotheses from a diverse set of systems—including traditional NMT models and advanced LLMs—operating at various “granularities” (at the sentence or the document level). Because the gold-standard reference is the result of post-editing, the TER metric (Olive, 2005; Snover et al., 2006) serves as a meaningful proxy for human post-editing effort, complementing and contrasting more advanced neural “black box” metrics like COMET (Rei et al., 2020), which was trained via regression to predict human quality judgment scores. We operationally define the predictive power of any given quality met-

ric or heuristic as its Kendall’s τ rank correlation with these two reference scores.

Our analysis yields three primary findings:

- On the source side, the predictive power of difficulty metrics is highly contingent on the ground-truth metric used; features that strongly correlate with COMET (e.g., segment length, neural predictors) show almost no correlation with TER.
- On the candidate side, modern QE metrics are significantly more aligned with, and predictive of, the quality of traditional neural MT outputs than those from general-purpose LLMs.
- While we confirm a statistically significant positional bias in document-level LLMs (i.e., the tendency for translation quality to degrade for segments occurring later in a document; Peng et al. 2024, 2025) its practical impact on translation quality appears to be negligible.

These findings highlight that the architectural shift towards LLMs alters the reliability of established quality prediction methods while simultaneously mitigating previous challenges in document-level translation. To facilitate further research, we make our dataset and code publicly available under CC-BY SA 4.0 license.¹

2. Data

The experiments in this paper are conducted on a unique dataset that is the byproduct of the creation of the French partition of the OLDI Seed Corpus (Marmonier et al., 2025). We first describe the source corpus and the process used to create the French reference translations, before detailing the resulting multi-candidate dataset used for our analyses.

2.1. The OLDI Seed Corpus and its French Partition

The OLDI Seed Corpus (originally the NLLB Seed Corpus) was created to provide a small but high-quality parallel dataset to bootstrap MT capabilities for low-resource languages (NLLB Team et al., 2022; Maillard et al., 2023). The source material consists of approximately 6,000 English sentences sampled from a curated list of 203 core Wikipedia articles, ensuring broad topic coverage and allowing these segments to be reassembled into documents.

As a contribution to the WMT 2025 Open Language Data Initiative (OLDI) shared task (Dale et al., 2025), a French partition of this corpus was created

¹<https://github.com/mmarmonier/Hindex>

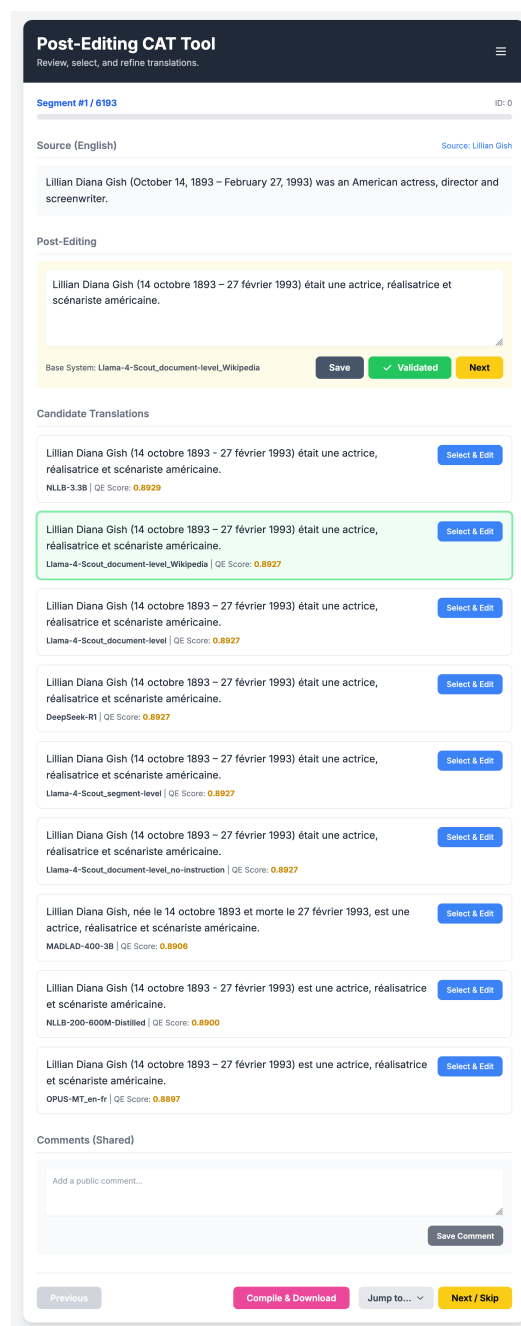


Figure 1: Post-editing interface used to compile the French partition of the OLDI Seed Corpus. Image borrowed from (Marmonier et al., 2025).

(Marmonier et al., 2025). The primary goal was to establish a high-quality pivot resource to facilitate the future creation of parallel corpora for the under-resourced regional languages of France whose speakers are far more likely to be fluent in French than in English. Given the high quality of modern English-French MT systems, the French partition was produced using a post-edited MT (MTPE) workflow. This process was carried out by two native French speakers with C2-level proficiency in English, using a custom-built interface (see Figure 1).

2.2. The Multi-Candidate Byproduct Dataset

Our current work makes use of the full set of data generated during the creation of the aforementioned French partition. This “byproduct” dataset is uniquely suited for our experiments, as for each of the 6,193 English source segments, it contains:

- A set of nine different MT hypotheses generated by a diverse array of systems and prompting techniques.
- A single, final human post-edited reference translation, which represents the gold-standard output after meticulous review and correction.

The nine translation hypotheses were generated using a mix of traditional NMT models and prompted LLMs, which can be grouped into two main categories:

- NMT models (*sent2sent*): Four dedicated sequence-to-sequence Transformer models that translate sentence by sentence. **OPUS-MT**: A standard bilingual Transformer model trained on open parallel corpora (Tiedemann and Thottingal, 2020). **NLLB-3.3B**: The 3.3B parameter multilingual model from the NLLB family (NLLB Team et al., 2022). **NLLB-600M-Distilled**: The smaller, 600M parameter distilled version of NLLB. **MADLAD-400-3B**: The 3B parameter multilingual model from the MADLAD-400 project (Kudugunta et al., 2023).
- LLMs (*smallDoc2sent* or *doc2doc*): Five hypotheses were generated using two LLMs with various prompting strategies. **Llama-4-Scout_segment-level**: A 109B-parameter Llama 4 Scout model (Meta AI, 2025) translated segments individually, where the prompt included up to five preceding source sentences for local context (*smallDoc2sent*). **Llama-4-Scout_document-level**: The same 109B model translated entire documents in a single pass using a full set of translation guidelines. **Llama-4-Scout_document-level_no-instruction**: A contrastive ablation of the above, where the translation guidelines were removed from the prompt. **Llama-4-Scout_document-level_Wikipedia**: The document-level setup, but with the prompt augmented by the full text of the corresponding French Wikipedia article. **DeepSeek-R1**: One document-level hypothesis was generated using the 671B-parameter model DeepSeek-R1 (DeepSeek-AI et al., 2025).

To establish a baseline for the quality of the raw machine-generated hypotheses, we calculated

System	Mean	95% CI	Group
DeepSeek-R1	12.44	[12.03, 12.86]	A
Llama-4-Scout_segment-level	26.09	[25.63, 26.55]	B
MADLAD-400-3B	26.66	[26.19, 27.13]	B
Llama-4-Scout_document-level_no-instruction	26.68	[26.22, 27.15]	B
Llama-4-Scout_document-level	26.73	[26.26, 27.19]	B
Llama-4-Scout_document-level_Wikipedia	27.27	[26.75, 27.78]	B
NLLB-3.3B	27.73	[27.26, 28.21]	B
OPUS-MT_en-fr	31.80	[31.28, 32.32]	C
NLLB-200-600M-Distilled	32.75	[30.92, 34.57]	C

Table 1: **TER** reference scores with 95% confidence intervals (↓ lower is better).

System	Mean	95% CI	Group
DeepSeek-R1	0.9219	[0.9206, 0.9232]	A
Llama-4-Scout_segment-level	0.8982	[0.8968, 0.8997]	B
MADLAD-400-3B	0.8956	[0.8940, 0.8971]	B
Llama-4-Scout_document-level_no-instruction	0.8954	[0.8938, 0.8969]	B
Llama-4-Scout_document-level	0.8949	[0.8933, 0.8964]	B
NLLB-3.3B	0.8934	[0.8918, 0.8949]	B
Llama-4-Scout_document-level_Wikipedia	0.8917	[0.8899, 0.8936]	B
NLLB-200-600M-Distilled	0.8779	[0.8761, 0.8798]	C
OPUS-MT_en-fr	0.8730	[0.8709, 0.8750]	D

Table 2: **COMET** reference scores with 95% confidence intervals (↑ higher is better).

reference-based COMET² and TER³ scores for each of the nine systems against the final human post-edited text.⁴ During data collection, the DeepSeek-R1 model refused to translate a small number of documents (165 segments) pertaining to sensitive topics, due to its safety filters. To ensure a fair comparison, these segments were excluded from all evaluations, leaving a total of 6,028 segments.

Tables 1 and 2 present the average performance of each system. The results show a clear hierarchy of quality. DeepSeek-R1 stands out as the top-performing system, achieving significantly higher COMET and lower TER scores than all other models. The remaining systems fall into several distinct statistical groups, with the larger NMT models and all Llama 4 prompting variants forming a competitive middle tier, and the smaller OPUS-MT and distilled NLLB models occupying the lower performance tiers.

3. Source-Side Experiments

Our first set of experiments investigates source-side quality prediction, more commonly known as “translation difficulty estimation.” The aim is to predict the eventual quality of a translation or, con-

²wmt22-comet-da

³nrefs:1|case:lc|tok:tercom|norm:no|punct:yes|asian:no|version:2.5.1

⁴We note that while evaluating TER against a single post-edited reference inherently favors the candidate chosen as the starting point, the interface displayed all nine options simultaneously and post-editors actively borrowed from non-selected hypotheses (Marmonier et al., 2025), making TER meaningful across the entire pool.

versely, the post-editing effort it will require, a priori, without access to the translation itself, relying exclusively on the source segment.

3.1. Methodology

To assess the predictive power of various source-side features, we measure their correlation with the final, reference-based translation quality measured by COMET or TER. We evaluate twelve metrics drawn from different research paradigms, which can be categorized as follows:

- **Linguistic and Readability Metrics:** Following foundational approaches to this task (Mishra et al., 2013), we include nine traditional metrics; five standard readability formulas (Dale-Chall, SMOG index, Flesch Reading Ease, Coleman-Liau index, and Flesch-Kincaid Grade) and four linguistic complexity features (segment length, word rarity, degree of polysemy, and syntactic tree height). See Section 11 in the appendix for details.
- **Neural Metrics:** State-of-the-art performance is now typically achieved with dedicated neural models (Fernicola et al., 2023; Proietti et al., 2025). We include two such specialized predictors from the Sentinel family (sentinel-src-da and sentinel-src-mqm) (Proietti et al., 2025). Another prominent neural approach, which conceptually borders on the candidate-side, gauges difficulty via “glass box” capture of a translation model’s surprisal (Lim et al., 2024); to operationalize this as a single source-side heuristic for benchmarking, we calculated the MT surprisal for the NLLB-200 600M model on its own translations only, and assigned these scores to the source segments.

For our analysis, we use Kendall’s rank correlation coefficient (τ_b) to measure the association between each source-side metric and the ground-truth quality scores (TER and COMET). Kendall’s τ_b is a robust non-parametric test that measures the strength of association between two variables based on the similarity of their rankings. It functions by comparing the number of concordant pairs (pairs of segments that are in the same relative order for both the difficulty metric and the quality score) and discordant pairs (pairs that are in the opposite order). We conducted the analysis per system. Specifically, for each of the 6,028 segments, we ranked the list of source-side metric scores (e.g., *segment_length*) and compared this ranking against the corresponding list of ground-truth quality scores (TER or COMET) for a single translation system (e.g., NLLB-3.3B). This process was then repeated for every source metric and every system. Unlike Pearson’s r , which assumes

a linear relationship and normally distributed data, rank-based coefficients are more appropriate for MT quality scores, whose distributions often violate these assumptions. We prefer Kendall’s τ_b over Spearman’s ρ as it is generally considered more robust and reliable in the presence of tied ranks, not uncommon in this type of data.

3.2. Results and Analysis

The results of our correlation analysis are presented in Figures 2 and 3. A key finding is that the predictive power of source-side metrics is dependent on whether translation quality is measured with TER (post-editing effort) or COMET (regression on Direct Assessment scores from past WMT shared tasks).

When quality is measured by TER (Figure 2) the correlation patterns are generally modest. The most consistent positive predictor is *mt_surprisal* (τ between 0.14 and 0.16 across most systems). Notably, this surprisal—calculated here using the NLLB-200-distilled model for which it should therefore not be interpreted as a strictly source-side metric—correlates with the TER scores of translations from other systems, including the much larger MADLAD-400-3B and Llama-4. Other metrics show very weak correlations to TER. The Sentinel models, which are strong predictors of COMET, show a much weaker negative correlation with this edit rate (τ between -0.05 and -0.16), and *segment_length* and *tree_height* show almost no correlation with post-editing effort.

In contrast, when quality is measured by COMET (Figure 3), the predictive patterns are markedly different and the correlations appear much stronger. Here, the Sentinel neural metrics (*sentinel_da_score* and *sentinel_mqm_score*) are by far the most effective predictors, exhibiting consistent positive correlations across all systems (τ from 0.30 to 0.46). Similarly, *segment_length* (τ between -0.26 and -0.33) and *tree_height* (τ between -0.17 and -0.24) now show the most salient negative correlations.

We hypothesize that this apparent strength in the COMET-based correlations is attributable in part to system artifacts. (i) Data and Architectural Overlap: Both the Sentinel predictors and the COMET reference metric are derived from the same XLM-R-based architectures and are likely trained on overlapping data from WMT shared tasks. Their high correlation may therefore reflect this shared foundation rather than independent predictive power. (ii) Metric-Internal Bias: While *segment_length* and *tree_height* show weak correlation with TER, they are strong positive predictors for COMET. COMET might have derived a misleading heuristic from its training data: since longer source segments give more opportunity for errors in translation candidates




Figure 2: Kendall’s τ correlation between source-side metrics and translation quality as measured by TER. An asterisk (*) indicates a statistically significant correlation ($p < 0.05$).

that subsequently receive lower DA scores from human annotators, the metric might have learned to associate source length directly with low translation quality. TER does not reflect this bias, suggesting that length itself was not a primary driver of proportional post-editing effort.

Whatever its cause, this dichotomy demonstrates that the source-side features that best predict translation difficulty are contingent upon the specific reference-based metric used to define said quality.

4. Candidate-Side Experiments

Our second set of experiments shifts the focus from the source text to the translation candidates to address two key aspects of quality prediction: the predictive power of standard Quality Estimation (QE) metrics and of positional bias in LLM-based document-level MT.

We first conduct a “hindsight” QE evaluation. We assess the ability of two state-of-the-art reference-free QE metrics, COMET-QE⁵ (Rei et al., 2022) and MetricX-QE⁶ (Juraska et al., 2024), to predict the final quality of the nine system outputs. Their predictions are correlated against our gold-standard reference scores (TER and COMET) to measure

⁵wmt22-cometkiwi-da

⁶metricx-24-hybrid-xl-v2p6

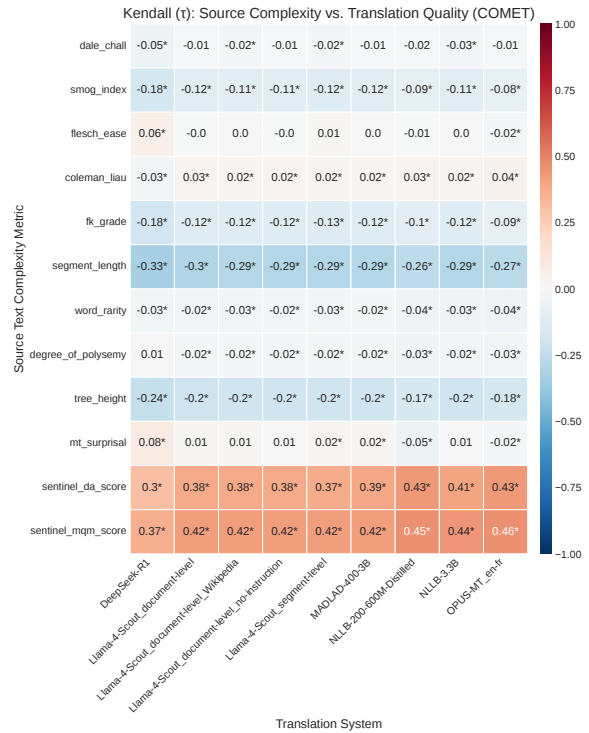


Figure 3: Kendall’s τ correlation between source-side metrics and translation quality as measured by COMET. An asterisk (*) indicates a statistically significant correlation ($p < 0.05$).

their performance across systems. Second, we investigate positional bias, a known challenge (Peng et al., 2024, 2025) where translation quality degrades as a model processes a long document.

4.1. Hindsight Quality Estimation

The original data collection interface (cf. Figure 1) was not anonymized; it displayed both the system name and a COMET-Kiwi (QE) score for each of the nine candidates. This introduces a strong potential for anchoring bias, as post-editors might have been influenced by the visible QE score when selecting a hypothesis to edit. Surprisingly, however, the data strongly suggests that post-editors were not captive to this bias. Table 3 shows the average rank of each system as determined by the COMET-Kiwi QE score from the original dataset. This ranking does not align with the final, reference-based quality rankings from Tables 1 and 2. For instance, DeepSeek-R1, the clear top system by both TER and COMET, was only ranked 5th on average by the QE model. Conversely, NLLB-3.3B was ranked 2nd by the QE model but was a mid-to-low performer in the final quality assessment. This discrepancy indicates that the human post-editors frequently overrode the QE metric’s ranking to select a candidate they deemed a better starting point.

System Name	Avg Rank	% 1st	# 1st	% 1st Solo	# 1st Solo	% Last	# Last
Llama-4-Scout_document-level_no-instruction	4.12	14.58%	903	1.63%	101	12.18%	754
NLLB-3.3B	4.12	19.62%	1215	10.45%	647	10.04%	622
MADLAD-400-3B	4.13	21.36%	1323	13.00%	805	11.06%	685
Llama-4-Scout_document-level	4.18	14.53%	900	1.37%	85	12.98%	804
DeepSeek-R1	4.27	23.53%	1457	16.83%	1042	15.52%	961
Llama-4-Scout_segment-level	4.35	18.57%	1150	11.30%	700	13.69%	848
Llama-4-Scout_document-level_Wikipedia	4.44	14.06%	871	4.25%	263	14.79%	916
NLLB-200-600M-Distilled	5.24	14.58%	903	8.96%	555	21.73%	1346
OPUS-MT_en-fr	5.47	15.36%	951	10.17%	630	27.92%	1729

Table 3: Analysis of COMET-Kiwi (QE) score distribution displayed to post-editors. Shows the average rank (lower is better), percentage/count of segments where a system was ranked 1st, percentage/count where it was ranked 1st with no ties (% 1st Solo), and percentage/count where it was ranked last.

Our hindsight QE evaluation confirms this hypothesis. The results are summarized in Figure 4. The predictive power of the QE metrics is found to not be uniform across systems, as both QE metrics correlate more strongly with the quality of the weaker MT systems than with the top-performing one. For instance, when predicting reference COMET, the correlation for *COMET_QE* is considerably higher for one of the weakest systems, OPUS-MT ($\tau = 0.53$), than for the strongest, DeepSeek-R1 ($\tau = 0.33$). This suggests that current QE models are better at identifying errors in mediocre translations than at distinguishing fine-grained quality differences among high-quality candidates. We hypothesize, based on observations made by the post-editors of the French OLDI Seed partition, that this weakness is rooted in a lack of the factual knowledge required to judge terminological accuracy in the encyclopedic domain to which the OLDI Seed Corpus belongs.

We note again that the high correlation between *COMET_QE* and reference COMET is likely to be an expected artifact of their shared base model (derived from XLM-R) and overlapping training data.⁷ A more telling benchmark is their performance against reference TER, where the same behavior is observed.

When the results are aggregated by system type as in Figure 5, an explanatory pattern emerges: the predictive power of the QE metrics is consistently stronger for the traditional NMT models than for the LLM-based systems. For example, *COMET_QE*’s average correlation with reference COMET is highest for the NMT group ($\tau = 0.49$), compared to the LLM group ($\tau = 0.43$). This significant trend⁸ holds across all conditions, as *MetricX_QE_Score* also shows its strongest correlation with reference TER for the NMT group ($\tau = 0.20$ vs. $\tau = 0.15$ for the LLM group).

The analysis by translation granularity suggests

⁷The same observation applies to MetricX, though to a lesser extent.

⁸We confirmed the statistical significance of these differences in correlation strength using bootstrap resampling ($p < 0.05$).

that this is an architectural effect, not simply a matter of context size. The *smallDoc2sent* LLM, despite its limited taxing of the context window, patterns with the *doc2doc* LLMs rather than the *sent2sent* NMTs. This finding suggests that the predictions of current QE models are more closely aligned with the outputs of specialized NMT architectures than with those of general-purpose LLMs, explaining why post-editors frequently overrode QE recommendations to select LLM hypotheses as starting points.

4.2. Positional Bias in Document-Level Translation

Our final experiment investigates positional bias in document-level MT, a phenomenon where a segment’s translation quality degrades the later it appears within a document, and its candidate-side predictive power.

To measure a segment’s position from a model-centric perspective, we calculated its cumulative token rank: the total number of tokens (from the prompt, thought tokens in the case of DeepSeek-R1, and prior generation) processed by a given *doc2doc* model before it began translating the segment in question, with each model’s specific tokenizer. We then correlated this rank with translation quality.

Correlating this rank directly with raw quality scores reveals a very weak but statistically significant negative trend (e.g., Figure 6, top; $\tau = -0.0325$ for DeepSeek-R1). However, this simple correlation is hard to interpret. The inherent difficulty of the source text acts as a major confounding factor, making it difficult to disentangle the effect of position from the effect of the source material itself.

To isolate the impact of position, we normalize the scores by calculating a “delta score”: the difference between a *doc2doc* system’s score and the average score of all baseline *sent2sent* systems for the same segment. This score delta better represents the performance difference attributable specifically to the *doc2doc* processing strategy, providing a clearer signal of positional




Figure 4: Kendall’s τ correlation between two reference-free QE metrics (*COMET_QE*, *MetricX_QE*) and our gold-standard reference scores. The top plot shows correlations against reference TER, while the bottom plot uses reference COMET. An asterisk (*) indicates statistical significance ($p < 0.05$).

degradation. Figure 6 illustrates this process for the DeepSeek-R1 system and its COMET scores. The raw COMET scores show a weak but statistically significant negative correlation with token rank ($\tau = -0.0325$, $p = 0.0002$), confirming the presence of a positional bias. After applying our normalization, this negative trend persists ($\tau = -0.0295$, $p = 0.0006$), demonstrating that the performance degradation is not merely an artifact of source difficulty but is linked to the segment’s position within the document.

The full results of this positional bias analysis are presented in Figure 7. The normalization reveals two key insights into the behavior of the different `doc2doc` systems. First, for the Llama-4-Scout model settings, applying the delta-score normalization noticeably strengthens the negative correlation between position and quality. For example, the COMET score correlation for the base document-level model nearly doubles, from $\tau = -0.022$ to $\tau = -0.043$. This confirms that for these systems, a subtle positional bias is indeed present and is more clearly isolated once the confounding effect of source difficulty is removed. In contrast, the correlation for the top-performing DeepSeek-R1 remains remarkably stable ($\tau \approx -0.03$) with or without normalization. This suggests that DeepSeek-R1’s performance is less sensitive to the intrinsic difficulty

of the source text to begin with; its robustness to source-side challenges means that the normalization has little effect.

However, the most critical finding across all systems is that the absolute magnitude of the correlation remains very low ($|\tau| < 0.05$). While we can confirm that positional bias is a statistically detectable phenomenon, its practical impact appears to be minimal for this task, suggesting it does not represent a bottleneck for document-level translation quality anymore with current advanced, long-context models.

5. Related Work

Early efforts at predicting translation quality from the source text alone framed the task as estimating translation difficulty, relying on readability metrics (used in translation studies research for the same purpose), linguistic features such as sentence length, degree of polysemy, and syntactic complexity to approximate the cognitive effort and expected error rates of human translators (Mishra et al., 2013). Related work has recently shifted from these feature-based methods toward supervised neural models. This new paradigm involves training encoders to regress directly on human quality judgments, like Direct Assessment (DA; cf. Gra-

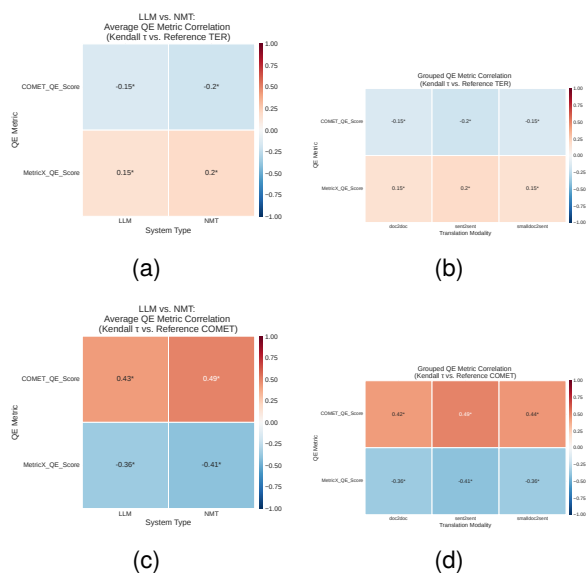


Figure 5: Kendall's τ correlation between QE metrics and gold-standard scores, averaged across system groups. The top row shows correlations against reference TER; the bottom row uses reference COMET. The left column groups by system type (LLM vs. NMT), while the right column groups by translation granularity.

ham et al. 2013) or MQM scores (Lommel et al., 2013), using only the source text as input (Fericola et al., 2023; Proietti et al., 2025). Other neural methods have also proven effective, for instance by showing that MT surprisal strongly correlates with human-assessed translation difficulty (Lim et al., 2024).

On the candidate side, QE evolved from early work on confidence estimation for MT (Jeffering et al., 2003; Gandrabur and Foster, 2003; Blatz et al., 2004) into a formal research program with the launch of the WMT QE shared task in 2012 (Callison-Burch et al., 2012), with successive WMT editions broadening its scope. The field has now shifted towards systems that directly regress to DA scores, rather than predicting post-editing effort, enabling the pre-training of QE models on annotator data from WMT's shared tasks (Rei et al., 2020).

The shift towards document-level MT has reintroduced challenges related to input length that affect translation quality (Maruf et al., 2021). While general length biases, where performance degrades on documents longer than those seen during training, are a known issue (Zhuocheng et al., 2023, *inter alia*), a more subtle problem is that of position bias. Recent studies have shown that the quality of translation can decrease for sentences appearing later in a document (Peng et al., 2024, 2025), suggesting that a sentence's absolute position within the input sequence impacts its final translation. This issue echoes challenges from early neural MT at

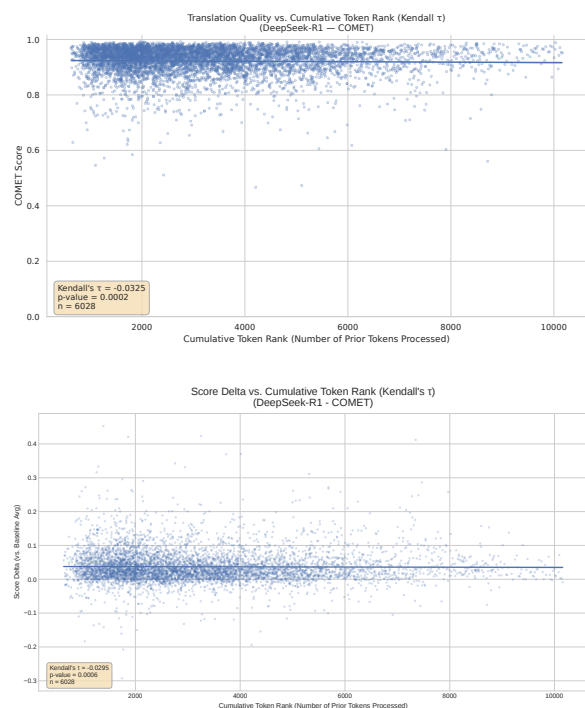


Figure 6: Scatter plots illustrating the effect of positional bias on the DeepSeek-R1 system. The top plot shows the weak negative correlation between raw COMET scores and the cumulative token rank. The bottom plot uses our normalized “delta score,” which controls for source-text difficulty.

the segment level, where strategies such as reversing the source sentence were developed to address position-dependent performance issues in segment-level models (Sutskever et al., 2014). These issues were solved by the advent of attentive models (Bahdanau et al., 2015).

6. Conclusion

This paper presented a series of quality prediction experiments using a unique dataset generated from a real-world, multi-candidate post-editing workflow. Our “hindsight” analyses yield three primary findings.

First, on the source side, we find that the predictive power of translation difficulty metrics is contingent on the ground-truth metric used to measure quality. Heuristics like segment length and the neural Sentinel predictors show strong correlations with COMET. However, we hypothesize that these are in part artifacts stemming from shared model architectures and internal metric biases, as these same features show weaker to no correlation with TER, our proxy for post-editing effort.

Second, on the candidate side, our analysis of the original post-editing process reveals a signif-

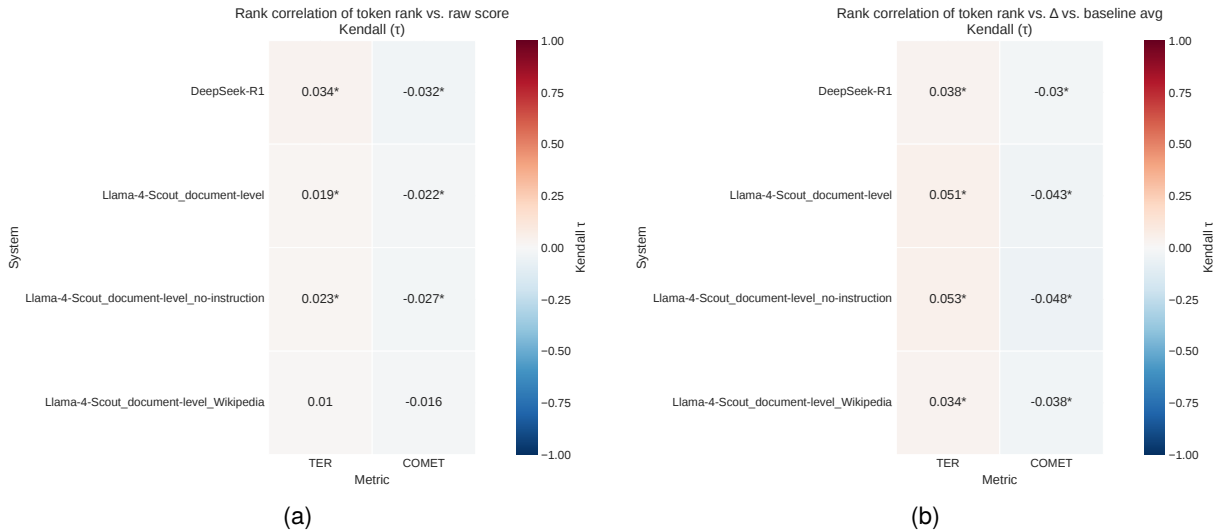


Figure 7: Kendall’s τ correlation between a segment’s cumulative token rank and its quality for all `doc2doc` systems. Plot (a) uses raw quality scores, while plot (b) uses our normalized delta scores.

icant mismatch between the QE model rankings provided to the annotators and their final, human-adjudicated quality, demonstrating that post-editors actively overrode the QE model’s guidance. This indicates that reference-free QE metrics are not uniformly aligned with human judgment. We further show that their predictions are more strongly aligned with traditional NMT architectures than with general-purpose LLMs, particularly for top-performing systems, a finding which may help explain this observed discrepancy.

Third, while we confirm the existence of a statistically significant positional bias in document-level LLMs, its practical impact on translation quality appears to be negligible in current long-context models.

Taken together, these findings suggest that the architectural shift towards LLMs alters the reliability of established quality prediction methods while also mitigating previous challenges in document-level MT.

7. Limitations

Our study’s findings are subject to several limitations, primarily related to the dataset and the experimental scope.

All experiments are conducted on a single language pair, in a single translation direction (English-French) and a single, specific domain (encyclopedic content). The conclusions—particularly regarding the alignment of predictive models with NMT vs. LLMs and the negligible impact of positional bias—are robust for this dataset but may not generalize to other language pairs, other domains (e.g., creative, conversational, or legal texts), or other LLMs not included in our study.

Finally, a conceptual limitation lies in our operational definition of “quality prediction” itself. This study measures predictive power by correlating various heuristics and QE models against two automated, reference-based metrics: COMET and TER. These ground-truth metrics are, themselves, only proxies for the concepts they aim to capture and fallible (cf. Zouhar et al. 2024, inter alia). A deep methodological critique of these foundational metrics is beyond the scope of this paper. We contend that our approach remains sound, as automated metrics are the *de facto* standard in MT evaluation and the language service industry. Furthermore, our analysis benefits from presenting results against these two distinct and complementary definitions of quality. This dual approach prevents either metric from being taken as an absolute “truth” in isolation.

8. Acknowledgements

This work was funded by the French *Agence Nationale de la Recherche* (ANR) under the project TraLaLaM (“ANR-23-IAS1-0006”), as well as by Inria under the “*Défi*”-type project COLaF. This work was partly funded by Rachel Bawden and Benoît Sagot’s chairs in the PRAIRIE institute, funded by the French national agency ANR, as part of the “Investissements d’avenir” programme under the reference ANR-19-P3IA-0001 and by Benoît Sagot’s chair in its follow-up, PRAIRIE-PSAI, also funded by the ANR as part of the “France 2030” strategy under the reference ANR23-IACL-0008.

This work was also granted access to the HPC resources of IDRIS under allocations 2025-AD011015117R2 and 2025-AD011017228 made by GENCI.

9. Bibliographical References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015)*, San Diego, CA, USA. Oral presentation.
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- John Blatz, Erin Fitzgerald, George Foster, Simona Gandrabur, Cyril Goutte, Alex Kulesza, Alberto Sanchis, and Nicola Ueffing. 2004. [Confidence estimation for machine translation](#). In *COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics*, pages 315–321, Geneva, Switzerland. COLING.
- Chris Callison-Burch, Philipp Koehn, Christof Monz, Matt Post, Radu Soricut, and Lucia Specia. 2012. [Findings of the 2012 Workshop on Statistical Machine Translation](#). In *Proceedings of the Seventh Workshop on Statistical Machine Translation*, pages 10–51, Montréal, Canada. Association for Computational Linguistics.
- Aaron Chatterji, Thomas Cunningham, David J. Deming, Zoë Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. 2025. [How People Use ChatGPT](#). Working Paper 34255, National Bureau of Economic Research.
- Meri Coleman and Ta Lin Liao. 1975. A computer readability formula designed for machine scoring. *Journal of Applied Psychology*, 60(2):283–284.
- David Dale, Laurie V. Burchell, Jean Maillard, Idris Abdulmumin, Antonios Anastasopoulos, Isaac Caswell, and Philipp Koehn. 2025. [Findings of the WMT 2025 Shared Task of the Open Language Data Initiative](#). In *Proceedings of the Tenth Conference on Machine Translation (WMT 2025)*, pages 495–502, Suzhou, China. Association for Computational Linguistics.
- Edgar Dale and Jeanne S. Chall. 1948. A formula for predicting readability: Instructions. *Educational Research Bulletin*, pages 37–54.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanbiao Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yudian Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. [DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning](#).
- Francesco Ferrnocola, Silvia Bernardini, Federico Garcea, Adriano Ferraresi, and Alberto Barrón-Cedeño. 2023. [Return to the source: Assessing machine translation suitability](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 79–

- 89, Tampere, Finland. European Association for Machine Translation.
- Rudolph Flesch. 1948. A new readability yardstick. *Journal of Applied Psychology*, 32(3):221–233.
- Simona Gandrabur and George Foster. 2003. Confidence estimation for translation prediction. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 95–102.
- Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2013. Continuous Measurement Scales in Human Evaluation of Machine Translation. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 33–41, Sofia, Bulgaria. Association for Computational Linguistics.
- Juraj Juraska, Daniel Deutsch, Mara Finkelstein, and Markus Freitag. 2024. MetricX-24: The Google Submission to the WMT 2024 Metrics Shared Task. In *Proceedings of the Ninth Conference on Machine Translation*, pages 492–504, Miami, Florida, USA. Association for Computational Linguistics.
- J. Peter Kincaid, Robert P. Fishburne Jr, Richard L. Rogers, and Brad S. Chissom. 1975. Derivation of New Readability Formulas (Automated Readability Index, Fog Count and Flesch Reading Ease Formula) for Navy Enlisted Personnel. Technical report, Naval Technical Training Command Millington TN Research Branch.
- Zheng Wei Lim, Ekaterina Vylomova, Charles Kemp, and Trevor Cohn. 2024. Predicting human translation difficulty with neural machine translation. *Transactions of the Association for Computational Linguistics*, 12:1479–1496.
- Arle Richard Lommel, Aljoscha Burchardt, and Hans Uszkoreit. 2013. Multidimensional quality metrics: a flexible system for assessing translation quality. In *Proceedings of Translating and the Computer 35*, London, UK. Aslib.
- Jean Maillard, Cynthia Gao, Elahe Kalbassi, Kaushik Ram Sadagopan, Vedanuj Goswami, Philipp Koehn, Angela Fan, and Francisco Guzman. 2023. Small Data, Big Impact: Leveraging Minimal Data for Effective Machine Translation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2740–2756, Toronto, Canada. Association for Computational Linguistics.
- Sameen Maruf, Fahimeh Saleh, and Gholamreza Haffari. 2021. A survey on document-level neural machine translation: Methods and evaluation. *ACM Computing Surveys*, 54(2).
- G. Harry McLaughlin. 1969. Smog grading—a new readability formula. *Journal of Reading*, 12(8):639–646.
- Meta AI. 2025. The Llama 4 herd: The beginning of a new era of natively multimodal AI innovation.
- George A. Miller. 1994. WordNet: A lexical database for English. In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Abhijit Mishra, Pushpak Bhattacharyya, and Michael Carl. 2013. Automatically predicting sentence translation difficulty. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 346–351, Sofia, Bulgaria. Association for Computational Linguistics.
- 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, Sermarley 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.
- Joseph Olive. 2005. Global autonomous language exploitation (GALE). *DARPA/IPTO Proposer Information Pamphlet*.
- Ziqian Peng, Rachel Bawden, and François Yvon. 2024. À propos des difficultés de traduire automatiquement de longs documents. In *35èmes Journées d’Études sur la Parole (JEP 2024)*, volume 1 : articles longs et prises de position, pages 2–21, Toulouse, France. ATALA & AFPC.
- Ziqian Peng, Rachel Bawden, and François Yvon. 2025. Investigating length issues in document-level machine translation. In *Proceedings of Machine Translation Summit XX: Volume 1*, pages 4–23, Geneva, Switzerland. European Association for Machine Translation.
- Maja Popović. 2015. chrF: character n-gram F-score for automatic MT evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.

Lorenzo Proietti, Stefano Perrella, Vilém Zouhar, Roberto Navigli, and Tom Kocmi. 2025. [Estimating machine translation difficulty](#).

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.

Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022. [CometKiwi: IST-Unbabel 2022 Submission for the Quality Estimation Shared Task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Matthew Snover, Bonnie Dorr, Rich Schwartz, Linea Micciulla, and John Makhoul. 2006. [A Study of Translation Edit Rate with Targeted Human Annotation](#). In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 223–231, Cambridge, Massachusetts, USA. Association for Machine Translation in the Americas.

Robyn Speer. 2022. [rspeer/wordfreq: v3.0](#). Zenodo, Version v3.0.2.

Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2, NeurIPS'14*, page 3104–3112, Cambridge, MA, USA. MIT Press.

Jörg Tiedemann and Santhosh Thottingal. 2020. [OPUS-MT – Building open translation services for the World](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 479–480, Lisboa, Portugal. European Association for Machine Translation.

Nicola Ueffing, Klaus Macherey, and Hermann Ney. 2003. [Confidence measures for statistical machine translation](#). In *Proceedings of Machine Translation Summit IX: Papers*, New Orleans, USA.

Zhang Zhuocheng, Shuhao Gu, Min Zhang, and Yang Feng. 2023. [Addressing the length bias challenge in document-level neural machine translation](#). In *Findings of the Association for*

Computational Linguistics: EMNLP 2023, pages 11545–11556, Singapore. Association for Computational Linguistics.

Vilém Zouhar, Pinzhen Chen, Tsz Kin Lam, Nikita Moghe, and Barry Haddow. 2024. [Pitfalls and outlooks in using COMET](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 1272–1288, Miami, Florida, USA. Association for Computational Linguistics.

10. Language Resource References

Sneha Kudugunta, Isaac Caswell, Biao Zhang, Xavier Garcia, Christopher A. Choquette-Choo, Katherine Lee, Derrick Xin, Aditya Kusupati, Romi Stella, Ankur Bapna, and Orhan Firat. 2023. [MADLAD-400: A Multilingual And Document-Level Large Audited Dataset](#).

Malik Marmonier, Benoît Sagot, and Rachel Bawden. 2025. [A French Version of the OLDI Seed Corpus](#). In *Proceedings of the Tenth Conference on Machine Translation (WMT 2025)*, pages 1048–1060, Suzhou, China. Association for Computational Linguistics.

11. Appendix: Readability and Linguistic Complexity Metrics

This appendix details the mathematical formulations of the standard readability metrics referenced in Section 3.1. All metrics were computed using the `textstat` Python library.⁹

To present the formulas concisely, we define the following variables for a given text:

- W : total number of words
- S : total number of sentences
- Sy : total number of syllables

Flesch Reading Ease (FRE)

The FRE score (`flesch_ease` in the figures) evaluates readability based on the average number of syllables per word and words per sentence (Flesch, 1948). Higher scores indicate easier reading. Note that the constants used in this formula (206.835, 1.015, and 84.6) are specifically calibrated for the English language and defined by the library.

$$FRE = 206.835 - 1.015 \left(\frac{W}{S} \right) - 84.6 \left(\frac{Sy}{W} \right)$$

⁹<https://github.com/textstat/textstat>

Flesch-Kincaid Grade Level (FKGL)

The FKGL (*fk_grade* in the figures) translates the Flesch Reading Ease score into a U.S. grade-level equivalent (Kincaid et al., 1975).

$$FKGL = 0.39 \left(\frac{W}{S} \right) + 11.8 \left(\frac{S_y}{W} \right) - 15.59$$

SMOG Index

The Simple Measure of Gobbledygook (SMOG; *smog_index* in the figures) estimates the years of education needed to fully understand a piece of writing, relying heavily on the count of polysyllables (P), defined as words with three or more syllables (McLaughlin, 1969).

$$SMOG = 1.043 \sqrt{P \times \frac{30}{S}} + 3.1291$$

Coleman-Liau Index (CLI)

Unlike Flesch and SMOG, the CLI (*coleman_liau* in the figures) relies on character counts instead of syllables, making it computationally robust for automated processing (Coleman and Liau, 1975). Let L be the average number of letters per 100 words, and S_{100} be the average number of sentences per 100 words.

$$CLI = 0.0588L - 0.296S_{100} - 15.8$$

Dale-Chall Readability Formula

This formula (represented as *dale_chall* in the figures) calculates a score based on average sentence length and the percentage of “difficult” words (D)—defined as words not present on a specific list of common words familiar to most fourth-grade students (Dale and Chall, 1948).

$$\text{Score} = 0.1579 \left(\frac{D}{W} \times 100 \right) + 0.0496 \left(\frac{W}{S} \right)$$

If the percentage of difficult words is greater than 5%, an adjusted score is calculated by adding 3.6365 to the final result.

Linguistic Complexity Features

In addition to the readability metrics, we calculated four linguistic complexity features referenced in Section 3.1. These were computed using a combination of the `spaCy`,¹⁰ `NLTK` (Bird and Loper, 2004), and `wordfreq` (Speer, 2022) Python libraries. Let W_{space} be the set of whitespace-separated tokens, and W_α be the subset of tokens containing only alphabetic characters.

¹⁰<https://github.com/explosion/spaCy>

Segment Length Segment length is defined strictly as the total number of whitespace-separated tokens in the source segment:

$$\text{Length} = |W_{space}|$$

Word Rarity Word rarity evaluates the average infrequency of the vocabulary used in the segment. It is calculated as the average negative natural logarithm of the word frequencies extracted via the `wordfreq` library. A smoothing factor of 10^{-9} is added to prevent undefined logarithms for out-of-vocabulary words.

$$\text{Rarity} = -\frac{1}{|W_\alpha|} \sum_{w \in W_\alpha} \ln(\text{freq}(w) + 10^{-9})$$

Degree of Polysemy The degree of polysemy captures lexical ambiguity by averaging the total number of WordNet synsets (Miller, 1994) available for each alphabetic token, retrieved using the `NLTK` library.

$$\text{Polysemy} = \frac{1}{|W_\alpha|} \sum_{w \in W_\alpha} |\text{synsets}(w)|$$

Syntactic Tree Height To measure structural complexity, we calculate the syntactic tree height. Using the `en_core_web_lg` dependency parser from the `spaCy` library, we extract the dependency tree of the segment. The height is defined recursively as the maximum depth of the tree starting from the root node.

12. Appendix: Supplementary Results using the ChrF Metric

The main body of this paper evaluates translation quality primarily through TER (as a proxy for post-editing effort) and COMET (as a proxy for human judgment). Due to space constraints in the main text, we present the supplementary results using the ChrF metric (Popović, 2015) in this appendix.

The following figures replicate the primary correlation experiments from Sections 3 and 4, substituting the ground-truth reference metrics with ChrF to demonstrate how the predictive power of our source-side and candidate-side features behaves under this character-level surface-form evaluation paradigm.

As illustrated in the supplementary Figures 8, 9, 10, 11 and 12, the predictive behavior of both source-side and candidate-side metrics against ChrF closely mirrors their behavior against TER, confirming the observations and metric dichotomies detailed in Sections 3 and 4.




Figure 8: Kendall’s τ correlation between source-side metrics and translation quality as measured by ChrF. An asterisk (*) indicates a statistically significant correlation ($p < 0.05$).

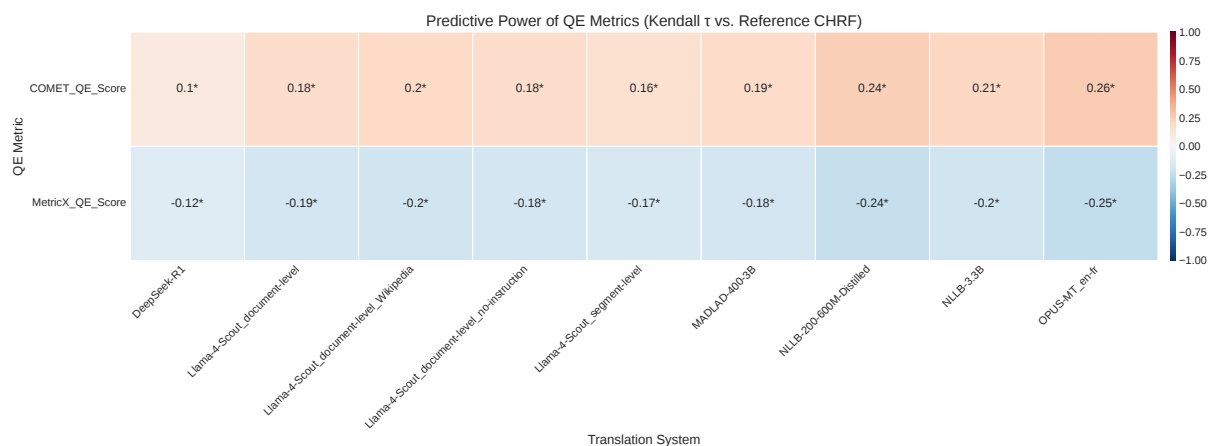


Figure 9: Kendall's τ correlation between two reference-free QE metrics (*COMET_QE*, *MetricX_QE*) and our gold-standard reference scores measured by ChrF. An asterisk (*) indicates statistical significance ($p < 0.05$).




Figure 10: Kendall's τ correlation between QE metrics and gold-standard scores measured by ChrF, averaged across system groups. The left figure (a) groups by system type (LLM vs. NMT), while the right figure (b) groups by translation granularity.

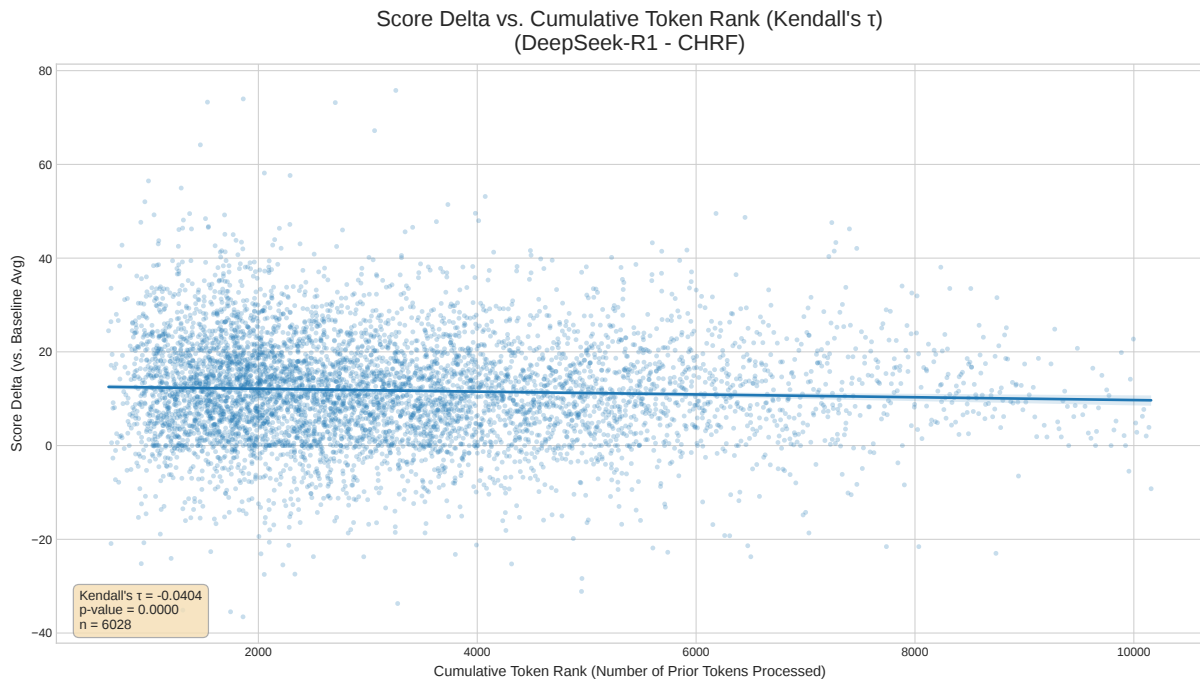


Figure 11: Scatter plots illustrating the effect of positional bias on the DeepSeek-R1 system. The plot shows the weak negative correlation between delta-normalized ChrF scores and the cumulative token rank, controlling for source-text difficulty.

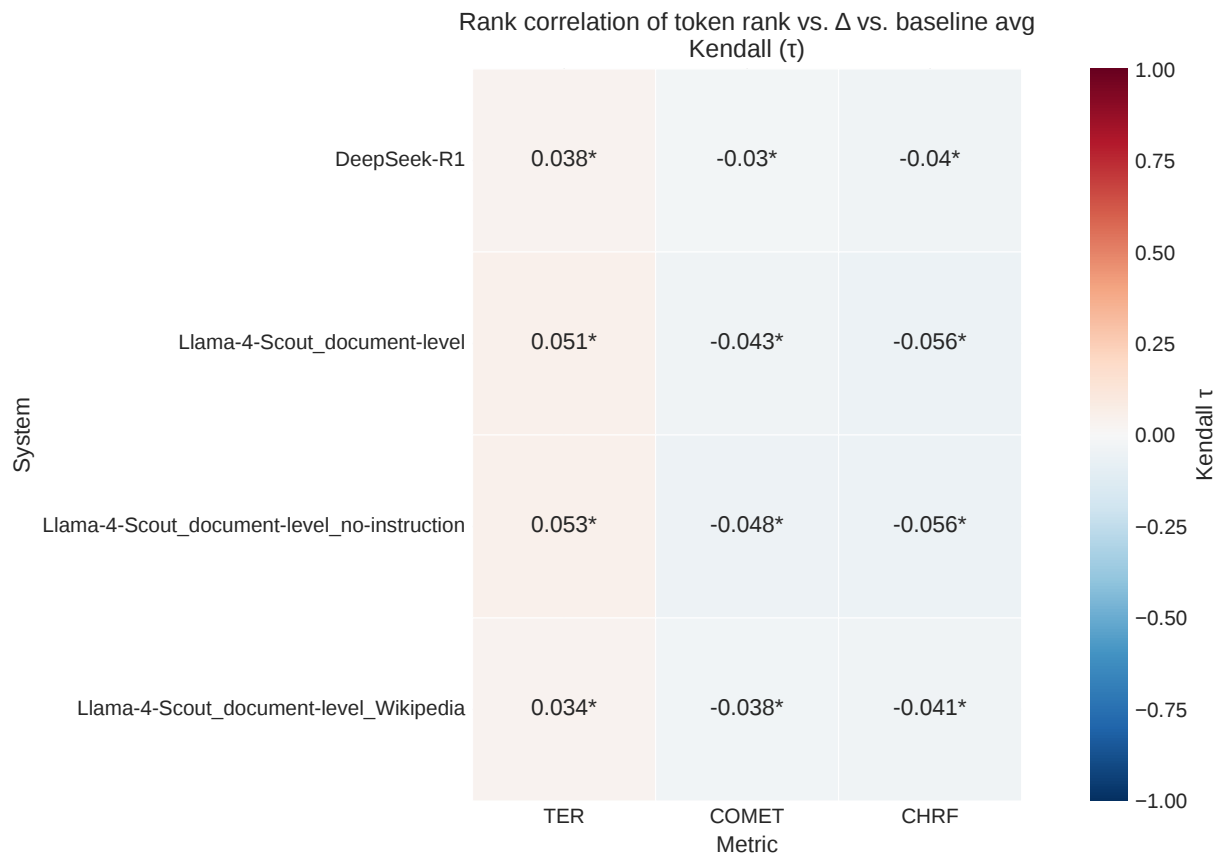


Figure 12: Kendall's τ correlation between a segment's cumulative token rank and its quality for all doc2doc systems. The plot uses our normalized delta scores and includes ChrF scores.