

Pterodactyl: Pushing the Limits of Efficient Automated Translation Evaluation

Evelyn Y. Garland^{1,2} and Carola F. Berger^{1,3}

¹1020 Group, ²Acta-Transphere, ³CFB Scientific Translations LLC

Correspondence: project1020group@gmail.com

Abstract

In this paper, we present 1020 Group’s submission to the WMT26 shared task on automated quality evaluation. Our system, Pterodactyl, was engineered specifically for subtasks 2 and 3, that is, segment-level quality score prediction and the detection of error-free segments. Pterodactyl is a highly customizable, ultra-low-cost system optimized for binary classification in translation triage and data curation. It is based on an LLM-as-a-judge approach, with prompt engineering and model selection guided by ROC analysis. We also developed several variants of Pterodactyl that were specifically adapted for subtasks 3.

1 Introduction

In recent WMT competitions, reference-free quality estimation (QE) systems based on an LLM-as-a-judge approach have gained increasing traction for automated translation quality evaluation (Lavie et al., 2025; Freitag et al., 2024; Zerva et al., 2024). However, the development and deployment of such systems can be complex and costly. Here, we develop an efficient, production-oriented strategy for pushing the cost-performance frontier of automated translation evaluation: achieving better task-specific performance with faster and cheaper models and prompt configurations.

Below, we present 1020 Group’s submission to the WMT26 shared task on automated quality evaluation. Our system, Pterodactyl, and its variants are intended as highly customizable, ultra-low-cost, and robust options for performing segment-level score prediction (WMT26 automated quality evaluation subtask 2) and detecting error-free segments (WMT26 automated quality evaluation subtask 3). The system is highly customizable thanks to an efficient, streamlined engineering process that can be implemented to select and build fit-for-purpose QE systems for different application scenarios. When possible, we avoid separately optimizing every

component when a shared screening criterion might suffice. Pterodactyl and its variants are also ultra-low-cost, in terms of both time and tokens, because our approach allows us to minimize cost and latency when two or more options show similar performance on test data. Finally, our approach allows engineering a system that remains relatively stable in different testing scenarios and under variations in the gold standard caused by inter-annotator disagreement.

2 The Pterodactyl Premise

Our approach is based on the LLM-as-a-judge paradigm, whereby our prompt engineering methodology and the selection of the most suitable models were guided by Receiver Operating Characteristic (ROC) analysis (Fawcett, 2006), supplemented by additional considerations. Our goal was to engineer a robust system that is simple to implement and adapt as well as sufficiently low-cost to be used in large production runs.

2.1 Basic Considerations

To develop our system, we started with commercially available general-purpose models and performed prompt engineering based on the considerations detailed below.

Our starting point was the basic premise that a QE system that is able to separate good and bad translations well is a promising candidate for WMT26 subtasks 2 and 3. We have shown in (Garland and Berger, 2026) that ROC analysis provides an efficient screening tool to find such candidates among a pool of systems. Prompt and model selection for subtasks 2 and 3 can therefore be partially unified with this ROC screening method.

We tested the following additional considerations during our prompt engineering iterations for both subtasks 2 and 3:

- Whether newer (larger) frontier models are

- better for translation quality estimation
- Whether more reasoning improves estimation performance

2.2 Further Considerations for Subtask 3

For subtask 3, further considerations come into play. Here, translations are to be triaged into error-containing and error-free segments, with a human gold standard obtained via contrastive error span annotation (cESA) (Zouhar et al., 2026). Submissions for subtask 3 are evaluated and ranked with reference to this cESA gold standard, as measured by the Matthews Correlation Coefficient (MCC). Since cESA is a new human annotation method for which no extensive development data is available, our strategy was to optimize our system for robustness with respect to the exact score threshold that separates the error-free class from the error-containing one.

What do we mean by robustness? Past WMT tasks used standard error span annotation (ESA) (Kocmi et al., 2024), with slightly different guidelines and anchor specifications (see, e.g. (Lavie et al., 2025) and references therein). It is therefore a priori unknowable if and how the WMT26 gold standard differs from prior years’ ESA gold standards. Thus, an error-classifying threshold that may be optimal for our system for WMT25 data may not be entirely optimal for WMT26. As a consequence, we did not optimize our prompts to obtain maximum MCC when computed on earlier WMT development

data, but rather to obtain a robustness plateau, that is, a range within which small variations in the error-classifying threshold lead to relatively minor changes in the MCC value. This can be observed in Figure 1.

Figure 1 shows two different systems with similar ROC performance, but quite different MCC versus threshold behavior. Example Systems A and B have comparable ROC curves and ROC-AUC values, but their MCC versus QE threshold behavior differs. While System A has a slightly lower best MCC value compared to System B, the peak of System A’s MCC distribution is decidedly broader than that of System B.

3 Our System Engineering Approach

3.1 Sample Data and Baselines

For the prompt engineering iterations described in the next subsection, we used data from past Conferences on Machine Translation (WMT) that were available in a public GitHub repository <https://github.com/google-research/mt-metrics-eval>, licensed under Apache v2.0. Specifically, we used the English into Chinese data with ESA annotations from WMT25 (Lavie et al., 2025), and the English into German data from WMT23 in the same repository that were supplemented with ESA annotations by (Kocmi et al., 2024; Zouhar et al., 2024). These ESA annotations for English into German WMT23 data are available at <https://github.com/wmt-conference/>




Figure 1: ROC curves (left) and MCC versus error-classifying threshold curves (right) for two example systems.

errors span annotation. For our prompt iterations, we split each of these datasets into a development set containing approximately 70% of the en→zh and en→de datasets, respectively, and a test set containing the remaining 30%.

Our main gold standard was defined in a manner similar to WMT26 subtask 3: An error-free segment is defined as having an ESA score ≥ 85 and no error span annotation. Where two sets of human annotations were available, we took the average of the ESA scores ≥ 85 , and either no error span annotation in annotation set 1 or no error span annotation in set 2. It turned out that our relative screening method via ROC analysis is relatively robust to changes in the definition of the gold standard.

And finally, to study an example of an ultra-low-resource language for subtask 3, we used the English into Northern Sámi recipe wmt26-eng-sme from the WMT26 General MT task datasets. From this recipe, we created a somewhat artificial development set as follows: We randomly selected 2000 longer segments from the recipe, then randomly selected 1000 segments from these as an error-free gold standard, and modified the other 1000 segments with random punctuation errors, omissions and additions of random words in Northern Sámi, as well as random source text insertions to obtain 1000 error-containing gold standard segments. Of course, this is nowhere near equivalent to a real development dataset, but given the fact that none of the authors speaks Northern Sámi, this is the closest available approximation.

As our baselines, we ran CometKiwi22 (Rei et al., 2022) and xCOMET-XL (Guerreiro et al., 2023) on these development and test data.

For testing and developing prompts for English into Northern Sámi, we also utilized the open-source LLM NorMistral-11B (Samuel et al., 2025b,a), with prompts in the Norwegian language instead of English.

3.2 Iterative Prompt Engineering

Our prompt engineering approach was iterative, discarding prompts that are no better than our baselines, while attempting to improve the promising ones that exhibited the best ROC-AUC performance. Once this iteration converged to a few prompts with similar ROC performance, we selected those that showed a robust MCC performance as detailed in Section 2.2. As our final system and system variants, we chose the ones

that exhibited the minimum cost and latency. This highly efficient approach can be automated and easily adapted to other scenarios outside the WMT realm.

We experimented with a variety of prompts. On the basis on linguistic and other technical considerations, we identified six distinct main strategies. After performing ROC analysis on the results, we discarded those that were clearly inferior. On the prompts that seemed promising, we iterated a number of times, incorporating various refinements. Our findings are reported in the next section.

4 Submitted Systems

Using the aforementioned iterative prompt engineering approach, we were able to answer some of the questions in Section 2.1 and converge on a final system. The one question we were unable to answer is whether the newest models, which were presumably trained on the largest and most recent training data, are able to handle ultra-low-resource languages better than slightly older models, given the same prompt. This led us to submit two system variants to subtask 2 (see Table 1), the primary system using the same cost-efficient general-purpose model for all languages, and a variant using one of the newest available models for the three language pairs English into Ladin, Ligurian, and Northern Sámi. The official results of subtask 2 will tell which option is better.

As a side note, while we used NorMistral-11B for testing and development, we did not end up incorporating it into our system framework. This is because we did not perform a separate prompt engineering optimization for Northern Sámi, in part because Norwegian seems to be a better prompting language for NorMistral-11B than English, and none of the authors speaks Norwegian. With the same prompt as for all the other languages, the commercial general purpose models we studied seem to perform better on our artificial mock dataset for Northern Sámi, but this may well be because the prompt was optimized for these models, not for NorMistral-11B, or perhaps because our artificial mock dataset is constructed somewhat arbitrarily, or a combination of both.

The final selected prompt instructs the LLM to evaluate the translation quality for each segment individually, without requiring the model to identify specific translation errors or to justify its score. Figure 2 shows the selected system’s ROC

curve compared to our baselines, CometKiwi22 and xCOMET-XL, on the WMT25 English into Chinese data. Figure 2 also lists the corresponding ROC-AUC values with their 95% confidence intervals. Figure 3 shows the selected system’s MCC versus threshold behavior.




Figure 2: ROC curve of the chosen prompt and model combination compared to the baselines CometKiwi22 and xCOMET-XL on the WMT25 en→zh dataset.




Figure 3: MCC versus error-classifying threshold curve for the selected system on the WMT25 en→zh dataset. The red diamond indicates the maximum MCC value.

We ran our scoring in batches of up to 20 segments produced by the same MT system for the same language pair, because this batching procedure leads to lower token usage and cost.

For subtask 3, we submitted two system variants that are derived from Pterodactyl and specifically optimized for subtask 3, employing a threshold-setting strategy based on patterns observed in the WMT25 data.

Table 1 gives an overview of our primary system and its variant that were submitted to subtasks 2 and 3, along with the computational cost involved in scoring the official WMT26 test data. Table 2 lists the variants we submitted specifically to subtask 3.

5 Conclusions and Outlook

Pterodactyl’s central design principle was deliberate minimalism: use no more model or engineering complexity than the task requires. Our submission to subtasks 2 and 3 of the WMT26 shared task on automated quality evaluation is based on an LLM-as-a-judge approach. Our prompt engineering and model selection process uses ROC analysis as a screening tool, making the development highly adaptable, robust, and low-cost. It remains to be seen, however, whether this minimal approach is an effective alternative to fine-tuning or building from scratch. The shared-task results will show whether Pterodactyl merely costs less—or whether it offers a genuinely competitive alternative to more resource-intensive system development.

The astute reader will of course observe that we have committed the cardinal sin of extrapolating from not necessarily representative development data into the great unknown of WMT26, which includes new language pairs and a novel human annotation scheme. But this leap of faith is not entirely unfounded, based on high-level patterns we observed across languages and domains from the WMT25 and WMT23 data, supplemented by our “artificial” Northern Sámi data. In an actual production environment, such an extrapolation is of course unnecessary, and the system can easily be adapted to real in-domain and in-language data for production scenarios, as described above.

6 Limitations

As already mentioned, the exact human gold standard for WMT26 is unknowable, because in prior WMT tasks, standard ESA was employed, whereas this year’s task utilizes contrastive ESA for the first time. However, this unknown factor is the same for all participants, and it remains to be seen what, if any, effects the new human annotation scheme has on the results. Another unknown is how this year’s test dataset, including the addition of new language pairs without human-annotated data, differs from those of previous years. Again, all participants are in the same boat here. Finally, our system was developed with rather limited resources, and we tested only on two language pairs with ESA annotations. Further studies are needed to show whether our one-size-fits-all system really does fit all.

	Pterodactyl	Austriadactylus
Description	Primary system, optimized for subtasks 2 & 3	Variant, optimized for subtasks 2 & 3
Foundational LLM	OpenAI gpt-5.4	OpenAI gpt-5.4
Reasoning effort	None	None
Language-specific exceptions	None	OpenAI gpt-5.6-terra, low reasoning, for Ladin, Ligurian, Northern Sámi
Computational cost	Total input 68.2	68.4
[million tokens]	Total output 10.7	11.5

Table 1: Systems submitted by 1020 Group to WMT26 subtasks 2 and 3, along with the computational cost involved in scoring the official WMT26 test data.

	Wukongopterus	Austriadraco
Description	Variant, optimized for subtask 3	Variant, differently optimized for subtask 3
Foundational LLM	OpenAI gpt-5.4	OpenAI gpt-5.4
Reasoning effort	None	None
Language-specific exceptions	None	None

Table 2: Systems submitted by 1020 Group to WMT26 subtask 3. No additional LLM-computational cost is involved, since the scores are taken from the main Pterodactyl scoring run and subsequently converted into a binary classification for subtask 3 via Python.

Acknowledgments

If we score higher, it is because we have the privilege of choosing which giants’ shoulders we stand on. We are deeply grateful to those who made that choice possible. We also thank Prof. Alon Lavie for encouraging us to participate as first-time entrants.

References

- Tom Fawcett. 2006. [An introduction to ROC analysis](#). *Pattern Recognition Letters*, 27(8):861–874.
- Markus Freitag, Nitika Mathur, Daniel Deutsch, Chi-Kiu Lo, Eleftherios Avramidis, Ricardo Rei, Brian Thompson, Frederic Blain, Tom Kocmi, Jiayi Wang, David Ifeoluwa Adelani, Marianna Buchicchio, Chrysoula Zerva, and Alon Lavie. 2024. [Are LLMs breaking MT metrics? Results of the WMT24 metrics shared task](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 47–81, Miami, Florida, USA. Association for Computational Linguistics.
- Evelyn Y. Garland and Carola F. Berger. 2026. [ROC analysis for evaluating translation quality estimation systems](#). *Preprint*, arXiv:2605.24721.
- Nuno M. Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André F. T. Martins. 2023. [xCOMET: transparent machine translation evaluation through fine-grained error detection](#). *Preprint*, arXiv:2310.10482.
- Tom Kocmi, Vilém Zouhar, Eleftherios Avramidis, Roman Grundkiewicz, Marzena Karpinska, Maja Popović, Mrinmaya Sachan, and Mariya Shmatova. 2024. [Error span annotation: A balanced approach for human evaluation of machine translation](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 1440–1453, Miami, Florida, USA. Association for Computational Linguistics.
- Alon Lavie, Greg Hanneman, Sweta Agrawal, Diptesh Kanojia, Chi-Kiu Lo, Vilém Zouhar, Frederic Blain, Chrysoula Zerva, Eleftherios Avramidis, Sourabh Deoghare, Archchana Sindhuja, Jiayi Wang, David Ifeoluwa Adelani, Brian Thompson, Tom Kocmi, Markus Freitag, and Daniel Deutsch. 2025. [Findings of the WMT25 shared task on automated translation evaluation systems: Linguistic diversity is challenging and references still help](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 436–483, Suzhou, China. 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.

David Samuel, Vladislav Mikhailov, Erik Velldal, Lilja Øvrelid, Lucas Georges Gabriel Charpentier, Andrey Kutuzov, and Stephan Oepen. 2025a. [Small languages, big models: A study of continual training on languages of Norway](#). In *Proceedings of the Joint 25th Nordic Conference on Computational Linguistics and 11th Baltic Conference on Human Language Technologies (NoDaLiDa/Baltic-HLT 2025)*, pages 573–608, Tallinn, Estonia. University of Tartu Library.

David Samuel, Lilja Øvrelid, Erik Velldal, and Andrey Kutuzov. 2025b. [Fluent alignment with disfluent judges: Post-training for lower-resource languages](#). *Preprint*, arXiv:2512.08777.

Chrysoula Zerva, Frederic Blain, José G. C. De Souza, Diptesh Kanojia, Sourabh Deoghare, Nuno M. Guerreiro, Giuseppe Attanasio, Ricardo Rei, Constantin Orasan, Matteo Negri, Marco Turchi, Rajen Chatterjee, Pushpak Bhattacharyya, Markus Freitag, and André Martins. 2024. [Findings of the quality estimation shared task at WMT 2024: Are LLMs closing the gap in QE?](#) In *Proceedings of the Ninth Conference on Machine Translation*, pages 82–109, Miami, Florida, USA. Association for Computational Linguistics.

Vilém Zouhar, Roman Grundkiewicz, Sara Rajaei, Parker Riley, Martin Popel, Rachel Bawden, Philipp Koehn, Marine Carpuat, and Tom Kocmi. 2026. [Contrastive ESA: Human evaluation of multiple translations at once](#). *Preprint*, arXiv:2607.26640.

Vilém Zouhar, Tom Kocmi, and Mrinmaya Sachan. 2024. [AI-assisted human evaluation of machine translation](#). *Preprint*, arXiv:2406.12419.