

Studying quantization trade-offs for efficient inference deployment in machine translation

Jim Zhao^{1,2*}, Sohir Maskey², Koen Oostermeijer², Douglas Orr³, Teryn Jones²

¹ University of Basel, ² Aleph Alpha Research, ³ Graphcore,

Correspondence: jim.zhao@unibas.ch

Abstract

Deploying large language models in realistic server environments poses challenges, as the system needs to provide high-quality responses with low latency. Quantization is a common approach to reduce the memory footprint and improve inference efficiency, yet its impact on latency and throughput is rarely evaluated under controlled, orchestration-level workloads. In this work we study the quantization trade-offs of EuroLLM (Martins et al., 2025) across three model sizes ranging from 1.7B to 22B for efficient deployment on a single A100 or H100 GPU. We demonstrate that combining a document-chunking strategy with W4A8 or W8A8 quantization improves the latency-throughput Pareto-curve under a wide range of workloads. Furthermore, since standard machine translation (MT) benchmarks rely on isolated sentences and fail to capture long-context dynamics, we introduce a document-level evaluation based on DocHPLT (O’Brien et al., 2025) to assess how text chunking strategies affect translation quality under quantization. Our results indicate that standard segment-level evaluation can potentially underestimate the interaction between quantization and long-context document translation, for some quantization formats, translation direction and models. Overall, our experiments show that the trade-off between inference efficiency and translation quality depends not only on the quantization format, but also on the choice of text chunking strategy.

1 Introduction

Large language models have revolutionized the field of machine learning and beyond, achieving close to or beyond human expert level across a wide range of domains, including coding (Coignion et al., 2024), medicine (Liu et al., 2025), and machine translation (Zheng et al., 2026). Deploying these models efficiently remains challenging.

Real-world serving systems must handle concurrent requests while balancing throughput and latency (Agrawal et al., 2024). These constraints are particularly relevant when models are deployed on a limited number of GPUs. Quantization is a common approach to reduce model memory requirements and can substantially improve inference efficiency (Jarmusch and Chandrasekaran, 2026).

Most prior work studies the quality retention and inference benefits of quantization formats in isolation, primarily focusing on memory or token-level metrics. In contrast, we analyze quantization trade-offs for MT inference under document-level serving constraints, focusing on end-to-end latency and throughput. We study these trade-offs on the EuroLLM (Martins et al., 2025) model family across three model sizes ranging from 1.7B to 22B parameters for deployment on a single A100 or H100 GPU, leading to the following contributions:

- We show via controlled closed-loop synthetic online benchmarks that combining W8A8 and W4A8 with a document-chunking strategy of 200-400 tokens provides the best efficiency gains for medium-sized documents across most of the considered deployment settings.
- Furthermore, we show that segment-level MT evaluation can potentially underestimate the quality degradation under quantization for certain translation directions and quantization formats.
- Concretely, for the Ukrainian → English translation direction we observe a stronger degradation of W4A16 with growing chunk size up to 1500 input prompt tokens for EuroLLM-1.7B and up to 4K input prompt tokens for EuroLLM-9B.

The remaining paper is structured as follows: We review related work in Section 2. In Section 3 we provide details about methods used in this work. In Section 4 we present our experimental results,

*Work done during internship at Aleph Alpha Research.

which we further discuss in Section 5. We discuss limitations in Section 6 and finally conclude in Section 7.

2 Related work

2.1 Quantization Paradigms

Quantization can improve inference efficiency by reducing the precision of model weights, activations, or the KV cache. Weight quantization reduces the model size and provides speed-ups under low batch sizes, where inference is often memory-bandwidth bound. As batch size increases, arithmetic intensity grows and inference can become compute-bound, reducing the relative benefit of weight-only quantization alone. Combining weight quantization with activation quantization overcomes this limitation by leveraging hardware-native low-precision compute formats, such as INT8, FP8, MXFP4 (Rouhani et al., 2023), NVFP4 (Xin et al., 2026), which enables arithmetic acceleration. Finally, KV quantization (Hooper et al., 2024) compresses the runtime memory footprint of the KV cache. This compression reduces memory bottlenecks under long sequences or large batch sizes, which might require CPU memory offloading.

Post-training quantization (PTQ) PTQ refers to methods, which quantize models after they have been fully trained. They typically use a small calibration dataset, although data-free approaches also exist (Nagel et al., 2019). Popular PTQ methods include GPTQ (Frantar et al., 2022), AWQ (Lin et al., 2025) or QQQ (Zhang et al., 2024b).

2.2 Efficient LLM Serving and Chunking

There is a rich line of work focused on optimizing the deployment and serving of large language models. At the scheduling level, Yu et al. (2022) introduced iteration-level scheduling, which is widely referred to as continuous batching, to efficiently process incoming requests of varying lengths. To address memory bottlenecks, Kwon et al. (2023) proposed PagedAttention to mitigate KV-cache fragmentation. These techniques are now standard in production-grade serving engines, such as vLLM (Kwon et al., 2023) and vendor specific frameworks like TensorRT-LLM (NVIDIA Corporation, 2023), which provide hardware-accelerated inference. Beyond engine-level memory management, inference efficiency can also be controlled at the input level through document chunking. While underexplored

in MT, chunking strategies are heavily researched in Retrieval-Augmented Generation (RAG). For instance, Qu et al. (2025) challenged the benefits of semantic chunking over more efficient fixed-size chunking in RAGs.

2.3 Document-level Machine Translation

While widely used MT benchmarks, such as WMT24++ (Deutsch et al., 2025) or Flores+ (Goyal et al., 2022) primarily evaluate translations at the segment level, document-level MT has a substantial research history. Läubli et al. (2018) demonstrate that evaluation protocols based on isolated sentences can lead to different conclusions from document-level evaluations when comparing professional and machine translations. Lopes et al. (2020) compare several document-level context-aware neural MT approaches, and Maruf et al. (2021) surveys methods and evaluation protocols for document-level MT. More recently, Wang et al. (2023b) evaluate GPT-3.5 (Ouyang et al., 2022) and GPT-4 (Achiam et al., 2023) for document-level translation and find that they outperform commercial MT systems under human evaluation, highlighting the potential of decoder-only LLMs for this setting. Similarly, Karpinska and Iyyer (2023) find that translating complete literary paragraphs with GPT-3.5 (Ouyang et al., 2022) yields higher-quality translations than translating each sentence independently across 18 language pairs. However, critical errors, including occasional content omissions, remain even when the model is given document-level context.

2.4 Evaluation of Quantized LLMs

Li et al. (2024) evaluate quantized LLMs over a wide range of tasks, including long-context tasks, such as key-value retrieval and multi-document question answering. Marchisio et al. (2024) study how quantization affects multilingual LLMs, showing that, in line with our finding, automatic benchmarks often underestimate the quantization effects and that languages are disparately affected by quantization. Finally, in the context of MT, Marie and Fujita (2025) study how different PTQ methods affect translation quality across languages. They do not examine how quantization interacts with document context or document-chunking strategies.

Our work connects these lines of research by jointly studying quantization, document

chunking, serving efficiency, and document-level translation quality.

3 Methods

3.1 Machine Translation Models

We consider the EuroLLM (Martins et al., 2025) model family across three model sizes ranging from 1.7B to 22B. EuroLLM utilizes a standard dense decoder-only Transformer architecture. Concretely, we use the instruction-tuned variants EuroLLM-1.7B-Instruct, EuroLLM-9B-Instruct-2512, and EuroLLM-22B-2512, where the first has a context length of 4096 and the latter two have a context length of 32K. All models use RoPE (Su et al., 2024) with $\theta = 10.000$ for EuroLLM-1.7B-Instruct and $\theta = 1.000.000$ for EuroLLM-9B-Instruct-2512 and EuroLLM-22B-2512. For brevity, we will refer to the EuroLLM models simply by their model size, e.g. EuroLLM-1.7B.

3.2 Hardware and Format Constraints

In this work, we consider three combinations of weight and activation quantization: W8A8 (8-bit weight and activation), W4A8 (4-bit weight and 8-bit activation), and W4A16 (4-bit weight-only). Further details on the concrete quantization scheme can be found in Appendix A.2. While there is an ongoing debate regarding whether floating-point or integer formats should be preferred for LLMs (Zhang et al., 2024a), we observe that W8A8 often remains close to the translation quality of the unquantized BF16 baseline in most cases while benefiting from widely supported INT8 tensor cores. This allows the deployment of the quantized model across a broad range of GPUs (including the RTX4090, A100, and H100) using highly optimized compute kernels.

Furthermore, while more sophisticated 4-bit formats like MXFP4 (Rouhani et al., 2023) and NVFP4 (Abecassis et al., 2025) show promise, they are only natively supported on newer Blackwell architectures (Jarmusch and Chandrasekaran, 2026). Finally, extreme low bit-widths (<4 bits) (Ma et al., 2024; Maskey et al., 2026; Wang et al., 2023a; Wu et al., 2025) are excluded from this study, as they are primarily optimized for highly memory-constrained, on-device deployment rather than achieving maximum throughput across the wide range of concurrent batch sizes expected in a server setting.

3.3 Post-Training Quantization

We perform post-training quantization (PTQ) using GPTQ (Frantar et al., 2022) and apply SmoothQuant (Xiao et al., 2023) before performing GPTQ to reduce activation outliers before activation quantization for W8A8 and W4A8, where we choose a smooth strength of 0.8 for W8A8 and 0.4 for W4A8 quantization. For W4A16 GPTQ is applied directly. We performed GPTQ on a subset of 1024 samples subsampled from OPUS (Tiedemann, 2016), where the target translations were generated by Seed-X (Cheng et al., 2025). All linear layers in the model, except for the embedding and head layer, are quantized using the same data format.

3.4 vLLM Engine

We use vLLM (Kwon et al., 2023) as our serving engine, which provides state-of-the-art kernel implementations, including W4A16, W4A8 and W8A8 kernels. Furthermore, it implements additional inference optimizations, such as PagedAttention (Kwon et al., 2023) and Continuous Batching (Yu et al., 2022). We can thus evaluate the benefit of quantized models for inference under near-production conditions, without inefficient kernel implementations potentially distorting the picture. We note that vLLM does not support W4A8 kernels on A100 by default, which is why we limit experiments using W4A8 on the H100 GPU. We used vLLM version 0.26.0 for our experiments.

3.5 Translation evaluation

To study the translation quality retention, we benchmarked the EuroLLM-1.7B and EuroLLM-9B model on BF16, W4A16 and W8A8 on a subset of DocHPLT (O’Brien et al., 2025) on a RTX4090 GPU. We select EN $\leftrightarrow$ UK as an exemplary translation pair supported in EuroLLM and covered by DocHPLT, leading to two translation directions in total.

3.5.1 Segment-level translation

For segment-level evaluation, we use chrF++ (Popović, 2017) to compare the translation quality retained by the quantized model. Segments are translated independently, then translations and references are concatenated by document before computing chrF++, which scores translations in the range [0, 100], where a higher score is better.

3.5.2 Document-level evaluation

We construct document-level inputs from high-confidence aligned segments in DocHPLT. DocHPLT (O’Brien et al., 2025) contains parallel documents that are split into partially aligned source and reference alignment records. To construct document-level inputs without breaking this alignment, we first filter documents by alignment records with $\text{Aligner-Score} \geq 0.5$ and $\text{Bicleaner-Score} \geq 0.8$, merge overlapping records into non-overlapping bilingual blocks, and preserve their original document order before chunking. This ensures that the blocks in each documents are sufficiently well aligned. We then greedily concatenate complete blocks in the document order until the tokenized prompt reaches or exceeds a target prompt-length threshold T . As blocks are never split to avoid breaking existing alignment, the realized prompt length may exceed T . In addition, short trailing chunks are also retained when a document produces multiple chunks to obtain a full document translation. We exclude a document at threshold T only when the complete document forms a single prompt shorter than $0.8T$, in order to prevent short documents from dominating estimates at larger thresholds. To control for changes in document composition, we evaluate all thresholds, models, and quantization formats on the same subset of documents eligible at the largest threshold. Each chunk is passed individually to the model and the generated outputs are concatenated in their original order to reconstruct the document translation. We compute chrF++ for each reconstructed document separately and report the unweighted mean across all eligible documents. The main reason for relying on this statistical measure is that neural metrics, such as xCOMET (Guerreiro et al., 2024), have a limited context window of 512 tokens for the combined source, reference, and translation text. Further details on the chunking and evaluation, including decoding parameters and prompt templates can be found in Appendix A.3.

3.6 Inference Efficiency

The inference efficiency benchmarks are split into two parts. We perform offline benchmarking in the first stage to evaluate the speed-ups from using hardware-accelerated GEMM kernels and reduced memory traffic in isolation. In a second stage, we study optimal chunking strategies of documents in a closed-loop online inference setting for different

number of concurrent users on medium-sized documents. In general, chunking documents into shorter parts allows parallel processing, but increases the total number of requests, which vLLM needs to schedule. By choosing the optimal chunk length, the end-to-end latency of the full document and the throughput of documents can be improved. However, sentences might lose context if split into too short chunks, leading to degraded translations. We evaluate the trade-offs of how the chunk size affects translation quality across different languages in the final subsection. We benchmark the models using randomly generated token IDs for both inference evaluations as input prompts and disable prefix caching to ensure comparability across model sizes and different quantization formats. We report the results in terms of generated tokens.

3.6.1 Offline Benchmarking

We first benchmark the inference efficiency of the quantization formats via offline benchmarking. While vLLM can handle requests of different lengths efficiently via continuous batching, we measure the inference statistics for fixed input and output sequence lengths and enforce the number of output tokens to the number of input tokens by ignoring the $\text{eos}_{\text{token}}$ for simplicity. This holds approximately true for MT tasks, but also highly depends on the language pair and the tokenizer used (Martins et al., 2025). We measure the wall time and output token throughput for a range of different sequence lengths and batch sizes up to 1024.

The roofline model is a useful tool to upper bound the benefits of different quantization formats for inference efficiency by putting the compute capacity and memory-bandwidth into relation. A full derivation of the roofline model and upper bound estimates on speed-ups can be found in Appendix A.1. During low-batch decoding, inference is generally dominated by streaming model weights from memory, such that reducing weight precision can provide substantial benefits. As the batch size increases, weights are reused across more tokens and arithmetic intensity increases, reducing the benefit of weight-only quantization as projection and MLP operations become compute-bound. In this regime, activation quantization becomes important because it enables hardware-accelerated low-precision computation. Long-context decoding introduces an additional limitation: attention and KV-cache traffic grow with context length but

are unaffected by weight quantization, causing the attainable end-to-end speed-up to decrease. These predictions represent idealized upper bounds and do not account for dynamic quantization overhead, kernel utilization, or scheduling and KV-cache capacity effects during online serving.

3.6.2 Closed-Loop Online Document-Level Benchmark

To benchmark the inference efficiency of quantized models more realistically, we perform a controlled closed-loop synthetic online benchmark, which consists of an additional orchestration layer to incorporate chunking strategies of full documents and to measure E2E latencies on the document level. We simulate document lengths corresponding to 1200 tokens under a growing number of concurrent users of up to 500 users and different chunk sizes ranging from 50 to 400 tokens.

4 Results

4.1 Offline Benchmarking

The results of the offline benchmarks can be found in Figure 1 with further results in Appendix A.5. Across the three size regimes we observe that quantization formats have marginal benefits for small model sizes $< 2\text{B}$ parameters. For medium-sized models of around 9B, W4A8 and W8A8 show efficiency gains across all batch sizes, while W4A16 becomes as fast as BF16 at large batch sizes as the model becomes compute-bound. Finally for large MT models of 22B, we observe a difference between the H100 and A100 due to the higher compute-to-memory-bandwidth ratio on the H100. While H100 has a ridge point at $\approx I_{\text{ridge}}^{\text{H100}} = 295 \text{ FLOP/byte}$, A100 has a ridge point at $\approx I_{\text{ridge}}^{\text{A100}} = 153 \text{ FLOP/byte}$ (see Appendix A.1 for more details). Thus, W4A16 remains competitive on the H100, while it saturates back to the BF16 baseline on the A100 at the highest batch sizes considered. We observe that for larger batch sizes, throughput actually can reduce, as is the case for the 22B model on A100. This observation is not predicted by a simple roofline model.

4.2 Closed-Loop Online Document-Level Benchmark

Based on the offline benchmark and quantization results, we limit the comparison to BF16, W8A8, and W4A8. The results can be found in Figure 2.

We make the following observations based on the results: Quantization is not always beneficial. In particular for the 1.7B model deployed on a H100, BF16 achieves higher throughput at lower latencies than W8A8 quantization. The reason for this is that the W4A8 and W8A8 schemes add overhead by applying dynamic activation quantization. In addition, the online benchmark further shows that document chunking can provide additional efficiency gains. The Pareto-optimal chunking strategy depends on the specific deployment setting and workload. At low concurrency, it is always preferable to use the smallest acceptable chunk length to allow for parallel processing, while for higher concurrency the Pareto-optimal chunk length varies mostly between 200 and 400 tokens. At the highest considered concurrency on the 22B model under BF16 precision, a trade-off between higher throughput and lower latency is also observable. While shorter chunk length allow for a better KV cache management and lead to a higher throughput, the tail latency also increases considerably. The quantized variants under the same workload do not show this behavior. We speculate that this is because they have more memory headroom by saving 20-30 GB of model memory. Apart from the 1.7B model deployed on a H100 GPU all other configurations provide efficiency gains by combining W4A8 or W8A8 quantization with a chunking strategy. While chunking of documents improves the inference efficiency, it might affect the translation quality by losing relevant context for correct translations. We discuss the effect of chunking strategies on translation quality in the next subsection.

4.3 Effect of Chunking Strategy on Translation Quality

Unlike the synthetic random-token workloads used in the previous section to cleanly isolate scheduling performance, evaluating translation quality requires real-world data. In this section we evaluate MT quality on document-level translations. Our results reveal that standard segment-level evaluation can underestimate the interaction between quantization and long-context document translation for certain quantization formats, translation directions and models.

The results for the UK $\rightarrow$ EN language pair can be found in Figure 3 with further results on EN $\rightarrow$ UK in Appendix A.4. A corresponding WMT24++ analysis, which can be found in Appendix A.4 shows a weaker and less consistent

Table 1: DocHPLT scores of EuroLLM-1.7B and EuroLLM-9B across different quantization formats in terms of chrF++. Target thresh refers to the target prompt-length threshold. Mean scores of chrF++ computed across subset of $N = 27$ documents of DocHPLT. Brackets show bootstrapped 0.95 confidence intervals. For more details, please refer to Section 4.3.

Model	Format	Evaluation (chrF++ $\uparrow$)		
		Indep-segment translation (UK $\rightarrow$ EN)	Document-level eval Target thresh=100 tok	Document-level eval Target thresh=1500 tok
EuroLLM-1.7B	BF16	65.7, CI[58.2, 70.6]	67.6, CI[59.2, 72.6]	35.6, CI[28.7, 41.1]
	W8A8	58.6, CI[51.9, 64.3]	65.3, CI [57.2, 70.3]	30.7, CI [24.5, 36.0]
	W4A16	63.6, CI[56.5, 69.0]	66.2, CI[58.7, 70.6]	6.3, CI[4.9, 8.5]
			Target thresh=100 tok	Target thresh=4000 tok
EuroLLM-9B	BF16	65.2, CI[60.3, 68.8]	69.6, CI[65.1, 72.4]	65.23, CI[60.0, 70.12]
	W8A8	63.4, CI[58.4, 67.1]	68.3, CI[63.4, 71.2]	60.22, CI[52.2, 66.2]
	W4A16	61.3, CI[56.2, 65.4]	64.8, CI[60.5, 68.0]	45.6, CI[36.8, 54.4]




Figure 1: Offline benchmark measuring average output throughput in tok/s against average wall time for a sequence length of 1000 tokens across three sizes regimes on a A100 (top row) and a H100 (bottom row). Further results on other sequence lengths can be found in Appendix A.5.

effect up to a token threshold of 1000. The figures show that even though EuroLLM-1.7B has a context length of 4096, it translates longer inputs considerably worse, even at full BF16 precision. This observation is considerably weaker for the BF16 checkpoint of EuroLLM-9B, which was explicitly trained for long contexts up to 32K. In addition, our evaluation shows that standard segment-level MT benchmarks, such as WMT (Deutsch et al., 2025) or Flores+ (Goyal et al., 2022) can potentially underestimate the degradation of quantized MT models under long context for certain translation directions and quantization formats. This is

highlighted in Table 1 as well as the right side of Figure 3, in which the W4A16 quantized checkpoints of EuroLLM-1.7B and EuroLLM-9B suggest a lower degradation on segment-level translations than what is observed if longer chunks are being translated directly. We note however that this effect is considerably weaker for W8A8 up to the considered sequence length as well as for the reverse translation direction EN $\rightarrow$ UK, as can be seen in Figure 5.




Figure 2: Closed-loop online benchmark for user concurrency from 1 to 500 users measuring throughput in documents/s against p99 full-Documnet latency for different chunking strategies on a document length of 1200 tokens across three model sizes on a A100 (top row) and a H100 (bottom row). Note that we did not benchmark W4A8 on the H100 for the 1.7B model size.




Figure 3: Chunk length vs. chrF++ Score for EuroLLM-1.7B (top row) and EuroLLM-9B (bottom row) across different quantization formats averaged across a subset of 27 documents from DocHPLT. Subplots on the right corresponds to paired per-document BF16 - quantized differences, where positive values mean degradation relative to BF16. Shaded area corresponds to bootstrapped 95% CI of document-level means. Note the different y limits for better visibility. Every point corresponds to individual document scores. Additional results on other language pairs be found in Appendix A.4.

5 Discussion

The closed-loop online benchmarks have shown that document chunking can improve model serving efficiency. Across model sizes, mainly medium-

sized and larger MT models benefit from quantization, while small models $< 2B$ only show marginal benefits or even slight slow down compared to the BF16 baseline due to the overhead introduced by

dynamic activation quantization. The efficiency gains across the full range of workloads arise from a combination of higher compute capacity by making use of INT8 tensor cores and reducing the model size, where the latter is beneficial both at low workloads when inference is memory bandwidth-bound as well as during capacity-bound serving under high user concurrency where memory savings benefit KV cache management and scheduling efficiency during online inference.

Furthermore, our experiments on how chunk length affects translation quality have shown that standard segment-level MT benchmarks can potentially underestimate the degradation for longer context translation for certain translation directions and quantization formats, which needs to be carefully evaluated case-by-case. While in some cases the quantized checkpoint follows the BF16 baseline closely, it can lead to a measurable degradation in MT quality in other cases, such as the W4A16 quantized checkpoints of EuroLLM-1.7B and EuroLLM-9B evaluated on UK $\rightarrow$ EN. We hypothesize that the different context length ranging from 4K to 32K of the models considered as well as the data that the models have been trained on play a major role. For instance, EuroLLM-1.7B is instruction-tuned only on segment-level translation pairs (Martins et al., 2025). In addition, the sentence-level calibration data of OPUS (Tiedemann, 2016) could have biased the quantized model checkpoints toward shorter translation thresholds. Although identifying the mechanistic reason underlying the sensitivity induced by quantization is still open and out of scope for this work, we hypothesize that a short Quantization-aware training (QAT) or distillation run with long context translation samples can help to close this gap. If the Pareto-optimal chunk length in terms of efficiency is longer than for translation quality, one could resolve this discrepancy by simulating shorter chunks using appropriate block-diagonal attention masks (though this adds additional overhead), but in the other case it results in a trade-off between achieving higher inference efficiency and better translation quality. We note however that we do not construct a joint quality–efficiency Pareto frontier in this work, since translation quality is evaluated only for a subset of the models and quantization formats considered in the inference benchmarks.

We summarize key takeaways of our work below:

1. **Quantization to W8A8 or W4A8 can improve the latency-throughput tradeoff, depending on hardware and model size.** Among the configurations supported in our vLLM setup, for models of 9B parameters and above, W8A8 on the A100 and W4A8 on the H100 can provide improved throughput at similar latency. At higher batch sizes, activation quantization is particularly important on the A100, whose lower compute-to-memory-bandwidth ratio causes inference to become compute-bound sooner.
2. **Quantization and document chunking should be optimized jointly for inference efficiency.** Across most evaluated deployment settings, combining W8A8 or W4A8 with chunk lengths of approximately 200–400 tokens improves the latency-throughput Pareto frontier for medium-length documents.
3. **Sentence-level MT benchmarks can underestimate long-context degradation.** Evaluations based on isolated segments can underestimate how translation quality changes with longer context and can therefore conceal degradation that occurs during document-level inference, in particular when deploying quantized models. This strongly depends on the considered quantization format and translation direction.

6 Limitations and Future Work

In this work, we only considered GPTQ (Frantar et al., 2022) combined with (Xiao et al., 2023) as a quantization method and it is possible that more sophisticated PTQ methods or QAT do not show degradation for longer contexts observed on the quantized EuroLLM models. We also did not control for the calibration data used for PTQ, and it remains to be verified whether longer context calibration data can reduce the degradation. We note that we apply dynamic activation quantization, which added an overhead in the inference benchmark, which could potentially be removed. While static activation quantization can help, it risks to potentially degrade the translation quality. Further, we point out that W4A8 was not tested on A100 due to missing vLLM support. It thus remains an open question, whether W4A8 will be preferable over W8A8 once corresponding software-support exists. Furthermore, we only evaluated how trans-

lation quality changes with chunk lengths on two translation directions, so they might be different for low-resource languages.

Our quality evaluation was also performed only across a subset of 27 UK $\rightarrow$ EN, respectively 13 EN $\rightarrow$ UK documents of DocHPLT, which needs to be extended to more documents in future work. While we ran quality experiments up to a threshold of 4K tokens, it is still not close to the 32K context length of EuroLLM-9B, so we were not able to characterize the behavior of EuroLLM-9B close to its context limit. Finally, in this work we focused only on weight and activation quantization. We leave the consideration of KV cache quantization for inference and quality retention for future work.

7 Conclusion

In this work, we have studied quantization trade-offs for efficient inference deployment for machine translation on the EuroLLM model family on a A100 and a H100 GPU. By evaluating the inference efficiency across three models, spanning 1.7B to 22B under closed-loop synthetic online benchmarks on vLLM, we find that for model sizes of 9B and above W8A8 generally provides the strongest Pareto performance in the evaluated configurations on the A100, while W4A8 provides the strongest Pareto performance on the H100. For small models < 2B W8A8 leads to marginal efficiency gains or even slight slowdown. In our synthetic document-serving benchmark we show that chunking medium-sized documents into chunks of 200-400 tokens is Pareto-optimal for inference efficiency. Finally, we observe that translation quality is better for shorter thresholds, while the robustness for longer sequences differs for model sizes, quantization format, and translation direction.

References

- Felix Abecassis, Anjulie Agrusa, Dong Ahn, Jonah Alben, Stefania Alborghetti, Michael Andersch, Sivakumar Arayandi, Alexis Bjorlin, Aaron Blakeman, Evan Briones, et al. 2025. Pretraining large language models with nvfp4. *arXiv preprint arXiv:2509.25149*.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Amey Agrawal, Nitin Kedia, Ashish Panwar, Jayashree Mohan, Nipun Kwatra, Bhargav Gulavani, Alexey Tumanov, and Ramachandran Ramjee. 2024. Taming {Throughput-Latency} tradeoff in {LLM} inference with {Sarathi-Serve}. In *18th USENIX symposium on operating systems design and implementation (OSDI 24)*, pages 117–134.
- Shanbo Cheng, Yu Bao, Qian Cao, Luyang Huang, Liyan Kang, Zhicheng Liu, Yu Lu, Wenhao Zhu, Jingwen Chen, Zhichao Huang, et al. 2025. Seed-x: Building strong multilingual translation llm with 7b parameters. *arXiv preprint arXiv:2507.13618*.
- Tristan Coignion, Clément Quinton, and Romain Rouvoy. 2024. A performance study of llm-generated code on leetcode. In *Proceedings of the 28th international conference on evaluation and assessment in software engineering*, pages 79–89.
- Daniel Deutsch, Eleftheria Briakou, Isaac Rayburn Caswell, Mara Finkelstein, Rebecca Galor, Juraj Juraska, Geza Kovacs, Alison Lui, Ricardo Rei, Jason Riesa, et al. 2025. Wmt24++: Expanding the language coverage of wmt24 to 55 languages & dialects. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 12257–12284.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. 2022. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. The Flores-101 evaluation benchmark for low-resource and multilingual machine translation. *Transactions of the Association for Computational Linguistics*, 10.
- Nuno M Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André FT Martins. 2024. xcomet: Transparent machine translation evaluation through fine-grained error detection. *Transactions of the Association for Computational Linguistics*, 12:979–995.
- Coleman Hooper, Sehoon Kim, Hiva Mohammadzadeh, Michael W Mahoney, Yakun S Shao, Kurt Keutzer, and Amir Gholami. 2024. Kvquant: Towards 10 million context length llm inference with kv cache

- quantization. *Advances in Neural Information Processing Systems*, 37:1270–1303.
- Aaron Jarmusch and Sunita Chandrasekaran. 2026. Microbenchmarking nvidia’s blackwell architecture: An in-depth architectural analysis. In *2026 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 1026–1036. IEEE.
- Marzena Karpinska and Mohit Iyyer. 2023. [Large language models effectively leverage document-level context for literary translation, but critical errors persist](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 419–451, Singapore. Association for Computational Linguistics.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pages 611–626.
- Samuel Läubli, Rico Sennrich, and Martin Volk. 2018. [Has machine translation achieved human parity? a case for document-level evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4791–4796, Brussels, Belgium. Association for Computational Linguistics.
- Shiyao Li, Xuefei Ning, Luning Wang, Tengxuan Liu, Xiangsheng Shi, Shengen Yan, Guohao Dai, Huazhong Yang, and Yu Wang. 2024. Evaluating quantized large language models. *arXiv preprint arXiv:2402.18158*.
- Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Guangxuan Xiao, and Song Han. 2025. Awq: Activation-aware weight quantization for on-device llm compression and acceleration. *GetMobile: Mobile Computing and Communications*, 28(4):12–17.
- Xiaohong Liu, Hao Liu, Guoxing Yang, Zeyu Jiang, Shuguang Cui, Zhaoze Zhang, Huan Wang, Liyuan Tao, Yongchang Sun, Zhu Song, et al. 2025. A generalist medical language model for disease diagnosis assistance. *Nature medicine*, 31(3):932–942.
- António Lopes, M. Amin Farajian, Rachel Bawden, Michael Zhang, and André F. T. Martins. 2020. [Document-level neural MT: A systematic comparison](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 225–234, Lisboa, Portugal. European Association for Machine Translation.
- Shuming Ma, Hongyu Wang, Lingxiao Ma, Lei Wang, Wenhui Wang, Shaohan Huang, Li Dong, Ruiping Wang, Jilong Xue, and Furu Wei. 2024. The era of 1-bit llms: All large language models are in 1.58 bits. *arXiv preprint arXiv:2402.17764*.
- Kelly Marchisio, Saurabh Dash, Hongyu Chen, Dennis Aumiller, Ahmet Üstün, Sara Hooker, and Sebastian Ruder. 2024. How does quantization affect multilingual llms? In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15928–15947.
- Benjamin Marie and Atsushi Fujita. 2025. The uneven impact of post-training quantization in machine translation. *arXiv preprint arXiv:2508.20893*.
- Pedro Henrique Martins, Patrick Fernandes, João Alves, Nuno M Guerreiro, Ricardo Rei, Duarte M Alves, José Pombal, Amin Farajian, Manuel Faysse, Mateusz Klimaszewski, et al. 2025. Eurollm: Multilingual language models for europe. *Procedia Computer Science*, 255:53–62.
- Sameen Maruf, Fahimeh Saleh, and Gholamreza Haffari. 2021. A survey on document-level neural machine translation: Methods and evaluation. *ACM Computing Surveys (CSUR)*, 54(2):1–36.
- Sohir Maskey, Constantin Eichenberg, Johannes Messner, and Douglas Orr. 2026. 1-bit wonder: Improving qat performance in the low-bit regime through k-means quantization. *arXiv preprint arXiv:2602.15563*.
- Markus Nagel, Mart van Baalen, Tijmen Blankevoort, and Max Welling. 2019. Data-free quantization through weight equalization and bias correction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1325–1334.
- NVIDIA Corporation. 2023. TensorRT-LLM: An open-source library for optimizing LLM inference. <https://github.com/NVIDIA/TensorRT-LLM>. Accessed: 2026-07-25.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Dayyán O’Brien, Bhavitvya Malik, Ona de Gibert, Pinzhen Chen, Barry Haddow, and Jörg Tiedemann. 2025. Dochplt: A massively multilingual document-level translation dataset. In *Proceedings of the Tenth Conference on Machine Translation*, pages 286–300.
- Maja Popović. 2017. chrF++: words helping character n-grams. In *Proceedings of the second conference on machine translation*, pages 612–618.
- Renyi Qu, Ruixuan Tu, and Forrest Bao. 2025. Is semantic chunking worth the computational cost? In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 2155–2177.
- Bitá Darvish Rouhani, Ritchie Zhao, Ankit More, Mathew Hall, Alireza Khodamoradi, Summer Deng, Dhruv Choudhary, Marius Cornea, Eric

- Dellinger, Kristof Denolf, et al. 2023. Microscaling data formats for deep learning. *arXiv preprint arXiv:2310.10537*.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063.
- Jörg Tiedemann. 2016. Opus-parallel corpora for everyone. In *Proceedings of the 19th Annual Conference of the European Association for Machine Translation: Projects/Products*.
- Hongyu Wang, Shuming Ma, Li Dong, Shaohan Huang, Huaijie Wang, Lingxiao Ma, Fan Yang, Ruiping Wang, Yi Wu, and Furu Wei. 2023a. Bitnet: Scaling 1-bit transformers for large language models. *arXiv preprint arXiv:2310.11453*.
- Longyue Wang, Chenyang Lyu, Tianbo Ji, Zhirui Zhang, Dian Yu, Shuming Shi, and Zhaopeng Tu. 2023b. Document-level machine translation with large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 16646–16661, Singapore. Association for Computational Linguistics.
- Xun Wu, Shaohan Huang, Wenhui Wang, Ting Song, Li Dong, Yan Xia, and Furu Wei. 2025. Bitnet distillation. *arXiv preprint arXiv:2510.13998*.
- Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. 2023. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International conference on machine learning*, pages 38087–38099. PMLR.
- Meng Xin, Sweta Priyadarshi, Jingyu Xin, Bilal Kartal, Aditya Vavre, Asma Kuriparambil Thekkumpate, Zijia Chen, Ameya Sunil Mahabaleshwarkar, Ido Shahaf, Akhiad Bercovich, et al. 2026. Quantization-aware distillation for nvfp4 inference accuracy recovery. *arXiv preprint arXiv:2601.20088*.
- Gyeong-In Yu, Joo Seong Jeong, Geon-Woo Kim, Soojeong Kim, and Byung-Gon Chun. 2022. Orca: A distributed serving system for {Transformer-Based} generative models. In *16th USENIX symposium on operating systems design and implementation (OSDI 22)*, pages 521–538.
- Yijia Zhang, Lingran Zhao, Shijie Cao, Sicheng Zhang, Wenqiang Wang, Ting Cao, Fan Yang, Mao Yang, Shanghang Zhang, and Ningyi Xu. 2024a. Integer or floating point? new outlooks for low-bit quantization on large language models. In *2024 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE.
- Ying Zhang, Peng Zhang, Mincong Huang, Jingyang Xiang, Yujie Wang, Chao Wang, Yineng Zhang, Lei Yu, Chuan Liu, and Wei Lin. 2024b. Qqq: Quality quattuor-bit quantization for large language models. *arXiv preprint arXiv:2406.09904*.
- Mao Zheng, Zheng Li, Tao Chen, Bo Lv, Mingrui Sun, Mingyang Song, Jinlong Song, Hong Huang, Decheng Wu, Hai Wang, et al. 2026. Hy-mt2: A family of fast, efficient and powerful multilingual translation models in the wild. *arXiv preprint arXiv:2605.22064*.

A Appendix

A.1 A concrete GEMM-only roofline model

We first consider an idealized model that includes only GEMM FLOPs and the movement of GEMM operands between HBM and the GPU. For

$$Y = XW, \quad X \in \mathbb{R}^{M \times K}, \quad W \in \mathbb{R}^{K \times N},$$

the number of FLOPs and transferred bytes are approximated by

$$\begin{aligned} F_{\text{gemm}} &= 2MKN, \\ Q_{\text{gemm}} &= b_a MK + b_w KN + b_a MN, \end{aligned}$$

where b_a and b_w denote the number of bytes per activation and weight, respectively. The arithmetic intensity is therefore

$$I_{\text{gemm}} = \frac{2MKN}{b_a MK + b_w KN + b_a MN}.$$

For a square hidden projection with $K = N = d$, this becomes

$$I_{\text{gemm}}(M; d) = \frac{2Md^2}{b_w d^2 + 2b_a Md}. \quad (1)$$

For $M \ll d$, weight traffic dominates and

$$I_{\text{gemm}}(M; d) \approx \frac{2M}{b_w}.$$

Consequently, with BF16 activations, the ideal arithmetic intensity is approximately M , $2M$, and $4M$ FLOP/byte for BF16, INT8, and INT4 weights, respectively.

We instantiate the analysis for the three EuroLLM models, see Table 2 for their architectural details.

The roofline model predicts the achievable GEMM performance, measured in FLOP/s, as

$$P_{\text{gemm}} = \min(P_{\text{peak}}, \text{BW}_{\text{HBM}} I_{\text{gemm}}),$$

where P_{peak} denotes the hardware’s theoretical peak compute throughput in FLOP/s and BW_{HBM} its peak HBM bandwidth in bytes/s. The product $\text{BW}_{\text{HBM}} I_{\text{gemm}}$ is the maximum compute throughput that can be sustained by the available memory bandwidth.

Using dense BF16 Tensor Core throughput, an H100 SXM has an approximate ridge point of

$$I_{\text{ridge}}^{\text{H100}} = \frac{989 \text{ TFLOP/s}}{3.35 \text{ TB/s}} \approx 295 \text{ FLOP/byte},$$

while an A100 80GB SXM has

$$I_{\text{ridge}}^{\text{A100}} = \frac{312 \text{ TFLOP/s}}{2.039 \text{ TB/s}} \approx 153 \text{ FLOP/byte}.$$

Here we assume that weight-only quantized kernels unpack and dequantize weights on-chip and use the BF16 Tensor Core compute ceiling. This is an optimistic upper bound.

With BF16 activations, the approximate number of GEMM rows (calculated via Equation (1)) required to reach the compute-bound regime are shown in Table 3.

During decoding, $M = B$, where B is the number of tokens processed together. At batch size one, all three EuroLLM models are therefore strongly memory-bandwidth bound. In this ideal model, INT8 and INT4 weights reduce weight traffic by factors of two and four and provide corresponding speedups.

As the batch size increases, the same weights are reused across more tokens. For example, the INT4 projection GEMMs become nominally compute-bound on H100 at batch sizes of approximately 104, 86, and 82 for EuroLLM-1.7B, EuroLLM-9B, and EuroLLM-22B, respectively. Beyond this point, further reducing weight traffic does not improve the idealized GEMM runtime.

During prefill, $M = BS$, where S is the prompt sequence length. Consequently, even at batch size one, the projection and MLP GEMMs become compute-bound after only a few hundred prompt tokens. The GEMM-only model therefore predicts that weight-only quantization is most useful for low-batch decoding and much less useful for long-sequence prefill.

These theoretical considerations also generalize directly to activation quantization, provided that the hardware and kernels support computation in the chosen activation format. The effect of activation precision can be modeled by substituting the corresponding number of activation bytes into b_a in Equation (1) and using the peak compute throughput of the corresponding hardware format. For example, if INT8 computation provides twice the peak throughput of BF16 while HBM bandwidth remains unchanged, the hardware ridge point doubles. Consequently, the GEMM remains memory-bound up to approximately twice the batch size, allowing further weight quantization to provide speedups over a correspondingly larger range of batch sizes. Reducing the activation size from two bytes to one

Table 2: Architectural details of EuroLLM models.

Model	L	d	d_{ff}	n_q/n_{kv}	d_{kv}	Parameters
EuroLLM-1.7B	24	2048	5632	16/8	1024	1.657B
EuroLLM-9B	42	4096	12288	32/8	1024	9.152B
EuroLLM-22B	54	6144	16384	48/8	1024	22.637B

Table 3: Approximate number of GEMM rows required to reach the compute-bound regime under BF16 activations, calculated via Equation (1).

Model	GPU	BF16 weights	INT8 weights	INT4 weights
EuroLLM-1.7B	H100	415	207	104
	A100	180	90	45
EuroLLM-9B	H100	345	172	86
	A100	165	83	41
EuroLLM-22B	H100	327	163	82
	A100	161	81	40

byte further decreases activation traffic, although this effect is small when weight traffic dominates.

Why decode can nevertheless benefit less than predicted. The GEMM-only model predicts no dependence of decode GEMMs on the existing context length. The projection and MLP GEMMs have $M = B$, irrespective of whether the context contains 1,000 or 30,000 tokens. Thus, under this model, decode latency would be constant with respect to context length. The observed context-length dependence comes primarily from attention and its KV-cache traffic.

For a model with L layers and KV dimension d_{kv} , reading the BF16 key and value cache during one decoding step requires approximately

$$Q_{\text{KV}} = 2BLSd_{\text{kv}}b_a,$$

where the factor two accounts for keys and values. Unlike weight traffic, this traffic grows linearly with both batch size and context length.

The corresponding attention FLOPs are approximately

$$F_{\text{attn}} \approx 4BLSd,$$

giving an idealized attention arithmetic intensity of

$$I_{\text{attn}} \approx \frac{2d}{b_a d_{\text{kv}}} = \frac{2g}{b_a}, \quad g = \frac{d}{d_{\text{kv}}}.$$

All three EuroLLM models have eight KV heads with head dimension 128, and therefore $d_{\text{kv}} = 1024$. Their GQA ratios and BF16 attention intensities are:

Model	$g = d/d_{\text{kv}}$	I_{attn} [FLOP/byte]
EuroLLM-1.7B	2	2
EuroLLM-9B	4	4
EuroLLM-22B	6	6

These values are far below the ridge point of either the A100 or H100. Decode attention is therefore strongly memory-bandwidth bound even when the projection and MLP GEMMs become compute-bound.

For the whole-model estimate, let N_W denote the number of weights that must be streamed during one decoding step. We exclude the input embedding table because inference reads only the embeddings of the current tokens, but include the full output projection:

Model	N_W	BF16 weight bytes	INT4 weight bytes
EuroLLM			
1.7B	1.395B	2.79 GB	0.70 GB
9B	8.628B	17.26 GB	4.31 GB
22B	21.851B	43.70 GB	10.93 GB

This assumes that all streamed matrices, including the output projection, are quantized. In practice, some implementations leave embeddings or the language model head in BF16.

In the fully memory-bound approximation, the upper-bound speedup from BF16 to INT4 weights

is

$$\text{speedup}_{W4} \leq \frac{2N_W + Q_{KV}}{0.5N_W + Q_{KV}}.$$

This approaches $4\times$ when weight traffic dominates, but approaches $1\times$ when KV-cache traffic dominates. Moreover, weights can be reused across the batch, whereas each sequence has its own KV cache. Large offline batches therefore make KV traffic important at much shorter context lengths.

The context length at which BF16 KV traffic equals the INT4 weight traffic is

$$S_{\text{cross}} = \frac{0.5N_W}{2BLd_{kv}b_a}.$$

For the concrete EuroLLM models, this gives:

Model	$B = 1$	$B = 4$	$B = 64$
EuroLLM-1.7B	7,094	1,774	111
EuroLLM-9B	25,077	6,269	392
EuroLLM-22B	49,395	12,349	772

For EuroLLM-1.7B, KV-cache traffic already exceeds the INT4 weight traffic at approximately 111 tokens for an offline batch of 64. For EuroLLM-9B, the corresponding crossover occurs at approximately 392 tokens. For EuroLLM-22B, its larger weight matrices delay the crossover to approximately 772 tokens.

At the native maximum context lengths, the memory-traffic upper bound predicts:

Model and context	$B = 1$	$B = 4$	$B = 64$
EuroLLM-1.7B, $S = 4096$	$2.90\times$	$1.91\times$	$1.08\times$
EuroLLM-9B, $S = 32768$	$2.30\times$	$1.48\times$	$1.04\times$
EuroLLM-22B, $S = 32768$	$2.80\times$	$1.82\times$	$1.07\times$

These numbers are upper bounds based only on total HBM traffic. They already show that a nominal $4\times$ reduction in weight bytes need not translate into a $4\times$ end-to-end decode speedup.

This explains why weight-only quantization does not always yield the speedup suggested by the reduction in model size. At small decode batches and short contexts, weight traffic dominates and the GEMM-only roofline predicts a large gain. At larger batches, weights are amortized across tokens and the GEMMs may become compute-bound. At long contexts, KV-cache reads and attention dominate the remaining memory traffic, so reducing weight bytes has progressively less effect.

Finally, the idealized model ignores the cost of unpacking and dequantizing weights, quantization scales and zero points, small- M GEMM under-utilization, kernel-launch overhead, normalization and positional-encoding kernels, paged-KV-cache indirections, communication, and sampling. It also assumes that every matrix is quantized and that all HBM traffic reaches peak bandwidth. These effects can make the realized speedup substantially smaller than the roofline upper bound.

A.2 Details on Quantization Formats

We provide additional details to the three quantization formats that we consider in this work: W8A8, W4A8, and W4A16 weight quantization. In W8A8, the weights are rescaled by the absmax of each row along the output dimension, while the activations are similarly quantized dynamically along the token dimension. As discussed in (Xiao et al., 2023), this maps well to hardware-accelerated GEMM kernels, as one can apply the scaling factors after the matrix multiplication:

$$\mathbf{Y} = \text{diag}(\Delta_{\mathbf{X}}^{\text{BF16}}) \cdot (\bar{\mathbf{X}}^{\text{INT8}} \cdot \bar{\mathbf{W}}^{\text{INT8}}) \cdot \text{diag}(\Delta_{\mathbf{W}}^{\text{BF16}}),$$

where $\mathbf{Y} \in \mathbb{R}^{d_{\text{in}} \times d_{\text{out}}}$ refers to the output matrix, $\bar{\mathbf{X}}^{\text{INT8}} \in \mathbb{R}^{d_{\text{in}} \times n}$ and $\bar{\mathbf{W}}^{\text{INT8}} \in \mathbb{R}^{n \times d_{\text{out}}}$ refer to the quantized activation and weight matrices, stored in INT8, and $\Delta_{\mathbf{X}}^{\text{BF16}} \in \mathbb{R}^{d_{\text{in}}}$ and $\Delta_{\mathbf{W}}^{\text{BF16}} \in \mathbb{R}^{d_{\text{out}}}$ refer to the activation scale vector, respectively the weight scale vector, stored in BFloat16.

For W4A8, weights are quantized with a group size of 128 following (Zhang et al., 2024b), while activations are dynamically quantized along the token dimension. Finally for W4A16 weight quantization, weights are also quantized with a group size of 128 while activations remain unchanged. All linear layers in the model, except for the embedding and prediction head layer, are being quantized using the same data format.

A.3 Details on document-level evaluation

We provide additional details on the document-level evaluation in this section. The prompt template used for the evaluation of both models can be found in Table 4, which we took from (Martins et al., 2025). We choose a temperature of $T = 0$, $\text{top}_p = 0.8$, and $\text{max_tokens} = 4096$ as decoding parameters and evaluate document chunking for target prompt-length thresholds in $\{100, 400, 800, 1000, 1500\}$ tokens for EuroLLM-1.7B and in $\{100, 400, 800, 1000, 2000, 4000\}$




Figure 4: Memory-bandwidth upper bound on the speedup from BF16 to INT4 weights for the three EuroLLM models with a BF16 KV cache. Increasing context length or batch size increases KV-cache traffic, which is unaffected by weight-only quantization. The relative memory-bound speedup is independent of the GPU’s absolute HBM bandwidth, although absolute latency is lower on the H100.

tokens for EuroLLM-9B on DocHPLT. On WMT24++ we evaluate document chunking for target prompt-length thresholds in $\{100, 200, 400, 500, 600, 800, 1000\}$. We note that we do not split segments and keep trailing chunks. This means that the chunk lengths can overshoot on some chunks and documents can end with a chunk much shorter than the threshold. Note that the target thresholds also include the rendered chat template. Across the four translation directions we have 27 matched documents for UK $\rightarrow$ EN, 13 documents for EN $\rightarrow$ UK from DocHPLT, as well as 19 matched documents for DE $\rightarrow$ EN and 11 documents for EN $\rightarrow$ DE from WMT24++. The common document subset was formed by performing an intersection of valid documents across all models, formats and target thresholds for each translation direction. We compute the chrF++ scores by setting `word_order = 2`. The transformers version used for the experiments was 5.14.1

A.4 Further Results on Chunk Length vs. Translation Quality

We provide further experimental results on how chunk length affects translation quality for English to Ukrainian translation on a subset of DocHPLT (Figure 5). Earlier experiments on EN $\leftrightarrow$ DE on WMT24++ can be found in Figure 6 and Figure 7. The results show that whether an increased degradation with context length depends not only on the quantization format and translation direction, but also on the model itself. Note that WMT24++ contains considerably shorter documents, such that translation quality is only evaluated up to a threshold of 1000 tokens limiting its power to say any-

thing for long document translations.

A.5 Further Results on Offline Benchmarking

We provide further results on offline benchmarking models from the EuroLLM family for sequence lengths of 100 tokens in Figure 8.

Table 4: Prompt templates for evaluating EuroLLM.

EuroLLM-Instruct	
Role	Prompt Template
System	Translate all user texts to {target_lang}.
User	{source_text}




Figure 5: Chunk length vs. chrF++ Score for EuroLLM-1.7B (top row) and EuroLLM-9B (bottom row) across different quantization formats averaged across a subset of 13 documents from DocHPLT. Shaded area corresponds to bootstrapped 95% CI of document-level means. Each point denotes a document-level chrF++ score in the left subplots and a paired BF16-quantized difference in the right subplots. Note the different y limits for better visibility.




Figure 6: Chunk length vs. chrF++ Score for EuroLLM-1.7B (top row) and EuroLLM-9B (bottom row) across different quantization formats averaged across a subset of 11 documents from WMT24++. Shaded area corresponds to bootstrapped 95% CI of document-level means. Each point denotes a document-level chrF++ score in the left subplots and a paired BF16-quantized difference in the right subplots. Note the different y limits for better visibility.




Figure 7: Chunk length vs. chrF++ Score for EuroLLM-1.7B (top row) and EuroLLM-9B (bottom row) across different quantization formats averaged across a subset of 19 documents from WMT24++. Shaded area corresponds to bootstrapped 95% CI of document-level means. Each point denotes a document-level chrF++ score in the left subplots and a paired BF16-quantized difference in the right subplots. Note the different y limits for better visibility.




Figure 8: Offline benchmark measuring average output throughput in tok/s against average wall time for a sequence length of *100 tokens* across three model sizes on a A100 (top row) and a H100 (bottom row).