

Towards Inducing Long-Context Abilities in Multilingual Neural Machine Translation Models

Varun Gumma* Pranjali A. Chitale* Kalika Bali

Microsoft Corporation

{varun230999, pranjaliचितले}@gmail.com

Abstract

Neural Machine Translation (NMT) models have traditionally used Sinusoidal Positional Embeddings (PEs), which often struggle to capture long-range dependencies and are inefficient for handling extended context or document-level translation tasks. This work addresses the challenge of transitioning pre-trained NMT models from absolute Sinusoidal PEs to Relative PEs, such as ROPE and ALiBI, without compromising performance. We demonstrate that parameter-efficient fine-tuning, using only a small amount of high-quality data, can successfully facilitate this transition. Experimental results indicate that switching from Sinusoidal to Relative PEs results in competitive translation quality on sentence-level evaluation benchmarks. Additionally, models trained with ROPE consistently outperform those using ALiBI and Sinusoidal PEs on document-level benchmarks across both string-based metrics and qualitative evaluations. Moreover, we find that a small amount of long-context data in a few languages is sufficient for cross-lingual length generalization, thereby inducing long-context capabilities.

1 Introduction

Neural Machine Translation (NMT) models have become critical tools in many Natural Language Processing (NLP) applications, especially for document-level translation. While NMT models perform well on sentence-level tasks, they often struggle with longer sequences, where preserving contextual coherence across sentences is essential. This limitation is largely due to their reliance on Absolute Positional Embedding (PE) methods, such as Sinusoidal PEs, which can limit the model’s ability to extrapolate effectively to longer contexts. Recent advancements in Relative PEs, such as ROTARY POSITIONAL EMBEDDINGS (ROPE) and AT-

TENTION WITH LINEAR BIASES (ALiBI), have demonstrated superior performance in handling longer sequences. These relative positional embedding techniques enable models to generalize better to long contexts, making them well-suited for document-level translation tasks. However, most state-of-the-art NMT models were originally trained using Sinusoidal embeddings, and retraining them with newer methods like ROPE or ALiBI is computationally expensive.

This presents a crucial question: *Can we modify the positional embeddings in pre-trained NMT models to improve long-context translation without significant retraining or performance degradation?* Addressing this problem would enable NMT models to translate long documents more effectively, improving translation coherence and quality across paragraphs and sections. Moreover, long-context NMT models could have broader applications, such as generating synthetic data for multilingual pre-training. Such data is often needed in scenarios where natural multilingual corpora are limited, especially for low-resource languages. However, current sentence-level NMT systems are suboptimal for this task as they lack the capacity to capture broader contextual dependencies in long documents.

In this work, we investigate whether it is possible to replace the positional embeddings in pre-trained NMT models with relative positional embeddings post-hoc. Our study explores both parameter-efficient fine-tuning (PEFT) methods and full fine-tuning (FFT) strategies to restore model performance while enabling long-context generalization. In the remainder of this paper, we present the following contributions:

- We conduct a comprehensive study on modifying positional embeddings in pre-trained transformer models post-hoc, aiming to retain performance on existing tasks while enabling

* Equal contribution

long-context capabilities.

- Experiments show that ROPE and ALIBI perform similarly on sentence-level benchmarks, but ROPE outperforms both ALIBI and SINE on document-level tasks, based on string-based evaluation metrics, which was then further validated by qualitative evaluation.
- Moreover, fine-tuning with ROPE using minimal long-context data across a few languages demonstrates effective cross-lingual length generalization, even without extensive long-context training data.
- Lastly, we open-source¹ our framework to support the efficient training and fine-tuning of long-context machine translation models. Further, our best-performing RoPE models are also available on HuggingFace².

Our findings indicate that Positional Embeddings in pre-trained NMT models can be modified and adapted without compromising performance, to efficiently induce long-context capabilities.

2 Related Works

Multilingual Neural Machine Translation (MNMT) (Zhang et al., 2019; Firat et al., 2016; Aharoni et al., 2019; Arivazhagan et al., 2019) has emerged as a leading methodology for creating translation systems capable of managing multiple languages. These MNMT systems leverage a unified encoder-decoder transformer framework (Vaswani et al., 2017), along with language-specific embeddings to incorporate linguistic nuances. Previous works (Fan et al., 2021; Costa-jussà et al., 2022) have developed many-to-many models at a global scale, while other works (Ramesh et al., 2022; Gala et al., 2023) have primarily focused on many-to-one and one-to-many multilingual models for Indic languages. An in-depth examination of MNMT methodologies is presented in the survey by Dabre et al. (2020).

Positional Embedding Methods Transformer-based models require explicit positional information, as all timesteps are processed in parallel. To

address this, Vaswani et al. (2017) proposed an empirical absolute PE approach, using a Fourier series with alternating sine and cosine terms. Vaswani et al. (2017) also experimented with learnable PEs, but these did not show significant performance differences compared to absolute Sinusoidal embeddings. Later works (Rosendahl et al., 2019; Shaw et al., 2018) demonstrated the advantages of relative PEs in transformers. Press et al. (2022) discussed ALIBI, a simple Relative PE strategy that adds relative token distances to the attention matrix, enabling models to generalize to longer sequences during inference, even when trained on shorter sequences. ROTARY POSITIONAL EMBEDDINGS (ROPE) (Su et al., 2024) introduced the concept of rotating vectors according to their positions to encode relative positional dependencies directly into the self-attention mechanism. Recent works (Sun et al., 2023; Chen et al., 2023; Peng et al., 2023; Ding et al., 2024; Ostmeier et al., 2024) have also proposed adaptations of ROPE for much longer sequences. In this study, we focus on the standard variant of ROPE, though our method is generalizable to other variants.

Parameter-Efficient Fine-tuning (PEFT) (Houlsby et al., 2019; Bapna and Firat, 2019; Hu et al., 2022) focuses on updating only a subset of model parameters to efficiently generalize to new tasks or domains, rather than fine-tuning all parameters. This approach mitigates the risks of catastrophic forgetting or overfitting in large models. Low-Rank Adaptation (LORA) (Hu et al., 2022), the most popular parameter-efficient fine-tuning strategy, hypothesizes that weight updates are low-rank and injects small trainable parameters into pre-trained weights to approximate low-rank updates.

Evaluation Metrics for Machine Translation Evaluation in machine translation is well explored, with string-based metrics like BLEU (Papineni et al., 2002), ChrF (Popović, 2015), and ChrF++ (Popović, 2017) being the most commonly used. However, prior studies (Kocmi et al., 2021) have pointed out limitations of BLEU, particularly its bias towards exact matches, which makes it less suitable for languages with flexible word order. Trained metrics, such as COMET (Rei et al., 2020, 2022) and BLEURT (Sellam et al., 2020), employ a regressor on top of an encoder to predict translation quality. These metrics have been shown to correlate more closely with human judgments, outperform-

¹ <https://github.com/VarunGumma/fairseq>

² <https://aka.ms/IndicTrans2-RoPE>

ing traditional string-based measures (Kocmi et al., 2021). However, their performance can vary significantly across different languages and may not always be reliable for Indian languages (Gala et al., 2023).

LLM-based Evaluators for Machine Translation Recently, large language models (LLMs) have emerged as general-purpose evaluators in multilingual scenarios (Hada et al., 2024b,a), exhibiting superior correlation with human judgment. This approach has also been applied to machine translation (MT). GEMBA (Kocmi and Federmann, 2023b) was one of the initial attempts to use LLMs for evaluating MT outputs. Follow-up works like EAPROMPT (Lu et al., 2024), AUTOMQM (Fernandes et al., 2023), and GEMBA-MQM (Kocmi and Federmann, 2023a) have utilized LLMs in a MULTIDIMENSIONAL QUALITY METRICS (MQM) framework (Lommel et al., 2014; Freitag et al., 2021), prompting models to identify and penalize erroneous spans of text, thereby offering more interpretable and standardized evaluations. These approaches demonstrate superior correlation with human judgments compared to other model-based metrics like COMET and BLEURT. Given that LLMs typically support longer context lengths, they are well-suited for evaluating tasks involving extended input sequences. Inspired by GEMBA-MQM (Kocmi and Federmann, 2023a), we employ GPT-4 as an evaluator to investigate the qualitative differences between baseline models.

3 Methodology

In this section, we discuss the dataset curation, model training and evaluation procedures.

3.1 Model

Our experimentation focuses on standard Encoder-Decoder transformer models. While several massively multilingual models, such as NLLB (Costa-jussà et al., 2022) and its distilled variants, have been released, our experimentation requires the availability of long-context or document-level training and evaluation datasets. Consequently, we focus on language-family specific models that allow for more controlled experimentation. For our experiments, we chose the distilled variants of INDIC-TRANS2 (Gala et al., 2023; Gumma et al., 2023). A key difference is that the distilled models released in Gala et al. (2023) we used were not trained with

BackTranslation (Sennrich et al., 2016) data, which we incorporate in our experimental setup. We experimented with models based on ROPE and ALIBI by replacing the Sinusoidal Positional Embedding in the pre-trained model with the corresponding relative Positional Embedding. This replacement led to a significant performance drop, prompting us to fine-tune the models using high-quality data (Mohiuddin et al., 2022; Gumma et al., 2023; Gala et al., 2023) to study their recovery in each setting.

3.2 Data

3.2.1 Train set

For maximum performance gains for our models, we choose a high-quality mixture of seen and unseen distributions for our training data. To this end, for the seen distribution, we select the BPCC-SEED, which is a high-quality human-verified subset of BPCC (BHARAT PARALLEL CORPUS COLLECTION), and has shown to be effective in improving model performance (Gala et al., 2023). As for the unseen distribution, we sample the top 150K pairs for each language from the BackTranslation (BPCC-BT) subset of BPCC based on the provided LASER (Costa-jussà et al., 2022) or LABSE (Feng et al., 2022) cosine-similarity scores³. In case similarity scores are unavailable for a translation pair, a random sample was chosen. Additionally, for the unseen data, we also include the ASIAN LANGUAGE TREEBANK (ALT) Corpus (Riza et al., 2016), and the COPARA ALIGNED CORPUS (E et al., 2023), which contains document-level translation pairs for four Dravidian languages, Telugu, Kannada, Malayalam and Tamil. The overall statistics of the dataset is available in table 1.

	BPCC SEED	BPCC BT	ALT	CoPARA	Total
<i>En-Indic</i>	2.29	2.76	0.04	0.02	5.11
<i>Indic-En</i>	2.29	3.75	0.04	0.02	6.09

Table 1: Number of training instances per subset (in Millions)

3.2.2 Dev/Test sets

Most existing benchmarks for Indian languages are composed of sentence-level test sets, which may not adequately reflect long-context performance or extrapolation capabilities. To address

³whichever was available for that language

	256-512	512-1024	1024-2048	2048-4096	Total
<i>eng_Latn-kan_Knda</i>	325	66	4	4	399
<i>eng_Latn-mal_Mlym</i>	313	75	4	3	395
<i>eng_Latn-tam_Taml</i>	331	77	5	3	416
<i>eng_Latn-tel_Telu</i>	321	71	5	4	401
Total	1290	289	18	14	1611

Table 2: Distribution of long-context data in the fine-tuning data mixture. Indic-En pairs also follow the same distribution. The language-pairs not mentioned do not have any data in the aforementioned lengths.

this, we construct document-level development and test sets⁴ using FLORES (Costa-jussà et al., 2022) and IN22-CONV (Gala et al., 2023). In this setup, sentences are grouped within the same context window, making them suitable for document-level evaluation. For FLORES, we group sentences by metadata, including URL, domain, and topic headers, to map sentence-level pairs into document-level pairs, where each document-level unit consists of three sentences. This method is viable since this benchmark was created by professional translators who translate three consecutive sentences from every selected context window, as described by Goyal et al. (2022). As a result, we obtain 286 merged sentence pairs per language, with a total of 5720 pairs. In case of IN22-CONV, the dataset was originally created by translating entire conversations from English into the target languages by professional translators, and used for sentence-level evaluation by Gala et al. (2023). We use metadata provided to trace back the conversations and treat the entire conversation as a sequence for long-context evaluation. This dataset has 44 conversations per language and covers 22 languages. The token length distribution of the synthesized IN22-CONV test set is shown in Figure 1. IN22-CONV serves as an accurate representation of a true document-level evaluation set due to its contextual dependencies across multiple turns in the conversation. Additionally, it supports all the languages under consideration, making it the primary benchmark for document-level evaluation.

3.3 Metrics and Evaluation

For consistency, we use the ChrF++ score (Popović, 2017) as the primary string-based metric, and fol-

⁴ <https://aka.ms/DocMTBenchmarks>

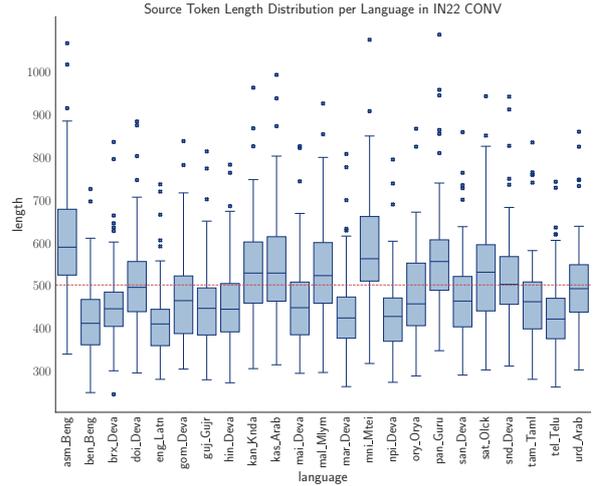


Figure 1: Distribution of source token lengths for the Indic side of conversations in the IN22-CONV Doc-Level test set tokenized using the INDICTRANS2 tokenizer. The red line is the average across all languages.

Test Set	Languages Covered
IN22	Assamese , Bengali , Bodo, Dogri, Konkani, Gujarati , Hindi , Kannada , Kashmiri, Maithili, Malayalam , Marathi , Manipuri, Nepali, Odia , Punjabi , Sanskrit, Santali, Sindhi, Tamil , Telugu , Urdu
FLORES	Assamese, Bengali, Gujarati, Hindi, Kannada, Kashmiri, Maithili, Malayalam, Marathi, Manipuri, Nepali, Odia, Punjabi, Sanskrit, Santali, Sindhi, Tamil, Telugu, Urdu
ALT	Bengali, Hindi

Table 3: List of languages in each test set. The top-12 languages are highlighted.

low the same evaluation practices as outlined in Gala et al. (2023) by normalizing the indic texts using INDICNLPLIBRARY⁵. For brevity, we report the aggregate score across all the languages in the respective test set, which are denoted in Table 3. For GPT-4O-based evaluation, we use the prompt template from GEMBA-MQM (Kocmi and Federmann, 2023a) as a starting point and make slight adjustments to elicit an implicit quality score between 0 and 1, in addition to identifying the error spans following the MQM guidelines. Another notable difference is that, since we are evaluating long-context translations, including few-shot examples would be inefficient due to the limited context

⁵ https://github.com/anoopkunchukuttan/indic_nlp_library

length. Therefore, we opt for a zero-shot evaluation setup. We report both the implicit score as rated by GPT-4o as well as calculated MQM-score (based on the scoring guidelines outlined in Kocmi and Federmann (2023a)). We start with an ideal score of 25, deducting points based on identified errors and their severity: **CRITICAL** errors are penalized by a factor of 5, **MAJOR** by 3, and **MINOR** by 1. The final score is normalized by dividing by 25, yielding a value between 0 and 1.

3.4 Finetuning Strategies

In our fine-tuning experiments, we swap the Sinusoidal Positional Embedding module with a ROPE, or ALiBi module, and fine-tune the model *as efficiently as possible* from that point to analyse the performance recovery.

FFT First, we choose a full-fine-tuning (FFT) approach which tweaks all the 211.78M parameters of the model. This would allow the embeddings, layer-normalizations, and the rest of the linear layers⁶ to be tuned as per the new relative Positional Embedding.

LORA In vanilla PEFT method, we add trainable parameters only to the linear layers resulting in 6.48M (3.06%) trainable parameters. Note that, unlike FFT, in this case, the embeddings and layer-normalizations do not have any trainable parameters. The main hypothesis in this experiment is we intend to understand if the model can adapt to the change in the Positional Embeddings by addition of few trainable parameters as FFT may be computationally expensive as well as result in catastrophic forgetting.

minLORA The relative Positional Embedding modules are added to the self-attention modules at every layers. Therefore, we intend to leverage this property and understand whether, we can further cut down on parameters and apply these additional parameters only to the affected modules. Hence, we test the hypothesis of tuning only the self-attention modules with LORA resulting in 2.35M (1.11%) trainable parameters.

⁶SELFATTENTION, CROSSATTENTION, FEEDFORWARD

4 Experiments

We use the FAIRSEQ⁷ framework (Ott et al., 2019) for all of our experiments. All the finetuning experiments were conducted on $8 \times$ H100 80Gb GPUs, and every model was trained till convergence with early-stopping. The exact hyperparameter sets are mentioned in Table 4.

Notably, for ROPE, we apply the rotation function to only the first half segment of the query and key feature vectors while leaving the rest as is. We empirically notice that the *partial-rotation*⁸ implementation gives slightly better performance as compared to a full-rotation.

Hyperparameter	Value
Global Max Tokens	32K
Optimizer	Adam
LR scheduler	inverse_sqrt
Gradient clip norm	1.0
Adam betas	(0.9, 0.98) / (0.9, 0.999)
Checkpoint Steps	500
Checkpoint Criteria	BLEU @ beam = 5
Patience	20
Max Positions	4096
Warmup steps	4000
Learning rate	1e-4
Dropout	0.2 / 0
Label smoothing	0.1 / 0
LoRA rank	16
LoRA alpha	32
LoRA dropout	0.1
LoRA Rank Scaled	True
FP16	True

Table 4: Default hyperparameter settings for finetuning. The cells with two numbers present the values for FFT/LORA.

5 Results and Discussions

In this section, we discuss the results that aim to understand whether Positional Embeddings can be efficiently swapped posthoc in a trained model.

5.1 Swapping out Positional Embeddings

Figure 2 demonstrate the performance of several experimental setups that involve swapping of Positional Embeddings and followed by fine-tuning (if applicable). The raw scores are presented in Table 6.

⁷Our variant implements, LORA, ROPE, ALiBi, and a slightly efficient MULTIHEADATTENTION as per the requirements of this work

⁸<https://github.com/lucidrains/x-transformers/issues/40>

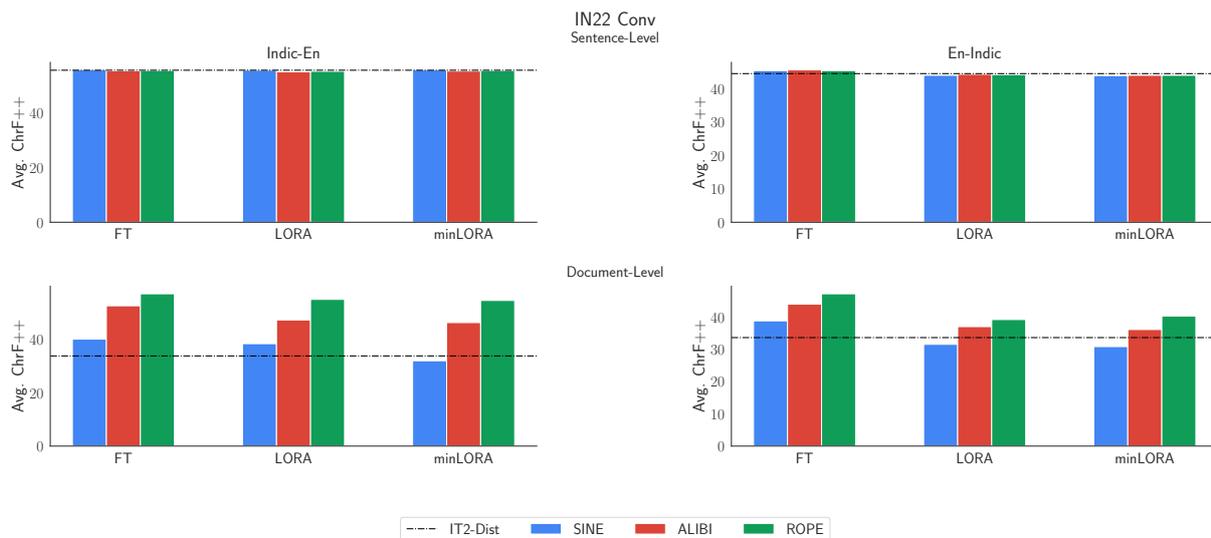


Figure 2: ChrF++ scores on Sentence-Level (top) and Document-Level (bottom) benchmarks. Results are presented for baselines using three fine-tuning setups (FFT, LORA, and min-LORA) that compare three types of positional embeddings (SINE, ALIBI, and ROPE), along with the pre-trained model performance baseline.

Fine-tuning is key to performance recovery.

When Positional Embeddings such as ROPE or ALIBI are introduced, we observe a significant drop in performance as the model is not trained to work with these. However, across various fine-tuning setups, we consistently observe that the performance can be regained through fine-tuning (see Figure 2). This demonstrates that it is feasible to modify Positional Embeddings at later stages and still retain the performance of the base model.

Efficacy of the Positional Embedding switch. A notable finding in this setting is that fine-tuning a model originally using Sinusoidal embeddings yielded performance comparable to models fine-tuned after the switch to using ROPE or ALIBI embeddings on sentence-level benchmarks as observed in Figure 2. This further reinforces the conclusion that Positional Embeddings in pre-trained transformer models can be safely tweaked even in later stages. Although Sinusoidal, ROPE, and ALIBI embeddings show similar performance in sentence-level test sets. The advantage of switching to ROPE or ALIBI can be observed in case of document-level test sets in line with the findings of Su et al. (2024) and Press et al. (2022).

Cross-lingual length-generalization. Our fine-tuning dataset primarily comprised sentence-level data, as detailed in Section 3.2.1, with only a limited amount of long-context or paragraph-level data available for four Dravidian languages accounting for less than 0.03% of the complete fine-tuning

data mixture (see Table 1). The token distribution across this dataset is illustrated in Table 2. While most languages lacked long-context data, evaluation on document-level benchmarks still demonstrated strong performance on these, as seen in Figure 2. This indicates that fine-tuning with a small amount of long-context data introducing relative positional embeddings can enable robust cross-lingual length generalization. This finding is particularly important because long-context data is often scarce across many languages. Nevertheless, effective adaptation to longer contexts was achieved through cross-lingual generalization, even with minimal long-context data in select languages.

Minimalistic LORA is on par with LORA

Across all the fine-tuning strategies evaluated, minimalistic LORA (min-LORA) fine-tuning, as well as FFT and LORA, produced comparable performance on sentence-level benchmarks. This suggests that extensive parameter adjustments are not necessary; adding LORA parameters solely to the self-attention modules, where the relative Positional Embedding modules are integrated, is sufficient for adaptation. This approach is highly parameter-efficient, requiring only 2.35M trainable parameters (1.11%). On the document-level benchmarks, we observe that both LORA and min-LORA are on par suggesting that min-LORA is adequate. However these parameter-efficient approaches slightly lag behind the respective FFT baseline.

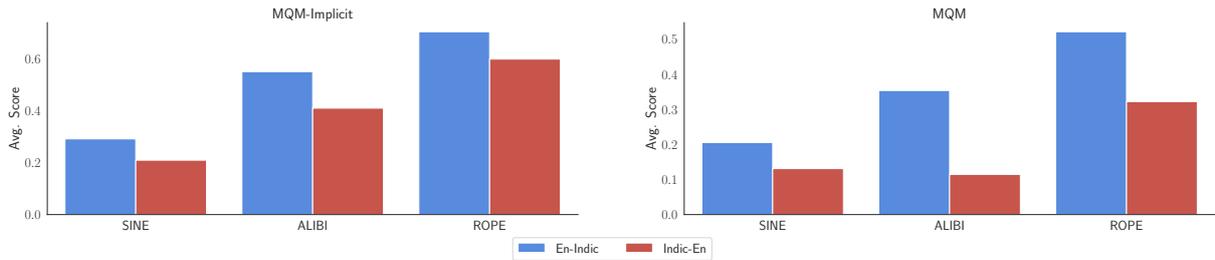


Figure 3: Implicit MQM scores and normalized formula-based MQM scores on the IN22-CONV benchmark, averaged across top-12 languages highlighted in table 3

Throughput ROPE has lower throughput compared to Sinusoidal embeddings due to its higher number of operations. Since the rotation function in ROPE is applied at every layer, it has a time complexity of $\mathcal{O}(n)$, whereas Sinusoidal embeddings are added only once, after the token embeddings, resulting in $\mathcal{O}(1)$ complexity. The exact decoding rates (in tokens per second) are illustrated in Figure 4. An efficient implementation of ALIBI (Press et al., 2022) achieves $\mathcal{O}(1)$ complexity, making it slightly slower than Sinusoidal but still significantly faster than ROPE. Therefore, for a basic sentence-level MT system, Sinusoidal PEs offer the best efficiency in terms of throughput and performance. However, for document-level or longer contexts, ROPE provides notable qualitative advantages, although it comes with a decrease in throughput.

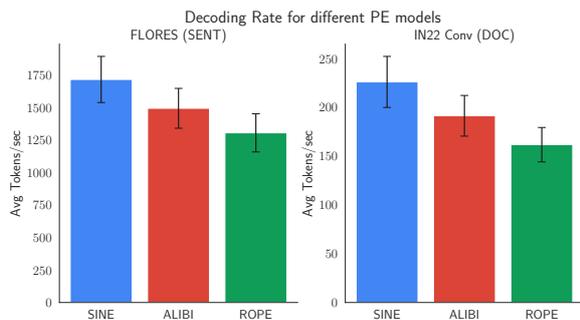


Figure 4: Throughput in terms of token/sec. Higher is better.

5.2 Qualitative Analysis

In the preceding section, we conducted a string-metrics-based evaluation on standard benchmarks such as IN22-CONV, ALT, and FLORES, covering both sentence-level and document-level variants. However, the effectiveness of these string-based metrics in evaluating translation quality, especially in longer-context scenarios where factors like se-

mantic coherence, co-reference resolution, and continuity are critical, remains under-explored.

Therefore, for a more fine-grained analysis we conduct GPT-4O based evaluation as described in Section 3. We focus exclusively on the FFT variants for GPT-4O evaluation, as they were the highest-performing systems. Figure 3 demonstrates two types of evaluations: an implicit score assigned by GPT-4O and a calculated score based on MQM annotations generated by GPT-4O. Both metrics clearly indicate that ROPE-based baselines outperform the SINE and ALIBI counterparts, with ALIBI proving superior to the SINE baseline. We also observe that, in a relative sense, the rankings derived from the implicit scores generally align with those based on the calculated scores. These results underscore the effectiveness of the proposed approach, showing that substituting positional embeddings had a notable positive impact on tasks involving longer contexts. We also present a table of errors identified by GPT-4O on the same sample conversation across the three systems being compared: ROPE, SINE, and ALIBI in tables 7 to 9 respectively. The MQM annotations clearly show that the translations produced by the ROPE model exhibit less severe errors, followed by ALIBI. In contrast, the SINE model has a higher number of errors with greater severity. This aligns with the earlier observations, confirming that the ROPE models are indeed superior to the SINE and ALIBI counterparts.

6 Additional Experiments

To assess the scalability of our approach, we extend our methodology to the larger 1B INDICTRANS2 model. For consistency, we reuse the same training data and hyperparameters as described in Section 3.2.1 and Table 4, respectively. In this case, we focus solely on FFT, as it yielded the highest performance recovery among the fine-tuning strategies

explored in Section 3.4. Table 5 presents the results of this experiment. Consistent with our previous findings, ROPE continues to exhibit the best length generalization, followed by ALiBI, as observed on the Document-Level IN22-CONV subset.

		FT			
DATASET		IT2	SINE	ALiBI	ROPE
Indic-En	IN22-GEN	63.8	63.6	62.8	63.3
	IN22-CONV	56.1	56.7	55.0	55.3
	FLORES	61.1	<u>61.0</u>	60.2	60.5
En-Indic	IN22-GEN	48.3	48.2	<u>48.4</u>	48.7
	IN22-CONV	44.8	<u>45.3</u>	45.9	45.9
	FLORES	<u>47.8</u>	48.3	48.3	48.3
Indic-En	ALT	66.9	<u>66.8</u>	66.5	66.7
	IN22-CONV	31.7	40.8	<u>51.0</u>	58.3
	FLORES	62.8	<u>62.6</u>	61.7	62.5
En-Indic	ALT	55.5	55.9	<u>55.8</u>	55.9
	IN22-CONV	28.6	35.5	<u>45.3</u>	48.6
	FLORES	49.1	49.8	<u>49.5</u>	49.8

Table 5: IT2-1B ChrF++ scores on sentence-level (top) and document-level (bottom) evaluation benchmarks. The best score is **bolded** and second best is underlined.

7 Conclusion

In this study, we demonstrated that standard NMT systems can be *efficiently retrofitted* with relative positional embeddings such as ROPE and ALiBI, enabling them to handle longer contexts. Moreover, we showed that this retrofitting can be achieved through parameter-efficient fine-tuning using minimal document-level data, even across a few languages, with evidence of cross-lingual transfer. Although our experiments focused on the machine translation (MT) task, we believe that this approach can generalize to other encoder-decoder and encoder-only models, potentially improving their long-context capabilities in a wide range of applications.

8 Limitations

Our work aligns with modular deep learning principles, exploring the post-hoc interchangeability of key components like Positional Embeddings. Although our preliminary findings suggest this is possible, several limitations should be noted:

- All our experiments have been limited to a single model primarily due to the complete availability of relevant information regarding

the INDICTRANS2 model and dataset. Further the lack of availability of long-context data is also a limiting factor in performing this experimentation across languages.

- We use string-based automatic evaluation and GPT-4O as an evaluator, but human evaluation may be needed to further validate the results and reinforce our findings.
- The benchmarks we test in this study are smaller in size due to the limited availability of long-context MT datasets for Indian languages.
- Lastly, we also acknowledge that no specific hyperparameter tuning was conducted for this task, and we utilize the standard set available in the literature.

9 Acknowledgements

We thank Raj Dabre (Senior Researcher, NICT), Jay Gala (Research Assistant, MBZUAI) and Amit Sharma (Principal Researcher, MSRI) for their for the discussions that helped us in refining the idea.

References

- Roei Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3874–3884, Minneapolis, Minnesota. Association for Computational Linguistics.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. Massively multilingual neural machine translation in the wild: Findings and challenges. *arXiv preprint arXiv: 1907.05019*.
- Ankur Bapna and Orhan Firat. 2019. [Simple, scalable adaptation for neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1538–1548, Hong Kong, China. Association for Computational Linguistics.
- Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. 2023. [Extending context window of large language models via positional interpolation](#).

- 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, Semarley 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](#).
- Raj Dabre, Chenhui Chu, and Anoop Kunchukuttan. 2020. [Multilingual neural machine translation](#). In *Proceedings of the 28th International Conference on Computational Linguistics: Tutorial Abstracts*, pages 16–21, Barcelona, Spain (Online). International Committee for Computational Linguistics.
- Yiran Ding, Li Lyna Zhang, Chengruidong Zhang, Yuanyuan Xu, Ning Shang, Jiahang Xu, Fan Yang, and Mao Yang. 2024. Longrope: Extending llm context window beyond 2 million tokens. *arXiv preprint arXiv: 2402.13753*.
- Nikhil E, Mukund Choudhary, and Radhika Mamidi. 2023. [CoPara: The first Dravidian paragraph-level n-way aligned corpus](#). In *Proceedings of the Third Workshop on Speech and Language Technologies for Dravidian Languages*, pages 88–96, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *Journal of Machine Learning Research*, 22(107):1–48.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2022. [Language-agnostic BERT sentence embedding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.
- Patrick Fernandes, Daniel Deutsch, Mara Finkelstein, Parker Riley, André Martins, Graham Neubig, Ankush Garg, Jonathan Clark, Markus Freitag, and Orhan Firat. 2023. [The devil is in the errors: Leveraging large language models for fine-grained machine translation evaluation](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 1066–1083, Singapore. Association for Computational Linguistics.
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016. [Multi-way, multilingual neural machine translation with a shared attention mechanism](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 866–875, San Diego, California. Association for Computational Linguistics.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, errors, and context: A large-scale study of human evaluation for machine translation](#). *Transactions of the Association for Computational Linguistics*, 9:1460–1474.
- Jay Gala, Pranjal A Chitale, A K Raghavan, Varun Gumma, Sumanth Doddapaneni, Aswath Kumar M, Janki Atul Nawale, Anupama Sujatha, Ratish Pudupully, Vivek Raghavan, Pratyush Kumar, Mitesh M Khapra, Raj Dabre, and Anoop Kunchukuttan. 2023. [Indictrans2: Towards high-quality and accessible machine translation models for all 22 scheduled indian languages](#). *Transactions on Machine Learning Research*.
- 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*, 10:522–538.
- Varun Gumma, Raj Dabre, and Pratyush Kumar. 2023. [An empirical study of leveraging knowledge distillation for compressing multilingual neural machine translation models](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 103–114, Tampere, Finland. European Association for Machine Translation.
- Rishav Hada, Varun Gumma, Mohamed Ahmed, Kalika Bali, and Sunayana Sitaram. 2024a. [METAL: Towards multilingual meta-evaluation](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2280–2298, Mexico City, Mexico. Association for Computational Linguistics.
- Rishav Hada, Varun Gumma, Adrian Wynter, Harshita Diddee, Mohamed Ahmed, Monojit Choudhury, Kalika Bali, and Sunayana Sitaram. 2024b. [Are large language model-based evaluators the solution to scaling up multilingual evaluation?](#) In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 1051–1070, St. Julian’s, Malta. Association for Computational Linguistics.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.

- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Tom Kocmi and Christian Federmann. 2023a. [GEMBA-MQM: Detecting translation quality error spans with GPT-4](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 768–775, Singapore. Association for Computational Linguistics.
- Tom Kocmi and Christian Federmann. 2023b. [Large language models are state-of-the-art evaluators of translation quality](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203, Tampere, Finland. European Association for Machine Translation.
- Tom Kocmi, Christian Federmann, Roman Grundkiewicz, Marcin Junczys-Dowmunt, Hitokazu Matsushita, and Arul Menezes. 2021. [To ship or not to ship: An extensive evaluation of automatic metrics for machine translation](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 478–494, Online. Association for Computational Linguistics.
- Arle Lommel, Hans Uszkoreit, and Aljoscha Burchardt. 2014. [Multidimensional quality metrics \(MQM\): A framework for declaring and describing translation quality metrics](#). *Revista Tradumàtica: tecnologies de la traducció*.
- Qingyu Lu, Baopu Qiu, Liang Ding, Kanjian Zhang, Tom Kocmi, and Dacheng Tao. 2024. [Error analysis prompting enables human-like translation evaluation in large language models](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 8801–8816, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Tasnim Mohiuddin, Philipp Koehn, Vishrav Chaudhary, James Cross, Shruti Bhosale, and Shafiq Joty. 2022. [Data selection curriculum for neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1569–1582, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Sophie Ostmeier, Brian Axelrod, Michael E. Moseley, Akshay Chaudhari, and Curtis Langlotz. 2024. [Liere: Generalizing rotary position encodings](#). *arXiv preprint arXiv: 2406.10322*.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. 2023. [Yarn: Efficient context window extension of large language models](#). *arXiv preprint arXiv: 2309.00071*.
- 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.
- Maja Popović. 2017. [chrF++: words helping character n-grams](#). In *Proceedings of the Second Conference on Machine Translation*, pages 612–618, Copenhagen, Denmark. Association for Computational Linguistics.
- Ofir Press, Noah Smith, and Mike Lewis. 2022. [Train short, test long: Attention with linear biases enables input length extrapolation](#). In *International Conference on Learning Representations*.
- Gowtham Ramesh, Sumanth Doddapaneni, Aravindh Bheemaraj, Mayank Jobanputra, Raghavan AK, Ajitesh Sharma, Sujit Sahoo, Harshita Diddee, Mahalakshmi J, Divyanshu Kakwani, Navneet Kumar, Aswin Pradeep, Srihari Nagaraj, Kumar Deepak, Vivek Raghavan, Anoop Kunchukuttan, Pratyush Kumar, and Mitesh Shantadevi Khapra. 2022. [Samanantar: The largest publicly available parallel corpora collection for 11 Indic languages](#). *Transactions of the Association for Computational Linguistics*, 10:145–162.
- Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022. [COMET-22: Unbabel-IST 2022 submission for the metrics shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- 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.
- Hammam Riza, Michael Purwoadi, Gunarso, Teduh Uliniansyah, Aw Ai Ti, Sharifah Mahani Aljunied, Luong Chi Mai, Vu Tat Thang, Nguyen Phuong Thai, Vichet Chea, Rapid Sun, Sethserey Sam, Sopheap Seng, Khin Mar Soe, Khin Thandar Nwet, Masao Utiyama, and Chenchen Ding. 2016. [Introduction of the asian language treebank](#). In *2016 Conference of The Oriental Chapter of International Committee for Coordination and Standardization of Speech Databases and Assessment Techniques (O-COCOSDA)*, pages 1–6.

- Jan Rosendahl, Viet Anh Khoa Tran, Weiyue Wang, and Hermann Ney. 2019. [Analysis of positional encodings for neural machine translation](#). In *Proceedings of the 16th International Conference on Spoken Language Translation*, Hong Kong. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468, New Orleans, Louisiana. Association for Computational Linguistics.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. [Roformer: Enhanced transformer with rotary position embedding](#). *Neurocomput.*, 568(C).
- Yutao Sun, Li Dong, Barun Patra, Shuming Ma, Shao-han Huang, Alon Benhaim, Vishrav Chaudhary, Xia Song, and Furu Wei. 2023. [A length-extrapolatable transformer](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14590–14604, Toronto, Canada. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Wen Zhang, Yang Feng, Fandong Meng, Di You, and Qun Liu. 2019. [Bridging the gap between training and inference for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4334–4343, Florence, Italy. Association for Computational Linguistics.

A Appendix

A.1 Scores

The raw scores corresponding to Figure 2 are shown in Table 6.

A.2 Qualitative Examples

The MQM annotations generated by GPT-4o for a test conversation across for each baseline is presented in Tables 7 to 9.

		FT				LoRA			minLoRA		
DATASET		IT2-Dist	SINE	ALiBi	ROPE	SINE	ALiBi	ROPE	SINE	ALiBi	ROPE
En-Indic	IN22-GEN	48.0	48.2	<u>48.7</u>	48.8	47.9	47.8	47.7	47.7	47.3	47.6
	IN22-CONV	44.6	<u>45.5</u>	45.7	45.5	44.1	44.4	44.3	44.0	44.1	44.1
	FLORES	47.7	48.4	48.5	<u>48.5</u>	47.9	47.8	47.9	47.8	47.7	47.7
Indic-En	IN22-GEN	62.0	62.2	61.8	<u>62.1</u>	61.8	61.1	61.5	61.7	61.1	61.5
	IN22-CONV	55.7	55.8	55.5	55.6	55.7	55.1	55.3	<u>55.8</u>	55.4	55.6
	FLORES	59.5	59.7	59.4	<u>59.6</u>	59.4	58.9	59.1	59.4	59.0	59.2
En-Indic	ALT	55.0	<u>55.2</u>	55.2	55.4	55.1	55.0	54.8	55.2	55.0	55.0
	IN22-CONV	33.7	38.9	<u>44.1</u>	47.3	31.6	37.1	39.3	30.9	36.2	40.4
	FLORES	49.0	49.7	<u>49.9</u>	50.4	49.7	49.0	49.5	49.6	48.9	49.5
Indic-En	ALT	<u>65.2</u>	65.3	64.8	65.1	65.2	65.0	65.0	65.2	64.9	65.0
	IN22-CONV	33.8	40.1	52.6	57.1	38.4	47.3	<u>55.0</u>	32.0	46.3	54.6
	FLORES	61.2	61.4	60.0	61.0	61.2	59.6	60.9	<u>61.3</u>	59.5	61.1

Table 6: IT2-Dist-200M ChrF++ scores on sentence-level (top) and document-level (bottom) evaluation benchmarks. The best score is **bolded** and second best is underlined.

Type	Category	Classification	Span Text	Explanation
Fluency	Awkward	MINOR	please sit down	Slightly awkward phrasing. 'Please have your seat' is more natural.
Fluency	Inconsistency	MINOR	on six months, on three months, or every month	Inconsistent phrasing in intervals. Use 'half-yearly, quarterly or monthly intervals.'
Accuracy	Mistranslation	MAJOR	And I also know its benefits.	Incorrectly translated. Should be 'I want to know the benefits.'
Accuracy	Omission	MAJOR	You can opt for that.	Omitted in the translation.
Accuracy	Mistranslation	MAJOR	Please write my phone number!	Redundant repetition at the end which was not present in the source.
Accuracy	Mistranslation	CRITICAL	after the completion of the first five years of the policy	Misleading regarding the conditions of the death benefit. Should align with 'On death after completion of five policy years but before the date of maturity...'

Table 7: ROPE Translation Errors with Classification

Type	Category	Classification	Span Text	Explanation
Fluency	Grammar	MINOR	The entire information about the process	Awkward sentence structure. <i>'The information about the process can be'</i> would be more fluent.
Accuracy	Mistranslation	MAJOR	Your name and age are Devang Saikia. I am 39 years old	This part is translated inaccurately. The source asks for name and age, but the translation incorrectly states them directly.
Accuracy	Omission	MAJOR	You can opt for that.	Translation omits <i>'What is the minimum basic sum assured in this policy?'</i> and other questions related to policy details.
Accuracy	Addition	MAJOR	What is the minimum age? What is the minimum age? What is the minimum age?	Text is repeated unnecessarily, creating confusion.
Accuracy	Mistranslation	MAJOR	You will have to surrender the policy regularly	Mistranslation; the context is about having completed certain years, not regular surrender.
Accuracy	Mistranslation	CRITICAL	premium annually on the death of the policy	Incorrect context; it should refer to payment intervals, not related to death of the policy.

Table 8: ALIBI Translation Errors with Classification

Type	Category	Classification	Span Text	Explanation
Accuracy	Mistranslation	CRITICAL	Thank you for giving my name and age.	Incorrect translation. Should be <i>'Thank you for letting us know your name and age.'</i>
Accuracy	Mistranslation	CRITICAL	What will be the premiums paid during the first five years of the policy?	Incorrect translation. Should be <i>'On death during the first five years, sum assured on death shall be payable.'</i>
Accuracy	Mistranslation	CRITICAL	But do I want to know the full amount of the policy before the end of the year?	Incorrect translation. Should be <i>'Can I surrender the policy before the maturity date? What will be the surrender value?'</i>
Accuracy	Mistranslation	CRITICAL	what will be the premiums paid before the end of the policy?	Incorrect translation. Should refer to death benefits, not about premiums.
Accuracy	Mistranslation	CRITICAL	I can give you the information on the phone.	Text unrelated to the source and incorrectly translated.
Accuracy	Mistranslation	CRITICAL	If it is less than three months old, then it will be enough to pay the premiums.	Incorrect translation. Policy is misunderstood. Should reflect surrender conditions.

Table 9: SINE Translation Errors with Classification