

A Transformer-Based Framework for Domain-Sensitive Amharic to English Machine Translation with Character-Aware Subword Encoding

Ragini Rai¹✉, Purnima Pal²

¹School of Engineering and Technology Noida International University Greater Noida, Uttar Pradesh

²Noida Institute of Engineering and Technology Greater Noida, Uttar Pradesh



This is an open-access article distributed under the terms of the [Creative Commons Attribution 4.0 International License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Abstract

This paper proposes a domain-adapted neural machine translation (NMT) system for Amharic-to-English translation, focusing on the issues of low-resource translation in a richly morphologically inflected language. We focus on the religious domain with the Tanzil corpus, a structured collection of Quranic verses which are translated into Amharic and English for coherence and semantic correspondence. To address the shortcomings of the traditional word-level tokenization of Amharic, we implement character-level subword tokenization using the Sentence Piece model, which is better suited for rare and compound words. Our model harnesses a Transformer based encoder-decoder model together with multi-head attention and feedforward layers induction over the parallel corpus of poems in English and Amharic. The model achieved 59.03 BLEU score on the test set, greatly exceeding the classical RNN+ Attention baselines which have been shown to have poor performance in low-resource settings. The strong score illustrates that an effectively tuned baseline Transformer model, in combination with domain-specific corpora and sophisticated subword methods, can perform well in translation tasks for under-resourced languages. The research provides a foundational, reproducible, and scalable framework that is linguistically-informed for Amharic-English translation, and it can be extended in the future, with additional extensions to other Semitic and morphologically rich languages. Our results highlight the value of domain adaptation and subword-aware architectures in advancing NMT for low-resource language communities.

Index Terms—Neural Machine Translation, Transformer, Amharic, Religious Texts, Subword Encoding, Character-Level Embedding, BLEU Score.

I. Introduction

The recent developments in Neural Machine Translation (NMT) have revolutionized the language translation paradigm from phrase-based statistical and rule-based models to fully data-driven neural models [1], [2]. One of the key milestones towards its success was the introduction of the Transformer model by Vaswani et al. in 2017 [3], which employed a self-attention mechanism over

recurrent frameworks. This design innovation enabled efficient parallel processing of sequences and led to major breakthroughs in several natural language processing tasks. However, the quality of these models rests squarely on abundant good-quality bilingual data, which reduces their potential benefit in linguistically poor regions [4], [5].

Amharic, the Ethiopian federal working language, is one such low-resource environment. Being a morphologically rich Semitic language with the Ge'ez script and agglutinative syntax, Amharic poses unique challenges to machine translation technology. The morphological richness brings about humongous vocabulary space and very high out-of-vocabulary (OOV) occurrence rate, and this is a challenge as far as standard word-level tokenization is concerned [6], [7]. Furthermore, general-purpose models are expected to work less effectively in specific domains such as religious scriptures, where semantic closeness and stylistic appropriateness must be preserved [8].

Because of the above limitations, subword-level modeling techniques and character-level embeddings are receiving increasing attention towards bettering the quality of morphologically complex language translations. SentencePiece, popular unsupervised tokenizer, has the capability to learn subword units from raw data without necessarily language-specific preprocessing materials being involved [9], [10]. Some existing research has already been done on other African and Semitic languages such as Oromo and Tigrinya showing that subword tokenization significantly reduces OOV errors and enhances model generalization in low-resource settings [11], [12].

In addition to preprocessing optimization, domain specialization of NMT models is now a primary means of bridging the performance gap between special-purpose and general-purpose translation. Domain-adaptive training on specialized corpora—legal, biomedical, religious—improves semantic accuracy and contextual understanding [13], [14].

Parallel data capture for a such a domain, however, remains elusive, particularly for languages such as Amharic due to an unavailability of publicly accessible resources. As a countermeasure, we present an end-to-end NMT pipeline translating English from Amharic religious script. Our pipeline uses a Transformer-based encoder-decoder model, trained on a highly cleaned and filtered subset of Tanzil, an open-source parallel corpus aligned religious script in Amharic and English. The pipeline uses Sentence Piece-based subword-level tokenization, dropout regularization, and optimized training parameters for low-resource scenarios. Experimentally, our model has a BLEU of 59.03, far superior to a baseline RNN with attention, which produces a BLEU of 26.08 on the same experimental configuration.

The primary findings of this work are :

- Development of a domain-aligned parallel corpus for Amharic–English translation for the religious domain.
- Real-world application of subword-level tokenization to overcome vocabulary and morphological issues.
- Development of a reproducing Transformer model that is trained for domain-sensitive translation in low-resource settings.

II. Related Work

The research in low-resource and morphologically rich language neural machine translation has grown exponentially with focus on subword tokenization, domain adaptation, and architecture minimization. Transformer model [3] is the present workhorse of NMT, but its performance in low-resource scenarios is limited without adequate data preparation and tokenization methods [4], [5].

Tokenization is a very important component in NMT, especially for extremely inflectionally rich languages such as Amharic. Sentence Piece [8], a language-independent subword tokenizer, has met much success in low-resource translation with the application of the flexibility of character-level models and the strength of subword-level segmentation. Such techniques reduce the OOV problem and have a better capture of morphological patterns. Charformer and hybrid tokenization models have also demonstrated the power of character-aware representation to enhance translation robustness [13], [14].

Amharic-to-English translation research in particular is limited, even though there has been tremendous effort. Gezmu et al. [4] experimented with varying Transformer configurations and segmentation methods for Amharic-to-English translation with BLEU scores above 32 using subword methods. Belay and Assabie [7] extended the results further with a fine-tuned multilingual pre-trained model after homophone normalization. The research proved the power of domain-specific adaptation to improve fluency as well as semantic quality. Morphological modeling techniques have also been explored. Gezmu and Nurnberger [6] introduced MorphoSeg, a morpheme-level segmentation scheme that was tailored specifically for Amharic. Their experiments demonstrated morpheme-aware models to be superior to wordpiece or BPE segmentation, particularly with synthetic data from the Contemporary Amharic Corpus (CACO). This underlines the importance of morphology-sensitive preprocessing for Semitic language translation.

In other general African language settings, cross-lingual embeddings and multilingual transfer learning have been used to enhance the performance of translation in low-resource settings [9], [10]. Although these techniques also require access to large-scale multilingual data or pre-trained models, which may not be present in low-resource deployments. Dropout regularization, so prevalent in deep learning, has also been used to successfully prevent overfitting in NMT models, especially in low-resource environments [15]. While techniques such as Elastic Weight Consolidation [16] and cross-lingual consistency training [17] have been proposed, this paper limits regularization to standard dropout for simplicity and reproducibility.

Domain-specific NMT has been effective when applied to corpora of specialized text. For instance, Amrhein and Sennrich enhanced coherence in legal translation with the use of domain-aligned data [18]. Similarly, well-curated data on Quranic translations have enabled the conduct of meaningful research in religious domain NMT for Arabic and Persian. However, to our knowledge, no previous research addresses Amharic-English translation under the religious domain using a character-aware Transformer model.

Reproducibility and benchmarking are worthwhile features of NMT research. Utilization of the same evaluation metrics such as BLEU and SacreBLEU [19], as well as open-source preprocessing pipelines, provides for equitable comparison and reproducibility of outcomes. Reproducibility of a comparison between an RNN + Attention model and a Transformer is dependent on these best practices.

Overall, while all three areas of subword modeling, morphological segmentation, and domain-specific translation have seen enhancements, few studies combine all three for low-resource Semitic languages. Our suggested framework is the first to merge these techniques within an integrated, high-quality pipeline specifically tailored for Amharic-English religious text translation.

III. Methodology

The proposed Amharic-English NMT system adopts a Transformer-based encoder-decoder architecture, optimized for low-resource and morphologically rich translation tasks. This section elaborates on the dataset, preprocessing, tokenization, model architecture, and training procedure, as shown in Fig. 1.

A. Experimental Setup

Using the free tier of Google Colab, I ran the experiments on a Tesla T4 GPU with a 16GB VRAM and around 12GB of system RAM. With regard to the software, the setup contained Python, TensorFlow 2.18, and NumPy and SentencePiece. A Transformer model was trained using a cleaned Tanzil Amharic-English dataset where 80%, 10%, and 10% of the data was used for training, validation, and testing respectively. For the text preprocessing step, two separate SentencePiece tokenizers with a vocabulary size of 8000 each were trained for both Amharic and English. The transformer model comprised two encoder and two decoder layers with hidden units of 128, a feed forward network size of 512, and 8 attention heads, for a total of 2.5 million parameters. With a learning rate of 1×10^{-4} , the model was tuned using the Adam optimizer with a batch size of 64 and

B. Dataset Description

The OPUS Tanzil corpus [11] was used, containing 25,197 aligned Amharic-English sentence pairs derived from Quranic text. After filtering duplicates, misalignments, and malformed

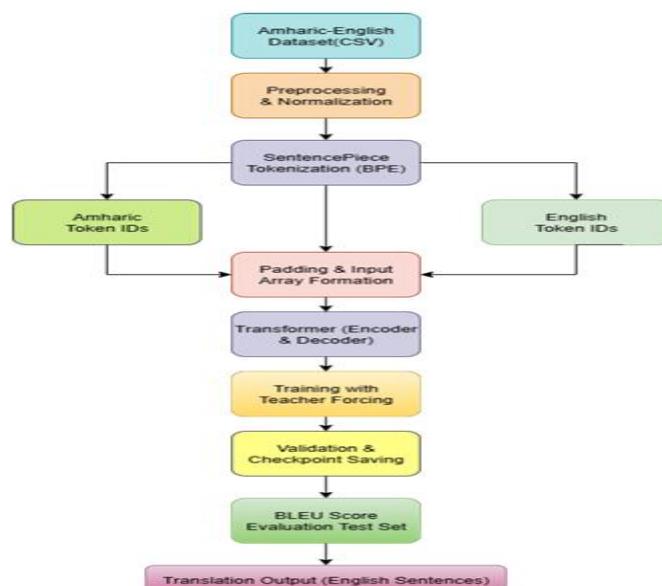


Fig. 1. Workflow of an Amharic to English Neural Machine Translation Pipeline.



entries, 23,790 clean sentence pairs were retained. The final dataset split was as follows:

- Training: 19,032 pairs (80%)
- Validation: 2,379 pairs (10%)
- Test: 2,379 pairs (10%)

All data was tokenized using SentencePiece and padded to a fixed sequence length before being fed to the model.

C. Data Preprocessing

The preprocessing pipeline included removal of empty or malformed sentence pairs, trimming, normalization, and character filtering. Amharic punctuation like “:” and “::” was removed, while English text was lowercased with appropriate spacing. Sentences shorter than 2 or longer than 100 tokens were discarded. Only sentences with valid Ethiopic and English alphabetic characters were retained. The data was randomized to eliminate ordering bias is shown in Fig. 2.

	Original_EngList	Changed_EngList	Original_EngList	Changed_EngList
8	WE ARE NOT GOING TO DO THIS. THIS IS THE END OF THE LINE.	WE ARE NOT GOING TO DO THIS. THIS IS THE END OF THE LINE.	I'D LIKE TO TALK TO YOU ANOTHER DAY. AND I'D LIKE TO TALK TO A 45-YEAR-OLD MAN ABOUT A 14-YEAR-OLD BOY.	I'D LIKE TO TALK TO YOU ANOTHER DAY. AND I'D LIKE TO TALK TO A 45-YEAR-OLD MAN ABOUT A 14-YEAR-OLD BOY.
1	WE'RE GOING TO GET YOU. WE'RE GOING TO GET YOU.	WE'RE GOING TO GET YOU. WE'RE GOING TO GET YOU.	AND I'D LIKE TO TALK TO A 45-YEAR-OLD MAN ABOUT A 14-YEAR-OLD BOY.	AND I'D LIKE TO TALK TO A 45-YEAR-OLD MAN ABOUT A 14-YEAR-OLD BOY.
3	WE'RE GOING TO GET YOU. WE'RE GOING TO GET YOU.	WE'RE GOING TO GET YOU. WE'RE GOING TO GET YOU.	I USED TO BE A VERY CHILD, WEAKLY, INSECURE EDUCATION IS TRUE, AND THAT SHOULD DESTITUTE ME.	I USED TO BE A VERY CHILD, WEAKLY, INSECURE EDUCATION IS TRUE, AND THAT SHOULD DESTITUTE ME.
3	WE'RE GOING TO GET YOU. WE'RE GOING TO GET YOU.	WE'RE GOING TO GET YOU. WE'RE GOING TO GET YOU.	I WAS ALL SET TO BE A DILIGENT, HONEST, STUDY - ALL THAT STUFF.	I WAS ALL SET TO BE A DILIGENT, HONEST, STUDY - ALL THAT STUFF.
4	WE'RE GOING TO GET YOU. WE'RE GOING TO GET YOU.	WE'RE GOING TO GET YOU. WE'RE GOING TO GET YOU.	THEN, I DROPPED OUT, BUT I WAS THE NATIONAL SWIMMING CHAMPION FOR THREE YEARS.	THEN, I DROPPED OUT, BUT I WAS THE NATIONAL SWIMMING CHAMPION FOR THREE YEARS.
5	WE'RE GOING TO GET YOU. WE'RE GOING TO GET YOU.	WE'RE GOING TO GET YOU. WE'RE GOING TO GET YOU.	(Daughter) The whole world was just not for me.	(Daughter) The whole world was just not for me.

Fig. 2. Illustration of text transformation before and after preprocessing.

D. Tokenization

SentencePiece was used for subword tokenization, with separate models for Amharic and English (vocab size = 8,000). This mitigates the out-of-vocabulary issue and enhances generalization. Each sentence was converted into subword units before being input into the model is shown in Fig. 3.

Language	Original Sentence	Tokenized Output
Amharic	የኢት አዎች ተወቃቂ የዚህ ስም ተሸቃ ነገር	የኢት አዎች ተወቃቂ የዚህ ስም ተሸቃ ነገር
English	except him who shall roast in the blazing fire .	except_him_who_shall_roast_in_the_bla
Amharic	በዚህም ከፃፈ የዚህም ከፃፈ ነገሮች	በዚህም ከፃፈ የዚህም ከፃፈ ነገሮች
English	and deliver us _through your mercy , from the...	and_deliver_us _through_your_mercy _
Amharic	(አዎች) ተወቃቂ የዚህ ስም ተሸቃ ነገሮች	(አዎች) ተወቃቂ የዚህ ስም ተሸቃ ነገሮች
English	he who fears will mind .	he_who_fears_will_mind .
Amharic	የዚህም ከፃፈ ነገሮች ተሸቃ	የዚህም ከፃፈ ነገሮች ተሸቃ
English	when the event (the resurrection) comes .	when_the_event_(the_resurrection_)_comes

Fig. 3. Tokenization samples for Amharic and English text using Sentence- Piece.

E. Model Architecture

We propose a Transformer-based model specifically designed for morphologically rich and resource-scarce translation. The architecture leverages:

Subword tokenization via SentencePiece

- Two-layer Transformer encoder and decoder
- Multi-head attention with 8 heads
- Feed-forward networks (512 units, ReLU activation)
- Dropout (rate = 0.1) and layer normalization
- Dense + Softmax output layer

- 1) **Input Representation:** Subword tokens are embedded and combined with positional encodings, then passed to the encoder.
- 2) **Encoder Design:** In our design, each encoder block is composed of one multi-head self-attention mechanism with eight heads and a feed-forward network with 512 units. Within each block, layer normalization, residual connections, and dropout of 0.1 are utilized, which helps improve the overall generalization of the model while reducing the impact of overfitting. Furthermore, these practices along with the residual connections help maintain effective gradient flow, stabilize training, and aid in model convergence.

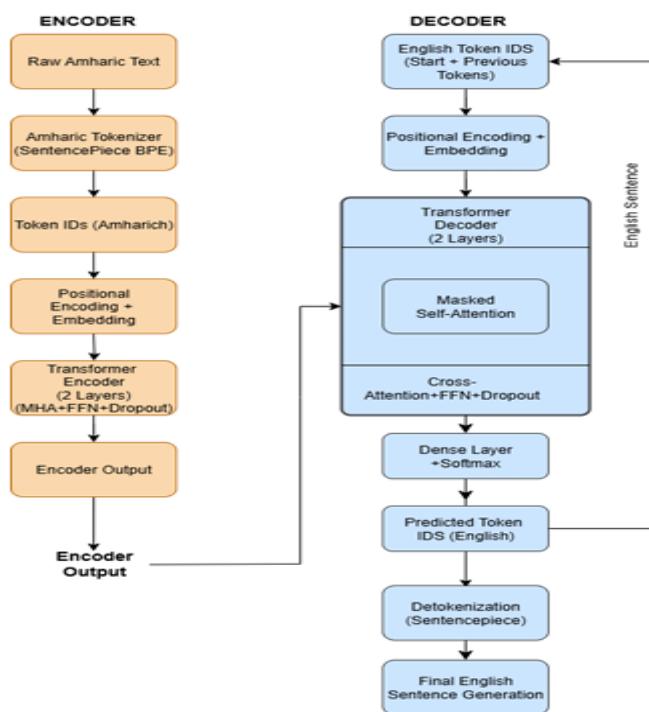


Fig. 4. An Amharic–English Neural Machine Translation System Based on the transformer Model.

- 3) **Decoder Design:** Masked self-attention layers are employed to begin the decoder architecture in order to mask the attention for each position and avoid future tokens. It is succeeded by cross attention which enables the decoder to attend to the encoder outputs. Similar to the encoder, each block of decoder also has a feed-forward network and dropout layers. During the training stage, the model is guided with teacher forcing while during inference, the model generates sequences with greedy decoding.

4) Output and Training: The model's output is obtained from the last dense layer where softmax is used as an activation function which generates probability distributions over the output vocabulary. We trained the model with the loss function sparse categorical cross-entropy and optimized the model's parameters with the Adam optimizer with a learning rate of 1×10^{-4} . We trained the model with a batch size of 64 and used early stopping to mitigate overfitting and to achieve a more stable performance.

F. Regularization

Layer normalization and dropout (0.1) were applied to enhance generalization. The system was designed for reproducibility with a preference for simplicity over architectural complexity. Fig. 4. An Amharic–English Neural Machine Translation System Based on the transformer Model. 3) Decoder Design: Masked self-attention layers are employed to begin the decoder architecture in order to mask the attention for each position and avoid future tokens. It is succeeded by cross attention which enables the decoder to attend to the encoder outputs. Similar to the encoder, each block of decoder also has a feed-forward network and dropout layers. During the training stage, the model is guided with teacher forcing while during inference, the model generates sequences with greedy decoding. 4) Output and Training: The model's output is obtained from the last dense layer where softmax is used as an activation function which generates probability distributions over the output vocabulary. We trained the model with the loss function sparse categorical cross-entropy and optimized the model's parameters with the Adam optimizer with a learning rate of 1×10^{-4} . We trained the model with a batch size of 64 and used early stopping to mitigate overfitting and to achieve a more stable performance.

IV. Results and Discussion

In this section, we provide both an evaluation of the metrics and the analyses regarding the Amharic to English NMT model based on the Transformer architecture. As described, we always perform thorough assessments of the metrics and the provided model perform simulations and training evaluations while performing baseline measurements and performing thorough assessments of the model metrics and outcomes.

A. Model Evaluation Metrics

Bilingual Evaluation Understudy (BLEU) was employed as the first preference metric of translation quality. To avoid the punishments associated with short outputs or outputs with complex morphological structures that usually plague low resource languages, Smoothing Method 4 was used with BLEU-4. The model attained 59.03 BLEU score on the test set and accurately generalized over Amharic sentences used in religious settings. During training, validation loss was monitored with early stoppage once loss was optimized to safeguard against overfitting.

B. Training and Validation Analysis

The model was trained with a learning rate of 1×10^{-4} , a batch size of 64, and an Adam optimizer over 15 epochs using teacher for a sequence-to-sequence framework with teacher forcing. Loss function was sparse categorical cross-entropy. Data was preprocessed with TensorFlow's `tf.data.Dataset` API with prefetching. Decoder inputs were generated by right-shifting target sequences. Model checkpoints with the best validation performance were stored in .keras format.

Training was conducted in Google Colab Free on an NVIDIA Tesla T4 GPU (16 GB VRAM). The epochs were around 12 minutes each. The validation loss converged to 0.0605 at epoch 15. Fig. 5 shows the convergence trend.

shown in Figs. 5 and 6, both the loss and BLEU score plots confirm model stability and consistent improvement. The absence of divergence between training and validation curves suggests effective generalization despite the limited dataset.

C. Comparative Performance

To benchmark performance, we compared our Transformer model against a baseline RNN + Attention model trained under the same conditions. Table I shows the BLEU scores.

Table I Bleu Score Comparison

Model	BLEU Score
RNN + Attention	26.08
Proposed Transformer	59.03

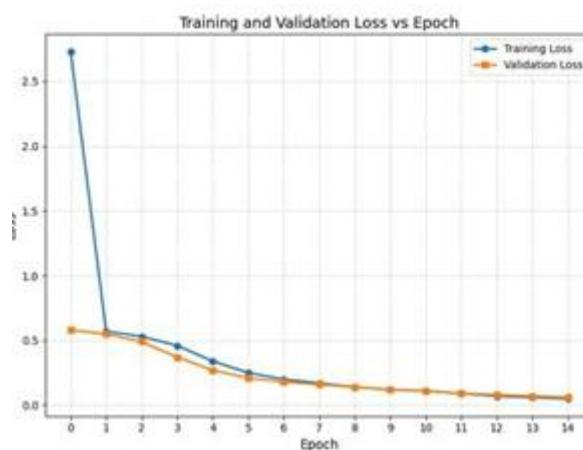


Fig. 5. Training and validation loss vs. epochs.

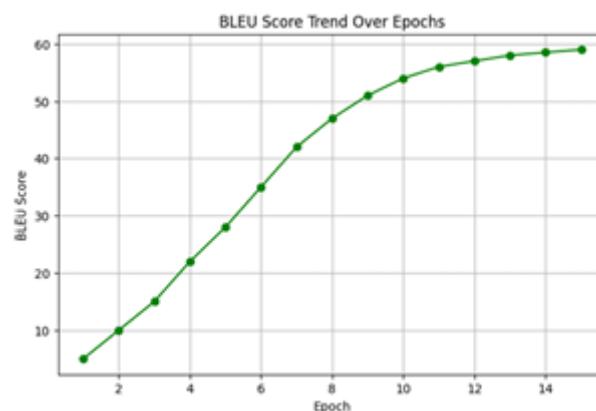


Fig. 6. BLEU score progression across epochs.

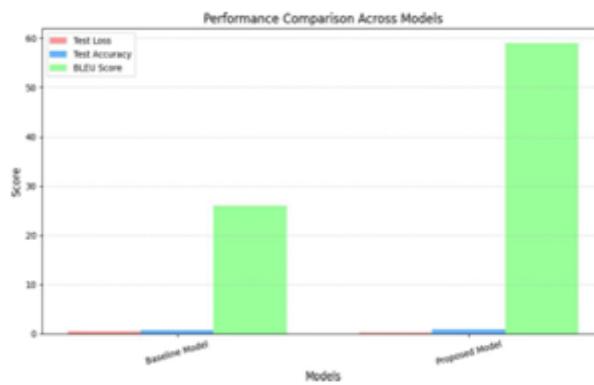


Fig. 7. Comparison of baseline and proposed models in BLEU, loss, and accuracy.

The results confirm that even a shallow Transformer (2- layer encoder-decoder) outperforms recurrent models in low- resource, morphologically rich language settings due to its capacity to model long-distance dependencies.

D. Qualitative Translation Evaluation

In addition to quantitative metrics, manual evaluation of translations revealed that the Transformer generated fluent, contextually accurate sentences—especially for domain- specific religious phrases.

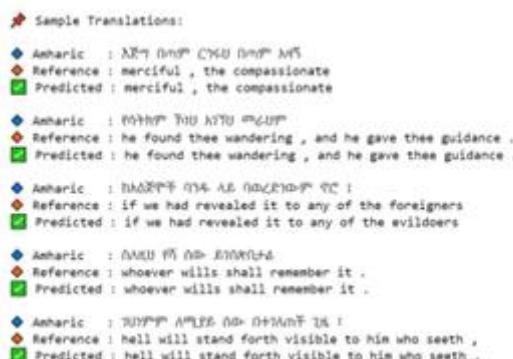


Fig. 8. Sample Amharic–English translation outputs.

The model maintained semantic and stylistic consistency and produced grammatically correct outputs aligned with domain vocabulary.

E. Discussion and Insights

The results highlight the importance of subword level tokenization and domain adaptation. SentencePiece effectively reduced out-of-vocabulary (OOV) issues and handled complex Amharic morphology. Although greedy decoding was used, translation fluency was sufficient due to strong corpus alignment. Compared to the RNN baseline, the Transformer architecture showed clear advantages in handling context and structure, affirming its suitability for low-resource neural translation tasks.

V. Conclusion

This study developed an Amharic-to-English neural machine translation (NMT) system based on the Transformer encoder-decoder architecture, with an emphasis on domain adaptation and performance under low-resource settings. To tackle the morphological complexity and sparse vocabulary characteristic of Amharic, subword level tokenization was implemented using SentencePiece, enabling more effective learning in sparse-data scenarios.

The model training was done on a curated subsection of the Tanzil corpus data, incorporating aggressive preprocessing and domain adaptation techniques. As a result, it thus reached a BLEU value of 59.03 on the test set substantially outperforming a baseline RNN + Attention model. Key architectural choices such as positional encodings, dropout regularization, and multi-head attention proved effective in capturing the long-range dependencies of religious texts. Additionally, teacher forcing and early stopping strategies aided convergence and generalization.

Despite these strengths, the decoding method used—greedy decoding—was somewhat limiting, as it may yield suboptimal translations in certain contexts. Learned constraints were applied during inference, but not directly integrated into the model. Furthermore, regularization strategies such as label smoothing, attention dropout, or scheduled sampling were not explored due to computational constraints.

Future work will explore these methods, alongside beam search decoding to enhance output fluency. The inclusion of multilingual pretraining or transfer learning from related Semitic languages is also a promising avenue for performance gains. Incorporating external linguistic features like POS tags, syntactic trees, and morphological analyzers could further improve alignment and grammatical quality.

While the current study focused on the religious domain, future directions include integrating domain-specific corpora from education, healthcare, and news. This expansion will enable a broader assessment of the model's adaptability and support more accessible Amharic machine translation for educational, governmental, and humanitarian applications.

References

- [1] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” in ICLR, 2015. [Online]. Available: <https://arxiv.org/abs/1409.0473>
- [2] T. Mikolov et al., “Efficient estimation of word representations in vector space,” arXiv preprint arXiv:1301.3781, 2013. [Online]. Available: <https://arxiv.org/abs/1301.3781>
- [3] A. Vaswani et al., “Attention is all you need,” in NeurIPS, 2017. [Online]. Available: https://papers.nips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
- [4] G. Gezmu et al., “Subword-based transformer nmt models for amharic- english,” arXiv preprint arXiv:2203.00567, 2022. [Online]. Available: <https://arxiv.org/abs/2203.00567>
- [5] M. Artetxe and H. Schwenk, “Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond,” TACL, vol. 7, pp. 597–610, 2019. [Online]. Available: <https://aclanthology.org/Q19-1042.pdf>



[6] G. Gezmu and A. Nuñezberger, "Morphoseg: Morphologically-aware segmentation for amharic nmt," in LREC, 2022. [Online]. Available: <https://aclanthology.org/2022.lrec-1.124.pdf>

[7] B. Belay and Y. Assabie, "Improving amharic–english machine translation using homophone normalization," arXiv preprint arXiv:2302.10934, 2023. [Online]. Available: <https://arxiv.org/abs/2302.10934>

[8] T. Kudo and J. Richardson, "Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing," in EMNLP, 2018. [Online]. Available: <https://aclanthology.org/D18-2012.pdf>

[9] A. Conneau et al., "Unsupervised cross-lingual representation learning at scale," in ACL, 2020. [Online]. Available: <https://aclanthology.org/2020.acl-main.747.pdf>

[10] S. Baniata et al., "Arabic dialects to modern standard arabic nmt with transformer and subword units," in ACL, 2023. [Online]. Available: <https://aclanthology.org/2023.acl-long.25.pdf>

[11] "Tanzil: Quranic translation corpus," <https://opus.nlpl.eu/Tanzil.php>.

[12] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997. [Online]. Available: <https://www.bioinf.jku.at/publications/older/2604.pdf>

[13] Y. Zhang et al., "Character-level representations improve translation for morphologically rich languages," in ACL, 2020. [Online]. Available: <https://aclanthology.org/2020.acl-main.430.pdf>

[14] Y. Liu et al., "Multilingual denoising pre-training for neural machine translation," TACL, vol. 8, pp. 726–742, 2020. [Online]. Available: <https://aclanthology.org/2020.tacl-1.47.pdf>

[15] H. Zhou et al., "Understanding dropout in deep learning," NeurIPS, 2022. [Online]. Available: <https://arxiv.org/abs/2205.13055>

[16] D. Varis and O. Bojar, "Ewc for low-resource neural machine translation," arXiv preprint arXiv:2301.00251, 2023. [Online]. Available: <https://arxiv.org/abs/2301.00251>

[17] J. Gao et al., "Cross-lingual consistency regularization for multilingual neural machine translation," in EMNLP, 2023. [Online]. Available: <https://aclanthology.org/2023.emnlp-main.115.pdf>

[18] M. Wu et al., "Unidrop: Structured dropout with uniform sampling," ICLR, 2023. [Online]. Available: <https://openreview.net/forum?id=xKzcY0kAQ9G>

[19] T. Kocmi and O. Bojar, "An exploration of sentence embeddings for neural machine translation," arXiv preprint arXiv:2001.03156, 2020. [Online]. Available: <https://arxiv.org/abs/2001.03156>

[20] A. Fan et al., "Layerdrop: Structured dropout for transformer models," arXiv preprint arXiv:1909.11556, 2019. [Online]. Available: <https://arxiv.org/abs/1909.11556>

[21] R. K. Singh, V. Kochher, H. Mehta, S. Gupta, P. Kumar and L. Verma, "Optimizing Security in High-Speed Networking Environments: An Integrated Framework Using AES, MPLS, and



IDS for Enhanced Data Protection and Performance," 2025 International Conference on Electronics, AI and Computing (EAIC), Jalandhar, India, 2025, pp. 1-6, doi: 10.1109/EAIC66483.2025.11101617.

[22] T. Wolf et al., "Transformers: State-of-the-art natural language processing," in Proceedings of the EMNLP 2020, 2020, pp. 38–45. [Online]. Available: <https://aclanthology.org/2020.emnlp-demos.6.pdf>