AGE: Amharic, Ge'ez and English Parallel Dataset

Henok Biadglign Ademtew
Ethiopian AI Institute
henokb2124@gmail.com

Mikiyas Girma Birbo
Maharishi International University
mbirbo@miu.edu

Abstract

African languages are not well-represented in Natural Language Processing (NLP). The main reason is a lack of resources for training models. Low-resource languages, such as Amharic and Ge'ez, cannot benefit from modern NLP methods because of the lack of high-quality datasets. This paper presents AGE, an open-source tripartite alignment of Amharic, Ge'ez, and English parallel dataset. Additionally, we introduced a novel, 1,000 Ge'ez-centered sentences sourced from areas such as news and novels. Furthermore, we developed a model from a multilingual pre-trained language model, which brings 12.29 and 30.66 for English-Ge'ez and Ge'ez to English, respectively, and 9.39 and 12.29 for Amharic-Ge'ez and Ge'ez-Amharic respectively.

1 Introduction

Language is the foundation on which communication rests, allowing us to share ideas and interact with one another (Adebara and Abdul-Mageed, 2022). One of the NLP applications is machine translation (MT), which helps facilitate human-machine and human-human communications (Abate et al., 2019). Data availability is one of the criteria to categorize one language as a high or low-resource language (Ranathunga et al., 2021). Recently, interest in low-resource MT has been increasing both within the MT research community (Haddow et al., 2022), as well as in native speaker communities (Nekoto et al., 2020). Modern NLP technologies, however, have primarily been developed in Western societies (Adebara and Abdul-Mageed, 2022). The current state-of-the-art(SOTA) MT models were trained on enormous datasets, including sentences in a source language and their corresponding target language translations, which is the most effective of these systems(Tonja et al., 2023).

To date, there is no publicly available MT system for Ge'ez language and it's not represented in the commercial MT systems such as Lesan[1], Google Translate[2], Microsoft Translator[3], and Yandex Translate[4]. It is also not included in large-scale pre-trained multilingual models like NLLB(Team et al., 2022), MT5(Xue et al., 2021), ByT5(Xue et al., 2022), and M2M-100(Fan et al., 2020). This makes it harder for people to learn and use the language. So, by focusing on Ge'ez, an ancient language with profound cultural and religious significance in Ethiopia, alongside Amharic, the country's official working language, and English, a global lingua franca, this dataset aims to bridge the gap between historical linguistic treasures and modern technological advancements. Through this work, we aim to provide a dataset for researchers and technologists aiming to advance machine translation capabilities, linguistic studies, and cultural preservation efforts. Furthermore, by enriching the available resources for Ge'ez, we contribute to the broader goal of advancing low-resource languages.

2 Related work

One of the major challenges in developing MT models for Ge'ez is the lack of public data. There were attempts to compile parallel corpora for Ge'ez to English and Ge'ez to Amharic MT tasks, but the development was unsatisfactory. (Mulugeta, 2015) researched Ge'ez-Amharic MT using SMT. He used IRSTLM for language modeling. The research was conducted on a dataset comprising 12,840 parallel Amharic-Ge'ez sentences, achieving an average translation accuracy with a BLEU score of 8.26 based on 10-fold cross-validation. (Abate et al., 2019) is the only publicly available dataset that was part of an effort to train Statistical

[1]https://lesan.ai
[2]http://translate.google.com/
[3]https://www.microsoft.com/en-us/translator/
[4]https://translate.yandex.com

Proceedings of the The Seventh Workshop on Technologies for Machine Translation of Low-Resource Languages (LoResMT 2024), pages 139–145

Machine Translation(SMT) for English-Ethiopian Languages and made 11,663 Ge'ez-English parallel sentences. They achieved English-Ge'ez and Ge'ez English translations with a BLEU score of 6.67 and 18.01, respectively. Using deep learning approaches, (Getachew and Yayeh, 2023) have explored bidirectional NMT from Ge'ez to English. The experiment was conducted by leveraging 16,569 parallel sentences from the Holy Bible and Battle of Saints and manually preparing daily conversational sentences. The results indicated that the transformer (**?**) model achieved BLEU scores of 27.19 for English to Ge'ez translation and 29.39 for Ge'ez to English translation. Another work by (Tegenaw et al., 2023) used NMT and transformers and attempted three experiments that used a pre-trained masked language model (MLM) utilizing a monolingual dataset of 33,004 sentences for each language. The experiments involved a parallel corpus for supervised learning without a pre-trained model and fine-tuning a pre-trained MLM with a bilingual dataset. The outcomes were evaluated using the BLEU score, achieving 31.65 in the second and 33.02 in the third experiment. Another recent work by (Wassie, 2023) improved translation by 4 BLEU using a new model but faced challenges with NLLB-200 for Ge'ez due to insufficient data. They also experimented with GPT-3.5's trial, which resulted in a 9.2 BLEU score, underperforming compared to their model's 15.2. They also highlighted the difficulties of training Ge'ez MT models.

A recurring issue noted in these experiments is the absence of data sharing with the public domain. As shown in Table 1, there is a lack of open-sourcing data and models, a significant obstacle to the representation of Ge'ez in NLP. This also indicates that despite extensive research in various studies, it's important for a unified effort among researchers to create and distribute resources open to the public. The collaborative effort would support further progress in expanding resources for the Ge'ez language.

3 Ge'ez Language

Ge'ez (ethiopge'ez), which is also known as Ethiopic, is one of the oldest Semitic languages (Tareke et al., 2002) and its alphabets is among the oldest alphabets still in use in the world of today. Furthermore, the Ge'ez language is among the four languages (Sabaean, Greek, and Arabic) that have been and continue to be used for ancient inscriptional arts. Ge'ez is currently not an actively spoken language nor a native tongue of any people. Its use is limited to the liturgical language of the Ethiopian Orthodox Tewahedo, Eritrean Orthodox Tewahedo, Ethiopian Catholic, and Eritrean Catholic Christians(Molla and Tabor, 2018). It is also used during prayer and at regularly scheduled public religious feast celebrations. The Bible dominates the literature, and it comprises the Deuterocanonical books. According to (Molla and Tabor, 2018), this language also has many medieval and early modern original texts. The majority of the essential works are correspondingly the literature of the Ethiopian Orthodox Tewahedo Church. These works include Christian Orthodox liturgy (service books, prayers, hymns), hagiographies, and a range of Patristic literature. Around 200 texts were written about home-grown Ethiopian saints from the fourteenth to the nineteenth century. The religious alignment of Ge'ez literature was due to traditional education being the obligation of priests and monks. More info about the alphabet on Ge'ez can be found in the appendix section.

4 Creation of the dataset

We introduce our newly Ge'ez-centered parallel dataset; **AGE** — **A**mharic, **G**e'ez , **E**nglish for machine translation.

4.1 Data Collection

Machine Translation (MT) necessitates using parallel sentences from source and target languages. We started by creating a novel parallel dataset comprising 1,000 sentence pairs. After cleaning, we extracted 17585 and 18676 sentence pairs for Amharic-Ge'ez and Ge'ez-English, respectively. The reason behind the number inconsistency is that some sources have either Amharic-Ge'ez or English-Ge'ez. For instance, with "Kufale" (The Book of Jubilees), our dataset comprised only sentence pairs in Ge'ez and English. The extracted pairs were collected from The Open Siddur Project[5], YouVersion[6], Ethiopic Bible[7], and Awde Mehret[8].

4.2 Data pre-processing

Our dataset, sourced from diverse sources, exhibited significant textual inconsistencies. We found

[5]https://opensiddur.org/
[6]https://www.bible.com/
[7]https://www.ethiopicbible.com/
[8]https://awdemehret.org/

Language	Author(s)	Sentences	Dataset	Model	Technique
Amharic, Ge'ez	(Mulugeta, 2015)	12, 840	✗	✗	SMT
Amharic, Ge'ez	(Kassa, 2018)	13,833	✗	✗	SMT
Amharic, Ge'ez	(Abel, 2018)	976	✗	✗	SMT
Ge'ez, English	(Abate et al., 2019)	11,663	✓	✗	SMT
Ge'ez, English	(Getachew and Yayeh, 2023)	16,569	✗	✗	NMT
Amharic, Ge'ez	(Tegenaw et al., 2023)	33,004	✗	✗	NMT
Amahric, Ge'ez	(Wassie, 2023)	4,000	✗	✗	MNMT

Table 1: Summary of related works for Ge'ez: Sentences show the number of sentences used during the experiment. Dataset and Model show the availability of datasets and models in publicly accessible repositories, and Technique shows the method used to build models.

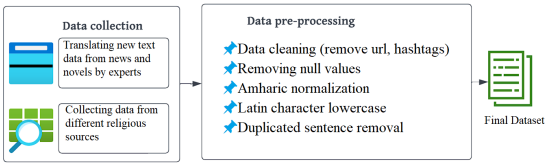

Figure 1: Data collection and pre-processing pipelines

Language Pair	Sentences	Token	Avg. Length
Amharic	17,584	17,585	13.35
Ge'ez		51,212	13.43
English	18,722	27,884	13.81
Ge'ez		54,160	13.43

Table 2: Overview of the Dataset Sizes and Characteristics

portions of the data excessively disordered and removed them from our collection. Figure 1 shows the general framework for the dataset development process. It had two primary tasks. The first task was data collection, which involved identifying the sources from which the tripartite parallel Amharic, Ge'ez, and English sentences were collected. The second task was translating a few collected sentences to Ge'ez. This task involved translators and reviewers. Three translators and three evaluators were assigned to handle a set of 1,000 sentences. We made an in-house tool to ease the translation and evaluation process, which significantly streamlined the entire workflow. We performed several preparatory actions to standardize all tokens in Amharic and English sentences gathered from multiple sources. These actions included cleaning the data (eliminating URLs, hashtags, and repeated sentences), normalizing Amharic homophone characters, and converting English characters to lowercase.

5 Baseline Experiments

As shown in table 1, prior research predominantly employed SMT(Josef and Ney, 2001), and a very few NMT using transformers (Vaswani et al., 2017). To extend these studies, we incorporated an approach by leveraging the NLLB-200 (Team et al., 2022), a pre-trained language model.

- **NLLB-200**: a sparsely gated 54B parameter Mixture-of-Experts(MoE) model. It has demonstrated SOTA results across many language pairs, improving the previous model's BLEU scores by 44%

Accessing the large NLLB-200 model requires a minimum of four 32GB GPUs just for inference, showcasing the need for significant computational resources. So, we used the NLLB-200 600M parameter variant, a dense transformer model distilled from NLLB-200 due to its much lower resource requirements, making it a more practical option for our computational constraints.

The work by (Adelani et al., 2022) to effectively adapt to large-scale pre-trained models and get improved performance suggests that these models have better capability for relatively smaller datasets. So, we split our dataset into TRAIN (80%), DEV (10%), and TEST split (10%). We fine-tune the model using the HuggingFace transformer tool(Wolf et al., 2020) with a learning rate 5e-5, a batch size of 4 per device, a maximum source length, and a maximum target length of 128, and a beam size of 10. All the experiments were performed on Google Colab Pro. Then, the quality of translation is assessed using the BLEU score(Papineni et al., 2002), a standard in the field for its objectivity and correlation with human judgment. Our baseline experiments focused on bidirectional translation tasks, Amharic-Ge'ez

Language pair	BLEU
Amharic-Ge'ez	9.39
Ge'ez-Amharic	12.29
English-Ge'ez	12.87
Ge'ez-English	30.66

Table 3: Baseline results of NLLB-200 600M

and English-Ge'ez translations, aiming to establish a foundational understanding of the NLLB-200 600M model's capabilities within the context of our dataset. Since our primary focus was developing machine translation for Ge'ez, we skipped training the model on bidirectional English-Amharic translation.

6 Results and Discussion

In this work, we adapted the NLLB-200 600M model to evaluate its performance in the Ge'ez language. Our results as shown in Table 3 reveal a clear gradient in BLEU score performance across various language pairs. For translations from Amharic to Ge'ez and vice versa, the model achieved BLEU scores of 9.03 and 12.26 for evaluation, with a slight increase in the prediction phase to 9.39 and 12.87, respectively. Our BLEU scores showed a dramatic increase in scores for the Ge'ez to English language pair. Notably, English translations demonstrated superior performance, with the Ge'ez to English pair achieving the highest scores of 30.35 in evaluation and 30.66 in prediction, indicating a robust model capability in this language direction.

The higher scores in recorded in translations involving English may be due to a combination of factors, including the richer linguistic resources available for English and the NLLB-200's pre-training, which includes 21.5 billion sentences in English (Team et al., 2022). The difference in scores between the language pairs involving Amharic and those involving English points to the challenges associated with being low-resource. The lower BLEU scores for Amharic to Ge'ez suggest inherent difficulties in capturing the nuances of Ge'ez, a Semitic language with complex morphology. (Tran et al., 2014) stated that translating into morphologically rich languages is a particularly difficult problem in machine translation due to the high degree of inflectional ambiguity in the target language, often only poorly captured by existing word translation models. On the other hand, better performance was registered for Ge'ez to English translation, which is an encouraging sign of the model's adaptability, especially considering that Ge'ez data wasn't included in NLLB-200's pretraining data. The model's success in this area showcases the potential of such systems when appropriately fine-tuned, even when working with languages traditionally underserved by NLP technologies. Finally, we will release our models and dataset for the public to use and expand on our work.

7 Conclusion and Future Works

This paper presents an attempt to prepare a standard parallel corpora for Ge'ez. One thousand newly translated sentences were gathered from nonreligious domains, and the rest text data was gathered from religious domains on the internet. Then, the data are further pre-processed and normalized to prepare a parallel dataset for the model training task. Using our dataset, we fintuned NLLB-200 model. The experimental results show that translating to and from English resulted in a better BLEU score than English to Ge'ez and Amharic to Ge'ez and vice versa. The abundance of English data in the pre-trained model and the morphological richness of Ethiopian languages significantly impact the model's performance during bidirectional training involving Ge'ez and Amharic and when these languages are target languages. To the best of our knowledge, this is the first ready-to-use Amharic, Ge'ez, English tripartite dataset. Our initiative to make the dataset and models open source will open doors for many researchers and developers. Future works include increasing both the quantity and diversity of the dataset. We also intend to incorporate the several Ge'ez data sources that are now absent from this dataset.

References

Solomon Teferra Abate, Michael Melese, Martha Yifiru Tachbelie, Million Meshesha, Solomon Atinafu, Wondwossen Mulugeta, Yaregal Assabie, Hafte Abera, Biniyam Ephrem, Tewodros Gebreselassie, et al. 2019. English-ethiopian languages statistical machine translation. In *Proceedings of the 2019 Workshop on Widening NLP*, pages 27–30.

Biruk Abel. 2018. Geez to amharic machine translation.

Ife Adebara and Muhammad Abdul-Mageed. 2022. Towards afrocentric nlp for african languages: Where we are and where we can go. *Preprint*, arXiv:2203.08351.

David Adelani, Jesujoba Alabi, Angela Fan, Julia Kreutzer, Xiaoyu Shen, Machel Reid, Dana Ruiter, Dietrich Klakow, Peter Nabende, Ernie Chang, Tajuddeen Gwadabe, Freshia Sackey, Bonaventure F. P. Dossou, Chris Emezue, Colin Leong, Michael Beukman, Shamsuddeen Muhammad, Guyo Jarso, Oreen Yousuf, Andre Niyongabo Rubungo, Gilles Hacheme, Eric Peter Wairagala, Muhammad Umair Nasir, Benjamin Ajibade, Tunde Ajayi, Yvonne Gitau, Jade Abbott, Mohamed Ahmed, Millicent Ochieng, Anuoluwapo Aremu, Perez Ogayo, Jonathan Mukiibi, Fatoumata Ouoba Kabore, Godson Kalipe, Derguene Mbaye, Allahsera Auguste Tapo, Victoire Memdjokam Koagne, Edwin Munkoh-Buabeng, Valencia Wagner, Idris Abdulmumin, Ayodele Awokoya, Happy Buzaaba, Blessing Sibanda, Andiswa Bukula, and Sam Manthalu. 2022. A few thousand translations go a long way! leveraging pre-trained models for African news translation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3053–3070, Seattle, United States. Association for Computational Linguistics.

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, Edouard Grave, Michael Auli, and Armand Joulin. 2020. Beyond english-centric multilingual machine translation. *Preprint*, arXiv:2010.11125.

Sefineh Getachew and Yirga Yayeh. 2023. Gex'ez-english bi-directional neural machine translation using transformer. In *2023 International Conference on Information and Communication Technology for Development for Africa (ICT4DA)*, pages 160–164.

Barry Haddow, Rachel Bawden, Antonio Valerio Miceli Barone, Jindřich Helcl, and Alexandra Birch. 2022. Survey of low-resource machine translation. *Preprint*, arXiv:2109.00486.

Franz Josef and Hermann Ney. 2001. Statistical machine translation.

Tadesse Kassa. 2018. Morpheme-based bi-directional ge'ez -amharic machine translation.

Ertiban Demewoz Molla and Debre Tabor. 2018. An analysis of ge'ez language heritage potential: traditional church schools and the practices of ethiopian orthodox tewahido churches.

Dawit Mulugeta. 2015. Geez to amharic automatic machine translation: A statistical approach.

Wilhelmina Nekoto, Vukosi Marivate, Tshinondiwa Matsila, Timi Fasubaa, Taiwo Fagbohungbe, Solomon Oluwole Akinola, Shamsuddeen Muhammad, Salomon Kabongo Kabenamualu, Salomey Osei, Freshia Sackey, Rubungo Andre Niyongabo, Ricky Macharm, Perez Ogayo, Orevaoghene Ahia, Musie Meressa Berhe, Mofetoluwa Adeyemi, Masabata Mokgesi-Selinga, Lawrence Okegbemi, Laura Martinus, Kolawole Tajudeen, Kevin Degila, Kelechi Ogueji, Kathleen Siminyu, Julia Kreutzer, Jason Webster, Jamiil Toure Ali, Jade Abbott, Iroro Orife, Ignatius Ezeani, Idris Abdulkadir Dangana, Herman Kamper, Hady Elsahar, Goodness Duru, Ghollah Kioko, Murhabazi Espoir, Elan van Biljon, Daniel Whitenack, Christopher Onyefuluchi, Chris Chinenye Emezue, Bonaventure F. P. Dossou, Blessing Sibanda, Blessing Bassey, Ayodele Olabiyi, Arshath Ramkilowan, Alp Öktem, Adewale Akinfaderin, and Abdallah Bashir. 2020. Participatory research for low-resourced machine translation: A case study in African languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2144–2160, Online. 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.

Surangika Ranathunga, En-Shiun Annie Lee, Marjana Prifti Skenduli, Ravi Shekhar, Mehreen Alam, and Rishemjit Kaur. 2021. Neural machine translation for low-resource languages: A survey. *ACM Computing Surveys*.

Gebru Tareke, Bahru Zewde, and David Pool. 2002. A history of modern ethiopia 1855-1991. *The International Journal of African Historical Studies*, 35:587.

NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, 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. *Preprint*, arXiv:2207.04672.

Ermias Tegenaw, Kris Calpotura, and Ashebir Dereje. 2023. Ge'ez to amharic translation with neural network-based technique.

Atnafu Lambebo Tonja, Tadesse Destaw Belay, Olga Kolesnikova, Seid Muhie Yimam, Abinew Ali Ayele, Grigori Sidorov, and Alexander Gelbukh. 2023. Amhen: Amharic-english large parallel corpus for machine translation.

Ke M. Tran, Arianna Bisazza, and Christof Monz. 2014. Word translation prediction for morphologically rich

languages with bilingual neural networks. In *Conference on Empirical Methods in Natural Language Processing*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Aman Kassahun Wassie. 2023. Machine translation for ge'ez language. *Preprint*, arXiv:2311.14530.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Huggingface's transformers: State-of-the-art natural language processing. *Preprint*, arXiv:1910.03771.

Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. ByT5: Towards a token-free future with pre-trained byte-to-byte models. *Transactions of the Association for Computational Linguistics*, 10:291–306.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

A Appendix

A.1 Loss Graph

Training loss curves for NLLB200 600M model between Amharic, Geez, and English, showing the progress of model learning over 10,000 steps. By the end of the 10,000 steps, the training loss for all models seems to converge, which means they may be approaching their optimal performance.

Figure 2: NLLB-200 Loss Graph

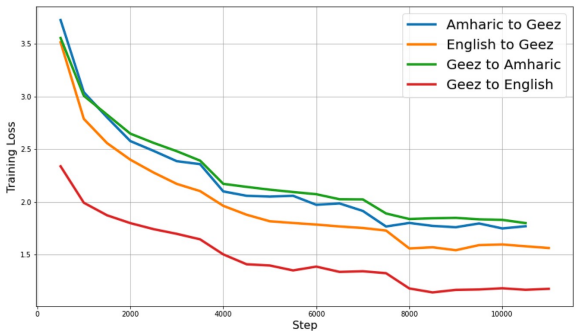

A.2 In-house tool web interface

Screenshot of a multilingual translation review interface showing sentence pairs in Amharic, Geez, and English, alongside user interaction options for approving or commenting on the translations for quality assurance and data curation purposes.

Figure 3: Reviewers interface of our in-house system.

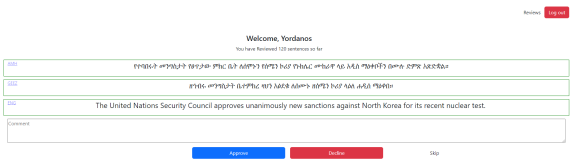

Figure 4: Interface of our in-house system receiving all the data.