

Extended Parallel Corpus for Amharic-English Machine Translation

Andargachew Mekonnen Gezmu, Andreas Nürnberger, Tesfaye Bayu Bati

Otto von Guericke Universität Magdeburg
 Universitätsplatz 2, Magdeburg, Germany
 {andargachew.gezmu, andreas.nuernberger}@ovgu.de
 tesfayebayu@hu.edu.et

Abstract

This paper describes the acquisition, preprocessing, segmentation, and alignment of an Amharic-English parallel corpus. It will be helpful for machine translation of a low-resource language, Amharic. We freely released the corpus for research purposes. Furthermore, we developed baseline statistical and neural machine translation systems; we trained statistical and neural machine translation models using the corpus. In the experiments, we also used a large monolingual corpus for the language model of statistical machine translation and back-translation of neural machine translation. In the automatic evaluation, neural machine translation models outperform statistical machine translation models by approximately six to seven Bilingual Evaluation Understudy (BLEU) points. Besides, among the neural machine translation models, the subword models outperform the word-based models by three to four BLEU points. Moreover, two other relevant automatic evaluation metrics, Translation Edit Rate on Character Level and Better Evaluation as Ranking, reflect corresponding differences among the trained models.

Keywords: Statistical Machine Translation, Neural Machine Translation, Less-Resourced Language

1. Introduction

To automate the intricate task of translation, researchers have followed different approaches. The earliest attempt was to use rule-based systems, which are criticized for being tedious and expensive to develop. Alternative empirical approaches such as Statistical Machine Translation (SMT) and Neural Machine Translation (NMT) came when parallel corpora were more and more available. Such methods take advantage of the authentic translations made by human translators in parallel corpora. They rely on machine learning to build translation models by taking parallel corpora as training data.

For empirical machine translation, we need a parallel corpus (bitext), a text that has a parallel translation in another language. Machine translation models are trained on a parallel corpus. Some international and governmental institutions provide such texts for public use. For instance, the Canadian Hansard corpus (Roukos S. et al., 1995) consists of parallel texts in English and French, drawn from official records of the proceedings of the Canadian Parliament. Similarly, the Europarl corpus (Koehn and Monz, 2005), extracted from the proceedings of the European Parliament, contains parallel corpora for twenty-one European languages. The United Nations (UN) Parallel Corpus (Ziems et al., 2016) is available in six official UN languages. The current version of the parallel corpus consists of manually translated UN documents between 1990 and 2014.

Other parallel corpora have been made from movie subtitles, like the OpenSubtitles corpus (Lison and Tiedemann, 2016), or from general web text, like the ParaCrawl corpus (Bañón et al., 2020). The Open Par-

allel Corpus (OPUS) (Tiedemann, 2012) collects parallel corpora from sources such as open-source software documentations and religious books.

Large numbers of parallel corpora are available for dominant international languages such as English, German, and French. Nevertheless, there is a scarcity of available parallel corpora for low-resource languages. The deficiency impedes the progress of machine translation for such languages.

We considered different sources to develop a parallel corpus for Amharic-English translation. The existing corpora were either small or had poor quality; they were mainly collected from the web. Although considering the web as a corpus, which is motivated for practical reasons of getting more extensive data with open access and low cost, may sound good, such sources are inaccurate. Moreover, as Amharic is not standardized, one may face many spelling variations in these sources and expect typographical errors. This calls for manual or automatic editing. Therefore, we collected our corpora from edited documents such as newspapers, magazines, and textbooks. We also normalized the text and made some automatic spelling error corrections.

Amharic is a Semitic language that serves as the official language of Ethiopia. Although it plays several roles in the government, it is considered a low-resource language because of its lack of essential tools and resources for natural language processing (Gezmu et al., 2018a; Tracey and Strassel, 2020).

Amharic uses a syllabic writing system, Ethiopic. Each Amharic letter systematically conflates a consonant and vowel (e.g., ብ /bə/ and ቡ /bu/). Sometimes consonants and vowels can be written as bare consonants (e.g., ብ /b/) or bare vowels (e.g., ኦ /a/ in አገር /agər/). Some phonemes with one or more homophonic

script representations and peculiar labiovelars sometimes compromise the consistency of the writing system. In Amharic orthography, there is no case difference. It is written from left to right. In present-day Amharic writings, words are delimited by plain space. Like other Semitic languages, Amharic words are highly inflectional and have a root-pattern morphology (Fabri et al., 2014). Therefore, Amharic lexicons cannot contain all word forms; the available bilingual lexicons contain only lemmas of common words. This hinders us to use sentence aligners that require bilingual lexicons.

Thus, in this research, we compiled a parallel corpus¹ for Amharic-English machine translation by extending the Ge'ez Frontier Foundation's news corpus made available for research purposes. We collected additional bilingual documents from various sources to compile the corpus. We also trained and evaluated NMT and SMT models using the corpus.

2. Related Work

There were attempts to compile parallel corpora for Amharic-English machine translation. The most notable ones are the “Amharic-English bilingual corpus”, “English-Ethiopian languages parallel corpora” (Abate et al., 2018), the “Low Resource Languages for Emergent Incidents: Amharic representative language pack” (LORELEI-Amharic) (Tracey and Strassel, 2020), and the OPUS collection (Tiedemann, 2012).

The European Language Resource Association (ELRA) hosts the Amharic-English bilingual corpus, containing a small parallel text from legal and news domains. Abate et al. (2018) compiled small-sized English-Ethiopian languages parallel corpora. Linguistic Data Consortium developed the LORELEI-Amharic corpus. Although LORELEI-Amharic is larger than Amharic-English bilingual corpus and English-Ethiopian languages parallel corpora, it is still not sufficient to train machine translation models with competitive performance (Koehn and Knowles, 2017; Lample et al., 2018). Besides, the parallel text was collected from discussion forums, newswires, and weblogs. Discussion forums and weblogs are susceptible to spelling mistakes. The problem worsens as there is no readily available spell checker to assist Amharic writers.

In the OPUS collection, there are parallel corpora for Amharic and English. For this language pair, however, some of the corpora have a few hundred parallel sentences (*e.g.*, Tatoeba, GlobalVoices, and TED2020); some use archaic language (*e.g.*, Tanzil and bible-uedin); and others contain misaligned parallel sentences (*e.g.*, MultiCCAligned and JW300).

The lack of clean, sizable, readily available, and contemporary-language parallel corpora impede the progress in Amharic-English machine translation.

¹The corpus is available at: <http://dx.doi.org/10.24352/ub.ovgu-2018-145>

There were few attempts in Amharic-English machine translation using small-sized corpora (Teshome and Besacier, 2012; Teshome et al., 2015; Ashengo et al., 2021). Still, their corpora are not readily available for the research community.

3. Corpus Preparation

We created a new parallel corpus by extending the existing news corpus made available for research purposes by Ge'ez Frontier Foundation. We collected, pre-processed, and segmented and aligned sentences of additional bilingual documents from various sources to compile the corpus.

3.1. Data Sources

We identified potential data sources that could serve as a basis for building a parallel corpus. We have considered newswires, magazines, and the Bible to get extensive data with open access.

Major newswires such as Deutsche Welle, BBC, and Ethiopian News Agency provide news articles in Amharic and English. Besides, the Ethiopian Herald and the Ethiopian Reporter publish bilingual news articles in Amharic and English. In these newswires, the translations are intended for the local public. Because of this, only a tiny portion of the English news articles are translated into Amharic, or vice versa. For instance, in the Ethiopian News Agency, approximately one news story out of ten has a rough translation (Argaw and Asker, 2005).

Watchtower (መጠበቂያ ግንብ in Amharic) and Awake (ንቁ magazines in Amharic) have been published since 2006. They are available for the public; they have adequate sentence-by-sentence translations. Watchtower mainly discusses religious issues. Unlike Watchtower, Awake contains articles on general interest topics such as nature, geography, and family life. So it corresponds more to news articles.

The Bible is the most translated and readily available book. It is translated with great care and has high coverage of vocabulary (Chew et al., 2006). Additionally, its content reflects the everyday living of human beings like love, war, and politics. However, older translations of the Bible used archaic languages. Fortunately, we found out the recent translations of the Bible use the contemporary language. For example, the Standard Version and the New World Translation use the modern-day language in both Amharic and English.

Therefore, we selected text from Awake and Watchtower magazines, the Bible, and newswires. Then, we preprocessed the text as a preparation step for the following sentence segmentation and alignment activities.

3.2. Preprocessing

The preprocessing of the text involves spelling correction and normalization. In addition, we removed boilerplates such as headers, footers (including footnotes), and verse numbers (in the Bible).

In the text, we observed different types of misspellings: misspellings result from missed out spaces, replacing letters with visually similar characters (*e.g.*, ቁ and ቁሁ), and typographical errors. We could not use our rule-based Amharic spelling corrector (Mekonnen, 2012) because of its limitations. Instead, we developed another spelling corrector (Gezmu et al., 2018b) that has a better performance measured with the benchmark test set² (Gezmu et al., 2021a). We employed the spelling corrector primarily to correct the first two types of spelling errors. Since the intensive manual intervention is needed to select the correct spelling from the plausible suggestions for typographical errors, we have not corrected the typographical errors in the current version of the corpus.

Different styles of punctuation marks have been used in Amharic text. For instance, for double quotation mark, two successive single quotation marks or similar symbols (*e.g.*, <<, >>, <, or >) are used; for end-of-sentence punctuation (:: “Amharic full stop”) two successive Amharic word separator (:) that give the same appearance are used. Thus, the normalization of punctuation is a nontrivial matter. We normalized all types of double quotes, all single quotes, question marks (*e.g.*, ? and ?), word separators (*e.g.*, : and :), full stops (*e.g.*, :: and ::), exclamation marks (*e.g.*, ! and !), hyphens (*e.g.*, :- and :-), and commas (*e.g.*, ቅ and ቅ).

3.3. Sentence Segmentation

Segmentation of sentences essentially involves the disambiguation of end-of-sentence punctuation. To do so, we identified end-of-sentence punctuation marks. We considered end-of-sentence punctuation (:: for Amharic and period for English) and question marks as a sentence boundary. The exceptions are abbreviations, initials of names, clitics, Uniform Resource Locators (URLs), e-mail addresses, and hashtags. Thus, to retain them we created a list of known abbreviations and clitics; and regular expressions for URLs, e-mail addresses, and hashtags. After sentence segmentation, we deleted duplicate sentences.

3.4. Sentence Alignment

Amharic has a rich morphology; it is practically impossible for Amharic lexicons to contain all word forms. Therefore, it is beneficial to use a sentence aligner that does not require any bilingual lexicon. Hence, we used the Bilingual Sentence Aligner³ (Moore, 2002) to align sentences in the bilingual documents.

Table 1 shows the number of sentences aligned in each bilingual document. The corpus is comprised of approximately 83% of the Watchtower magazine and Bible text that can be considered as a “belief and

Document	Number of sentence pairs
Awake	16491
Watchtower	72512
The Bible	48651
News articles	7710
Total	145364

Table 1: The number of sentences (segments) aligned in each bilingual document.

thought” domain (Burnard, 2007). The remaining 17% of the Awake magazine and news articles is in the “world affairs” domain (Burnard, 2007).

After merging and shuffling the aligned sentences, we divided them into the training, validation (development), and test sets. Table 2 shows the statistics of each dataset.

4. Baseline Systems

During recent years, there are many improvements over SMT, such as hierarchical phrase-based SMT (Chiang, 2007) and syntax-based SMT (Galley et al., 2004; Galley et al., 2006), and NMT like Universal Transformers (Dehghani et al., 2019). Nevertheless, we relied on baseline systems for both approaches to evaluate them objectively.

4.1. Baseline SMT System

Our phrase-based SMT baseline system had settings that were typically used by Ding et al. (2016), Williams et al. (2016), Koehn and Knowles (2017), and Sennrich and Zhang (2019). We used the Moses (Koehn et al., 2007) toolkit to train phrase-based SMT models. First, we used GIZA++ (Och, 2003) and the growdiag-final-and heuristic for symmetrization for word alignment. Then, we used the phrase-based reordering model (Koehn et al., 2003) with three different orientations: monotone, swap, and discontinuous in backward and forward directions conditioned on the source and target languages.

We used five-gram language models smoothed with the modified Kneser-Ney (Kneser and Ney, 1995). The system applied KenLM (Heafield, 2011) language modeling toolkit for this purpose. Initially, we have not used big monolingual corpora for language models. This is because they are no longer the exclusive advantages of phrase-based SMT, as NMT can also benefit from them (Sennrich and Zhang, 2019). Afterward, to prove this claim, we used the Contemporary Amharic Corpus⁴ (CACO) (Gezmu et al., 2018b) for English-to-Amharic translation.

The feature weights were tuned using Minimum Error Rate Training (MERT) (Och, 2003). We also used the k-best batch Margin Infused Relaxed Algorithm

²The test set is available at: <https://github.com/andmek/ErrorCorpus>

³The implementation is available at: <https://www.microsoft.com/en-us/download/details.aspx?id=52608>

⁴The corpus is available at: <http://dx.doi.org/10.24352/ub.ovgu-2018-144>

Dataset	Sentences	English Tokens	Amharic Tokens	English Types	Amharic Types
Test	2500	46154	34689	5842	11644
Validation	2864	53818	39980	6470	13068
Training	140000	2574538	1930220	33589	155824
Total	145364	2674510	2004889	45901	180536

Table 2: The number of sentences (segments), tokens, and types in each dataset.

(MIRA) for tuning (Cherry and Foster, 2012) by selecting the highest-scoring development run with a return-best-dev setting.

In decoding, we applied the default normal stack search algorithm.

4.2. Baseline NMT System

To train NMT models, we used the encoder-decoder architecture implemented with Transformers Vaswani et al. (2017). We tuned the hyperparameters of our NMT baseline system following Vaswani et al. (2017), Deng et al. (2018), Gezmu et al. (2021b), and Gezmu and Nürnberger (2022). The hyperparameters include the Adam optimizer (Kingma and Ba, 2015) with varied learning rate over the course of training, dropout (Srivastava et al., 2014) rate of 0.1, label smoothing (Szegedy et al., 2016) of value 0.1, batch size of 1024, six Transformer blocks with eight heads, filter size of 2048, and hidden size of 512. We used tensor2tensor (Vaswani et al., 2018) library to implement the system. The situation of training NMT models is complex because the training of NMT models is usually non-deterministic and hardly ever converges (Popel and Bojar, 2018). Most research in NMT does not specify any stopping criteria. Some mention only an approximate number of days elapsed to train the models (Bahdanau et al., 2015) or the exact number of training steps (Vaswani et al., 2017). We trained, thus, each NMT model for 250000 steps following Vaswani et al. (2017).

For decoding, we used a single model obtained by averaging the last twelve checkpoints. Following Wu et al. (2016), we used a beam search with a beam size of four and a length penalty of 0.6.

5. Experiments and Evaluation

We evaluated the performance of the SMT and NMT systems. We also made a comparison of word-based and subword-based NMT models. The experiments used the same datasets for each system.

5.1. Datasets and Preprocessing

We trained our models on the benchmark dataset – the Amharic-English parallel corpus explained in Section 3. The training set consists of 140000 sentence pairs; the validation and test sets have 2864 and 2500 sentence pairs.

We tokenized the English datasets with Moses’ tokenizer script; we modified Moses’ script to tokenize

the Amharic datasets. Next, to share named-entities between the languages, the Amharic datasets were transliterated with a transliteration scheme, Amharic transliteration for machine translation⁵, which is fully discussed in (Gezmu et al., 2021b).

We removed sentence pairs with extreme length ratios of more than one to nine and sentences longer than eighty tokens for the phrase-based SMT baseline. For word-based NMT models, we used a shared vocabulary of the top forty-four thousand most frequent tokens (tokens that appear five or more times in the corpus). We set this optimum vocabulary size because it will be too large to fit our GPU’s memory if we include less frequent tokens. Besides, we used the word-piece method (Schuster and Nakajima, 2012; Wu et al., 2016), which is similar to Byte Pair Encoding (Gage, 1994; Sennrich et al., 2016b), to segment words in the datasets for subword-based NMT models. We used the word-piece implementation in tensor2tensor⁶ library. Furthermore, since the vocabulary size in word-piece has an impact on the performance of the NMT models, we trained models with different vocabulary sizes (Wu et al., 2016; Denkowski and Neubig, 2017; Cherry et al., 2018; Ding et al., 2016).

5.2. Evaluation

Eventually, translation outputs of the test sets were detokenized and evaluated with a case-sensitive Bilingual Evaluation Understudy (BLEU) metric (Papineni et al., 2002). For consistency, we used the metric’s implementation made by Post (2018), sacreBLEU⁷. To fill the limitations of BLEU (Callison-Burch et al., 2006; Reiter, 2018), we also used Better Evaluation as Ranking (BEER) (Stanojevic and Sima’an, 2014) and Translation Edit Rate on Character Level (CharacTER) (Wang et al., 2016) metrics. Unlike BLEU and BEER, the smaller the CharacTER score, the better. Moreover, the Amharic outputs were not back transliterated to use these automatic metrics effectively.

6. Results

Table 3 shows the performance results of the SMT and NMT systems with BLEU, BEER, and Charac-

⁵The implementation is available at: <https://github.com/andmek/AT4MT>

⁶Available at: <https://github.com/tensorflow/tensor2tensor>

⁷Signature BLEU+case.mixed+numrefs.1+smooth.exp+tok.13a+version.1.4.9

Translation Direction	System	BLEU	BEER	CharacTER
Amharic-to-English	NMT-1K	32.2	0.575	0.536
	NMT-2K	32.2	0.575	0.536
	NMT-4K	32.8	0.577	0.530
	NMT-8K	33.0	0.576	0.527
	NMT-16K	32.9	0.574	0.528
	NMT-32K	32.2	0.570	0.539
	NMT-Word-Based	28.8	0.537	0.588
	SMT-MERT	26.0	0.514	0.629
	SMT-MIRA	23.2	0.494	0.705
English-to-Amharic	NMT-1K	25.5	0.558	0.520
	NMT-2K	25.7	0.554	0.525
	NMT-4K	26.1	0.557	0.517
	NMT-8K	26.4	0.555	0.521
	NMT-16K	26.7	0.555	0.520
	NMT-32K	26.7	0.552	0.523
	NMT-Word-Based	23.0	0.514	0.585
	SMT-MERT	20.0	0.502	0.643
	SMT-MIRA	19.2	0.484	0.704

Table 3: Performance results of SMT and NMT models.

TER metrics. The SMT system achieved better scores when feature weights were tuned using MERT than batch MIRA. Thus, we took the phrase-based SMT system tuned with MERT as our baseline. Likewise, the NMT baseline systems use vocabulary sizes of eight thousand (8K) and sixteen thousand (16K) in Amharic-to-English and English-to-Amharic translation directions. Example Amharic-to-English translation outputs are given at the appendix.

System	BLEU	BEER	CharacTER
SMT	20.0	0.502	0.643
SMT + CACO	21.2	0.508	0.628
NMT	26.7	0.555	0.520
NMT + CACO	27.8	0.563	0.501

Table 4: Performance results of English-to-Amharic translation using the CACO corpus.

In both translation directions, NMT models with subword units score the highest values of all. The NMT baseline models outperform the SMT baseline models by approximately six to seven BLEU. Among the NMT models, the subword-based models outperform the word-based models by three to four BLEU. The BEER and CharacTER metrics as well reflect corresponding differences.

Even though big monolingual corpora are not integral components of NMT, both SMT and NMT can benefit from them. Table 4 shows the results of English-to-Amharic translation using the CACO corpus for language model of the baseline SMT system, and back-translation of the baseline NMT system to produce synthetic training data (Sennrich et al., 2016a; He et al., 2016; Cheng et al., 2016; Qin, 2020). Both systems

gained more than one BLEU scores by using CACO. The baseline NMT model attained the optimum result when we randomly drew three times the size of the original training data from the CACO corpus and generated synthetic data by translating it to English.

7. Conclusions

We collected, preprocessed, segmented, and aligned Amharic-English parallel sentences from various sources. In doing so, we addressed different issues such as normalization and spelling correction. The corpus will be helpful for machine translation of a low-resource language, Amharic. Therefore, we freely released the corpus for research purposes. Also, we developed baseline SMT and NMT systems; we trained SMT and NMT models using the corpus. Additionally, we used a large monolingual corpus for the language model of SMT and back-translation of NMT in the experiments. As a result, NMT models outperform SMT models by approximately six to seven BLEU in the automatic evaluation. Besides, among the NMT models, the subword models outperform the word-based models by three to four BLEU. Moreover, two other relevant automatic evaluation metrics, CharacTER and BEER, reflect corresponding differences among the trained models.

We recommend future work to increase the size of the corpus by extracting text from scanned documents. In addition, we are engaged in doing additional experiments with other word segmentation methods for subword-based translations.

8. Acknowledgements

We want to thank Daniel Yacob, director of Ge'ez Frontier Foundation, for the Amharic news articles.

Appendix: Example Translation Outputs

The following examples show the reference translations and the translated English sentences using NMT and SMT systems.

Reference: About that time, my parents asked me to come back home.

NMT-Word-Based: About that time, I asked my parents to go to their home.

NMT-Subword-Based: About that time, my parents asked me to return home.

SMT-MERT: About that time, my parents to return to the house to house work.

SMT-MIRA: About that time, my parents to return to the house.

Reference: 2. Can we really live forever?

NMT-Word-Based: 2. Can we really live forever?

NMT-Subword-Based: 2. Can we really live forever?

SMT-MERT: 2. Really, we can live forever?

SMT-MIRA: 2. Really live forever?

Reference: Sandra quickly discovered that she had been scammed.

NMT-Word-Based: Sandra saw that she was dying right away.

NMT-Subword-Based: Sandra immediately recognized that she was mistaken.

SMT-MERT: Sandra she immediately.

SMT-MIRA: Sandra immediately.

Reference: Distressing circumstances can have a terrible impact on us.

NMT-Word-Based: When distressing situations strike, they may feel emotionally.

NMT-Subword-Based: Distressing situations can cause anxiety.

SMT-MERT: When distressing situations can emotional.

SMT-MIRA: When distressing situations emotional.

Reference: Olive oil is used copiously, as it is produced there on a large scale.

NMT-Word-Based: Olive oil is used in abundant value.

NMT-Subword-Based: The olive oil is so extensive that it pushes on the abundant possible.

SMT-MERT: As the bulk of olive oil for the benefit of the.

SMT-MIRA: Olive oil as a in the.

Reference: Six years later, the whole world economy collapsed.

NMT-Word-Based: Six years later, the entire world economy was destroyed.

NMT-Subword-Based: Six years later, the global economy sank into the world.

SMT-MERT: Six years later, the entire world econ-

omy, have been shattered.

SMT-MIRA: Six years later the global economy, have been shattered.

Reference: We also need to remember that Jesus said: “you must love your neighbor as yourself.”

NMT-Word-Based: We must remember that Jesus state: “you must love your neighbor as yourself.”

NMT-Subword-Based: Keep in mind that Jesus also said: “you must love your neighbor as yourself.”

SMT-MERT: Jesus said: “you must love your neighbor as yourself,” that we keep in mind.

SMT-MIRA: Jesus said: “you must love your neighbor as yourself” that we keep in mind.

9. Bibliographical References

- Abate, S. T., Woldeyohannis, M. M., Tachbelie, M. Y., Meshesha, M., Atinafu, S., Mulugeta, W., Assabie, Y., Abera, H., Seyoum, B. E., Abebe, T., Tsegaye, W., Lemma, A., Andargie, T., and Shifaw, S. (2018). Parallel corpora for bi-lingual english-ethiopian languages statistical machine translation. In Emily M. Bender, et al., editors, *Proceedings of the 27th International Conference on Computational Linguistics, COLING 2018, Santa Fe, New Mexico, USA, August 20-26, 2018*, pages 3102–3111. Association for Computational Linguistics.
- Argaw, A. A. and Asker, L. (2005). Web mining for an amharic - english bilingual corpus. In José Cordeiro, et al., editors, *WEBIST 2005, Proceedings of the First International Conference on Web Information Systems and Technologies, Miami, USA, May 26-28, 2005*, pages 239–246. INSTICC Press.
- Ashengo, Y. A., Aga, R. T., and Abebe, S. L. (2021). Context based machine translation with recurrent neural network for english-amharic translation. *Mach. Transl.*, 35(1):19–36.
- Bahdanau, D., Cho, K., and Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. In Yoshua Bengio et al., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Bañón, M., Chen, P., Haddow, B., Heafield, K., Hoang, H., Esplà-Gomis, M., Forcada, M. L., Kamran, A., Kirefu, F., Koehn, P., Ortiz-Rojas, S., Sempere, L. P., Ramírez-Sánchez, G., Sarrías, E., Strelec, M., Thompson, B., Waites, W., Wiggins, D., and Zaragoza, J. (2020). Paracrawl: Web-scale acquisition of parallel corpora. In Dan Jurafsky, et al., editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 4555–4567. Association for Computational Linguistics.
- Burnard, L. (2007). Reference guide for the british national corpus (xml edition): Design of the corpus [online]. accessed 12 april 2022.

- Callison-Burch, C., Osborne, M., and Koehn, P. (2006). Re-evaluation the role of bleu in machine translation research. In Diana McCarthy et al., editors, *EACL 2006, 11st Conference of the European Chapter of the Association for Computational Linguistics, Proceedings of the Conference, April 3-7, 2006, Trento, Italy*. The Association for Computer Linguistics.
- Cheng, Y., Xu, W., He, Z., He, W., Wu, H., Sun, M., and Liu, Y. (2016). Semi-supervised learning for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Cherry, C. and Foster, G. F. (2012). Batch tuning strategies for statistical machine translation. In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 3-8, 2012, Montréal, Canada*, pages 427–436. The Association for Computational Linguistics.
- Cherry, C., Foster, G. F., Bapna, A., Firat, O., and Macherey, W. (2018). Revisiting character-based neural machine translation with capacity and compression. In Ellen Riloff, et al., editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 4295–4305. Association for Computational Linguistics.
- Chew, P. A., Verzi, S. J., Bauer, T. L., and McClain, J. T. (2006). Evaluation of the Bible as a resource for cross-language information retrieval. In *Proceedings of the Workshop on Multilingual Language Resources and Interoperability*, pages 68–74, Sydney, Australia, July. Association for Computational Linguistics.
- Chiang, D. (2007). Hierarchical phrase-based translation. *Comput. Linguistics*, 33(2):201–228.
- Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., and Kaiser, L. (2019). Universal transformers. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Deng, Y., Cheng, S., Lu, J., Song, K., Wang, J., Wu, S., Yao, L., Zhang, G., Zhang, H., Zhang, P., Zhu, C., and Chen, B. (2018). Alibaba’s neural machine translation systems for WMT18. In Ondrej Bojar, et al., editors, *Proceedings of the Third Conference on Machine Translation: Shared Task Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pages 368–376. Association for Computational Linguistics.
- Denkowski, M. J. and Neubig, G. (2017). Stronger baselines for trustable results in neural machine translation. In Thang Luong, et al., editors, *Proceedings of the First Workshop on Neural Machine Translation, NMT@ACL 2017, Vancouver, Canada, August 4, 2017*, pages 18–27. Association for Computational Linguistics.
- Ding, S., Duh, K., Khayrallah, H., Koehn, P., and Post, M. (2016). The JHU machine translation systems for WMT 2016. In *Proceedings of the First Conference on Machine Translation, WMT 2016, collocated with ACL 2016, August 11-12, Berlin, Germany*, pages 272–280. The Association for Computer Linguistics.
- Fabri, R., Gasser, M., Habash, N., Kiraz, G., and Wintner, S. (2014). Linguistic introduction: The orthography, morphology and syntax of semitic languages. In Imed Zitouni, editor, *Natural Language Processing of Semitic Languages, Theory and Applications of Natural Language Processing*, pages 3–41. Springer.
- Gage, P. (1994). A new algorithm for data compression. *C Users Journal*, 12(2):23–38.
- Galley, M., Hopkins, M., Knight, K., and Marcu, D. (2004). What’s in a translation rule? In Julia Hirschberg, et al., editors, *Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, HLT-NAACL 2004, Boston, Massachusetts, USA, May 2-7, 2004*, pages 273–280. The Association for Computational Linguistics.
- Galley, M., Graehl, J., Knight, K., Marcu, D., DeNeefe, S., Wang, W., and Thayer, I. (2006). Scalable inference and training of context-rich syntactic translation models. In Nicoletta Calzolari, et al., editors, *ACL 2006, 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, Sydney, Australia, 17-21 July 2006*. The Association for Computer Linguistics.
- Gezmu, A. M. and Nürnberger, A. (2022). Transformers for low-resource neural machine translation. In Ana Paula Rocha, et al., editors, *Proceedings of the 14th International Conference on Agents and Artificial Intelligence, ICAART 2022, Volume 1, Online Streaming, February 3-5, 2022*, pages 459–466. SCITEPRESS.
- Gezmu, A. M., Nürnberger, A., and Seyoum, B. E. (2018a). Portable spelling corrector for a less-resourced language: Amharic. In Nicoletta Calzolari, et al., editors, *Proceedings of the Eleventh International Conference on Language Resources and Evaluation, LREC 2018, Miyazaki, Japan, May 7-12, 2018*. European Language Resources Association (ELRA).
- Gezmu, A. M., Seyoum, B. E., Gasser, M., and Nürnberger, A. (2018b). Contemporary Amharic corpus: Automatically morpho-syntactically tagged Amharic corpus. In *Proceedings of the First Workshop on Linguistic Resources for Natural Language Processing*, pages 65–70, Santa Fe, New Mexico, USA, August. Association for Computational Lin-

- guistics.
- Gezmu, A. M., Lema, T. T., Seyoum, B. E., and Nürnberger, A. (2021a). Manually annotated spelling error corpus for amharic. *CoRR*, abs/2106.13521.
- Gezmu, A. M., Nürnberger, A., and Bati, T. B. (2021b). Neural machine translation for Amharic-English translation. In Ana Paula Rocha, et al., editors, *Proceedings of the 13th International Conference on Agents and Artificial Intelligence, ICAART 2021, Volume 1, Online Streaming, February 4-6, 2021*, pages 526–532. SCITEPRESS.
- He, D., Xia, Y., Qin, T., Wang, L., Yu, N., Liu, T., and Ma, W. (2016). Dual learning for machine translation. In Daniel D. Lee, et al., editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 820–828.
- Heafield, K. (2011). Kenlm: Faster and smaller language model queries. In Chris Callison-Burch, et al., editors, *Proceedings of the Sixth Workshop on Statistical Machine Translation, WMT@EMNLP 2011, Edinburgh, Scotland, UK, July 30-31, 2011*, pages 187–197. Association for Computational Linguistics.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In Yoshua Bengio et al., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kneser, R. and Ney, H. (1995). Improved backing-off for m-gram language modeling. In *1995 International Conference on Acoustics, Speech, and Signal Processing, ICASSP '95, Detroit, Michigan, USA, May 08-12, 1995*, pages 181–184. IEEE Computer Society.
- Koehn, P. and Knowles, R. (2017). Six challenges for neural machine translation. In Thang Luong, et al., editors, *Proceedings of the First Workshop on Neural Machine Translation, NMT@ACL 2017, Vancouver, Canada, August 4, 2017*, pages 28–39. Association for Computational Linguistics.
- Koehn, P. and Monz, C. (2005). Shared task: Statistical machine translation between european languages. In Philipp Koehn, et al., editors, *Proceedings of the Workshop on Building and Using Parallel Texts@ACL 2005, Ann Arbor, Michigan, USA, June 29-30, 2005*, pages 119–124. Association for Computational Linguistics.
- Koehn, P., Och, F. J., and Marcu, D. (2003). Statistical phrase-based translation. In Marti A. Hearst et al., editors, *Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, HLT-NAACL 2003, Edmonton, Canada, May 27 - June 1, 2003*. The Association for Computational Linguistics.
- Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., Dyer, C., Bojar, O., Constantin, A., and Herbst, E. (2007). Moses: Open source toolkit for statistical machine translation. In John A. Carroll, et al., editors, *ACL 2007, Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, June 23-30, 2007, Prague, Czech Republic*. The Association for Computational Linguistics.
- Lample, G., Ott, M., Conneau, A., Denoyer, L., and Ranzato, M. (2018). Phrase-based & neural unsupervised machine translation. In Ellen Riloff, et al., editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 5039–5049. Association for Computational Linguistics.
- Lison, P. and Tiedemann, J. (2016). Opensubtitles2016: Extracting large parallel corpora from movie and TV subtitles. In Nicoletta Calzolari, et al., editors, *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016*. European Language Resources Association (ELRA).
- Mekonnen, A. (2012). Development of an amharic spelling corrector for tolerant-retrieval. In Janusz Kacprzyk, et al., editors, *International Conference on Management of Emergent Digital EcoSystems, MEDES '12, Addis Ababa, Ethiopia, October 28-31, 2012*, pages 22–26. ACM.
- Moore, R. C. (2002). Fast and accurate sentence alignment of bilingual corpora. In Stephen D. Richardson, editor, *Machine Translation: From Research to Real Users, 5th Conference of the Association for Machine Translation in the Americas, AMTA 2002 Tiburon, CA, USA, October 6-12, 2002, Proceedings*, volume 2499 of *Lecture Notes in Computer Science*, pages 135–144. Springer.
- Och, F. J. (2003). Minimum error rate training in statistical machine translation. In Erhard W. Hinrichs et al., editors, *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics, 7-12 July 2003, Sapporo Convention Center, Sapporo, Japan*, pages 160–167. ACL.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W. (2002). Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Popel, M. and Bojar, O. (2018). Training tips for the transformer model. *Prague Bull. Math. Linguistics*, 110:43–70.
- Post, M. (2018). A call for clarity in reporting BLEU scores. In Ondrej Bojar, et al., editors, *Proceedings of the Third Conference on Machine Translation: Research Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pages 186–191. Association for Computational Linguistics.

- Qin, T. (2020). *Dual Learning*. Springer.
- Reiter, E. (2018). A structured review of the validity of BLEU. *Comput. Linguistics*, 44(3).
- Schuster, M. and Nakajima, K. (2012). Japanese and Korean voice search. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2012, Kyoto, Japan, March 25-30, 2012*, pages 5149–5152. IEEE.
- Sennrich, R. and Zhang, B. (2019). Revisiting low-resource neural machine translation: A case study. In Anna Korhonen, et al., editors, *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 211–221. Association for Computational Linguistics.
- Sennrich, R., Haddow, B., and Birch, A. (2016a). Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Sennrich, R., Haddow, B., and Birch, A. (2016b). Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958.
- Stanojevic, M. and Sima'an, K. (2014). BEER: better evaluation as ranking. In *Proceedings of the Ninth Workshop on Statistical Machine Translation, WMT@ACL 2014, June 26-27, 2014, Baltimore, Maryland, USA*, pages 414–419. The Association for Computer Linguistics.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2818–2826. IEEE Computer Society.
- Teshome, M. G. and Besacier, L. (2012). Preliminary experiments on English-Amharic statistical machine translation. In *Third Workshop on Spoken Language Technologies for Under-resourced Languages, SLTU 2012, Cape Town, South Africa, May 7-9, 2012*, pages 36–41. ISCA.
- Teshome, M. G., Besacier, L., Taye, G., and Teferi, D. (2015). Phoneme-based English-Amharic statistical machine translation. In *AFRICON 2015*, pages 1–5. IEEE.
- Tiedemann, J. (2012). Parallel data, tools and interfaces in OPUS. In Nicoletta Calzolari, et al., editors, *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012, Istanbul, Turkey, May 23-25, 2012*, pages 2214–2218. European Language Resources Association (ELRA).
- Tracey, J. and Strassel, S. M. (2020). Basic language resources for 31 languages (plus English): The LORELEI representative and incident language packs. In Dorothee Beermann, et al., editors, *Proceedings of the 1st Joint Workshop on Spoken Language Technologies for Under-resourced Languages and Collaboration and Computing for Under-Resourced Languages, SLTU/CCURL@LREC 2020, Marseille, France, May 2020*, pages 277–284. European Language Resources Association.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. In Isabelle Guyon, et al., editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Vaswani, A., Bengio, S., Brevdo, E., Chollet, F., Gomez, A. N., Gouws, S., Jones, L., Kaiser, L., Kalchbrenner, N., Parmar, N., Sepassi, R., Shazeer, N., and Uszkoreit, J. (2018). Tensor2Tensor for neural machine translation. In Colin Cherry et al., editors, *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas, AMTA 2018, Boston, MA, USA, March 17-21, 2018 - Volume 1: Research Papers*, pages 193–199. Association for Machine Translation in the Americas.
- Wang, W., Peter, J., Rosendahl, H., and Ney, H. (2016). Character: Translation edit rate on character level. In *Proceedings of the First Conference on Machine Translation, WMT 2016, colocated with ACL 2016, August 11-12, Berlin, Germany*, pages 505–510. The Association for Computer Linguistics.
- Williams, P., Sennrich, R., Nadejde, M., Huck, M., Haddow, B., and Bojar, O. (2016). Edinburgh’s statistical machine translation systems for WMT16. In *Proceedings of the First Conference on Machine Translation, WMT 2016, colocated with ACL 2016, August 11-12, Berlin, Germany*, pages 399–410. The Association for Computer Linguistics.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, L., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., and Dean, J. (2016). Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144.
- Ziemski, M., Junczys-Dowmunt, M., and Poulisquen,

B. (2016). The united nations parallel corpus v1.0. In Nicoletta Calzolari, et al., editors, *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016*. European Language Resources Association (ELRA).

10. Language Resource References

Roukos S. et al. (1995). *Hansard French/English*. LDC Catalog No: LDC95T20. Linguistic Data Consortium.