

Lesan – Machine Translation for Low Resource Languages

Asmelash Teka Hadgu

ASME@LESAN.AI

Abel Aregawi

ABEL@LESAN.AI

Adam Beaudoin

ADAM@LESAN.AI

Lesan AI

Editors: Douwe Kiela, Marco Ciccone, Barbara Caputo

Abstract

Millions of people around the world can not access content on the Web because most of the content is not readily available in their language. Machine translation (MT) systems have the potential to change this for many languages. Current MT systems provide very accurate results for high resource language pairs, e.g., German and English. However, for many low resource languages, MT is still under active research. The key challenge is lack of datasets to build these systems. We present Lesan¹, an MT system for low resource languages. Our pipeline solves the key bottleneck to low resource MT by leveraging online and offline sources, a custom Optical Character Recognition (OCR) system for Ethiopic and an automatic alignment module. The final step in the pipeline is a sequence to sequence model that takes parallel corpus as input and gives us a translation model. Lesan’s translation model is based on the Transformer architecture. After constructing a base model, back translation is used to leverage monolingual corpora. Currently Lesan supports translation to and from Tigrinya, Amharic and English. We perform extensive human evaluation and show that Lesan outperforms state-of-the-art systems such as Google Translate and Microsoft Translator across all six pairs. Lesan is freely available and has served more than 10 million translations so far. At the moment, there are only 217 Tigrinya and 15,009 Amharic Wikipedia articles. We believe that Lesan will contribute towards democratizing access to the Web through MT for millions of people.

Keywords: machine translation, deep learning, low resource languages, amharic, tigrinya

1. Machine Translation Pipeline

Our machine translation pipeline shown in Figure 1 has the following three key components: OCR, automatic alignment and MT modeling that leverage current advances in deep learning.

Data sources Machine translation (MT) systems today provide very accurate results for high resource language pairs Barrault et al. (2019). However, these systems rely on large-scale parallel corpus to work. For pairs such as German and English, these parallel sentences can be easily obtained from the Web. However, for the majority of languages there is not enough data online to build these datasets. We leverage offline and online sources to mitigate this bottleneck.

Offline sources are gathered through partnerships with media companies, libraries and other archives that have books, magazines and other linguistic materials. These are typically translated books or magazines that have daily news of source and target language in different file formats and hard copies. Online sources include news sites, blogs, religious and legal sites that feature stories in both the source and target language.

1. <https://lesan.ai/>

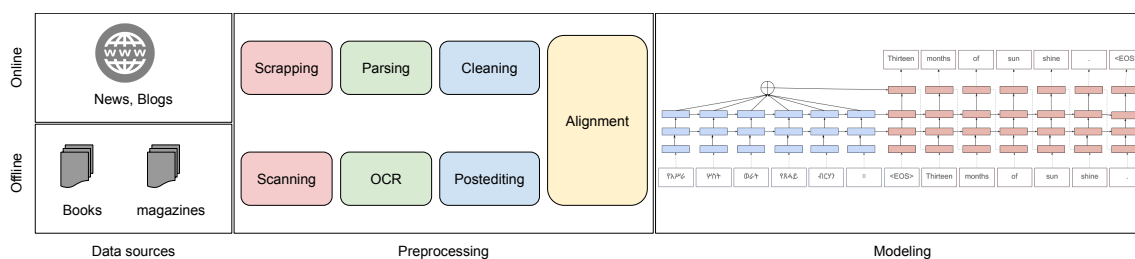


Figure 1: Lesan machine translation system pipeline

Data preprocessing After gathering large-scale unstructured data from online and offline sources, the next step is a preprocessing step. The goal of this module is to generate clean text (corpus) in machine readable format. Scanned books and magazines pass through an OCR system to convert them into machine readable format. We built a custom OCR system for Ethiopic, script used for Tigrinya, Amharic and several other languages. The output of our OCR system is passed through post-editing to correct errors in characters and words. Data crawled from the Web is parsed to extract text by removing HTML-tags and gathering metadata such as author, date that could be helpful in downstream steps. Text coming from different sources such as text parsed from web pages and text coming from the post-editing module passed through several NLP modules such as deduplication, language detection and sentence splitting.

Automatic alignment All source and target sentences that have been cleaned are then passed through an automatic alignment system to construct parallel corpus. A candidate generation step, produces candidate pairs, then each pair of source and target sentence is fed to our automatic matching system that determines whether an input pair is a translation or not. This system is similar to [Junczys-Dowmunt \(2018\)](#) where we compute cross-entropy scores according to two inverse translation models, and take the weighted average aggregated over a set of candidate pairs.

Sequence-to-sequence modeling with back translation The final step in the pipeline is a sequence to sequence model that takes the parallel corpus as input and gives us a translation model. Lesan’s translation model is based on the Transformer architecture [Vaswani et al. \(2017\)](#). After constructing a base model, back translation [Edunov et al. \(2018\)](#) is used to leverage monolingual corpora.

Our machine translation pipeline solves the key bottleneck to low resource MT by leveraging online and offline sources, a custom OCR system for Ethiopic and an automatic alignment module.

2. Evaluation

In this section we will describe human evaluation of Lesan compared to three state-of-the-art commercial systems: Google Translate ², Microsoft Translator ³ and Yandex Translate ⁴.

2. <https://translate.google.com/>

3. <https://www.bing.com/translator>

4. <https://translate.yandex.com/>

Google Translate uses neural MT (Wu et al., 2016). Yandex Translate uses a hybrid system of neural MT and statistical MT ⁵. All these services provide APIs to access their system. Lesan’s MT models are implemented using the OpenNMT Klein et al. (2017) toolkit.

Human Evaluation Experts were selected to evaluate the translation outputs of the systems. We assign the task of selecting what news sources to translate to the expert evaluators in order to avoid bias. The main requirements were: identify 20 news stories (1 - 2 paragraphs); select stories across different genres: politics, social, economy, entertainment, health, sport, technology etc. The output of translation systems next to the sources are given to expert evaluators to assign a score. The outputs are shuffled in such a way that one cannot tell which output is from which system. We chose a Likert scale ⁶ of 5 corresponding to a range from completely wrong (0) to accurate and fluent (4). These error categories are adapted from the severity levels for error categories ⁷ used in evaluation of translation quality developed by TAUS. The complete description of the scoring scheme is shown in Table 1. We have released the evaluation datasets ⁸ to foster research and progress on evaluating MT systems for low resource languages.

Table 1: Human evaluation guideline to evaluate performance of MT systems.

Scale	Value	Description
Wrong translation	0	This is for a completely wrong translation. The translation output does not make sense given the source text.
Major problem	1	There is a serious problem in the translation with some parts of the source missing or mistranslated and it would be hard to match translation output with source text without major modifications.
Minor problem	2	The translation has minor problems given the source text but requires some minor changes, e.g, changing a word or two to make it fully describe the source text.
Good translation	3	The translation describes the source text; however, there may be some problems with style such as punctuation, word order or appropriate wording.
Accurate and fluent	4	Great job! The output is a correct translation of the source text. It’s both accurate and fluent.

We report the normalized mean and standard deviation of the scores. The results are given in Table 2 for Amharic to and from English and in Table 3 for Tigrinya to and from Amharic and English. Across all directions Lesan outperforms these state-of-the-art systems.

3. Broader Impact

There are several applications of machine translation systems for broader impact. Let’s take the case of Wikipedia. Wikipedia currently has a total of more than six billion articles and

5. <https://tech.yandex.com/translate/doc/dg/concepts/how-works-machine-translation-docpage/>

6. https://en.wikipedia.org/wiki/Likert_scale

7. <https://www.taus.net/qt21-project#harmonized-error-typology>

8. <https://zenodo.org/record/5060303>

Table 2: Human evaluation comparing Lesan against three commercial MT systems for Amharic to and from English.

Direction	System	Sentence	Story
Am → En	Yandex	0.23 ± 0.30	0.19 ± 0.25
	Microsoft	2.13 ± 0.51	2.06 ± 0.50
	Google	2.58 ± 0.54	2.54 ± 0.48
	Lesan	2.68 ± 0.41	2.71 ± 0.55
En → Am	Yandex	0.28 ± 0.34	0.20 ± 0.29
	Microsoft	2.57 ± 0.43	2.54 ± 0.44
	Google	2.98 ± 0.30	2.88 ± 0.33
	Lesan	3.25 ± 0.38	3.17 ± 0.42

Table 3: Human evaluation comparing Lesan and Microsoft for Tigrinya to and from Amharic and English.

Direction	System	Sentence	Story
Am → Ti	Microsoft	1.92 ± 0.43	1.85 ± 0.54
	Lesan	1.94 ± 0.47	1.86 ± 0.53
Ti → Am	Microsoft	1.60 ± 0.44	1.44 ± 0.57
	Lesan	1.94 ± 0.50	1.77 ± 0.50
En → Ti	Microsoft	2.32 ± 0.6	2.17 ± 0.63
	Lesan	2.33 ± 0.63	2.19 ± 0.58
Ti → En	Microsoft	2.01 ± 0.63	1.89 ± 0.67
	Lesan	2.78 ± 0.31	2.63 ± 0.39

over 17 billion words in its English edition. Unfortunately, millions of people cannot access this because it’s not available in their language. For instance, at the moment there are only 217 articles on the Tigrinya Wikipedia and 15,009 articles on the Amharic Wikipedia ⁹. In future work, we would like to leverage Lesan’s MT system to empower human translators towards our mission of opening up the Web’s content to millions of people in their language.

Acknowledgments

We would like to thank Sergey Edunov from Facebook AI Research for valuable feedback on our machine translation pipeline.

References

- Loïc Barrault, Ondrej Bojar, Marta R. Costa jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. Findings of the 2019 conference on machine translation (wmt19). In *WMT (2)*, pages 1–61. Association for Computational Linguistics, 2019.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. Understanding back-translation at scale. *arXiv preprint arXiv:1808.09381*, 2018.
- Marcin Junczys-Dowmunt. Dual conditional cross-entropy filtering of noisy parallel corpora. *arXiv preprint arXiv:1809.00197*, 2018.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M Rush. Opennmt: Open-source toolkit for neural machine translation. *arXiv preprint arXiv:1701.02810*, 2017.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144, 2016.

9. https://en.wikipedia.org/wiki/List_of_Wikipedias

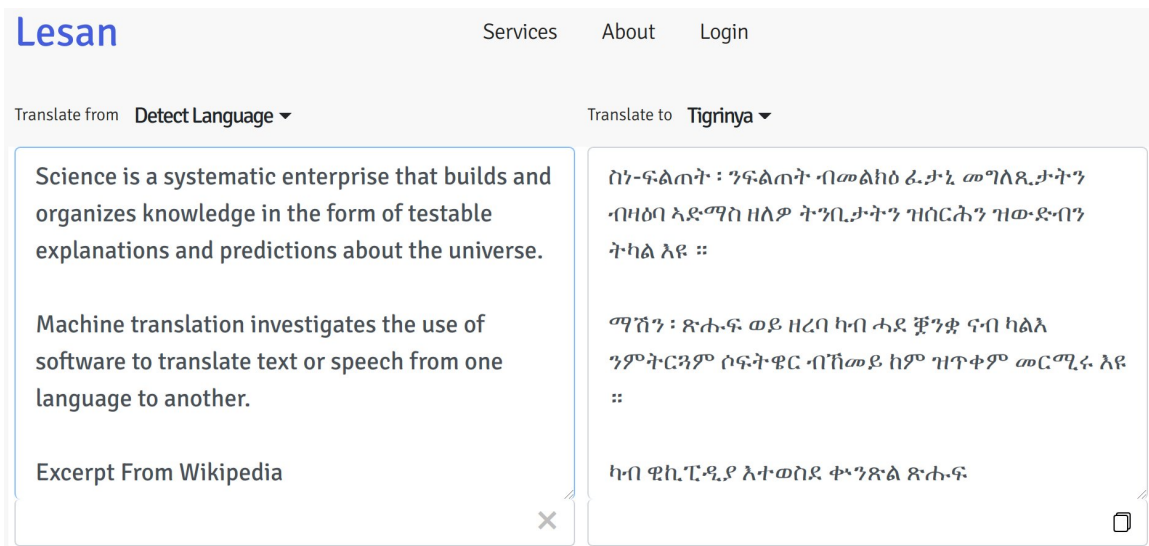


Figure 2: Web interface of Lesan - on the left is an excerpt from English Wikipedia and on the right the translation output to Tigrinya.