

Neural Machine Translation for Extremely Low-Resource African Languages: A Case Study on Bambara

Allahsera Auguste Tapo^{1,*}, Bakary Coulibaly², Sébastien Diarra², Christopher Homan¹, Julia Kreutzer³, Sarah Luger⁴, Arthur Nagashima¹, Marcos Zampieri¹, Michael Leventhal²

¹Rochester Institute of Technology

²Centre National Collaboratif de l'Éducation en Robotique et en Intelligence Artificielle (RobotsMali)

³Google Research, ⁴Orange Silicon Valley

*aat3261@rit.edu

Abstract

Low-resource languages present unique challenges to (neural) machine translation. We discuss the case of Bambara, a Mande language for which training data is scarce and requires significant amounts of pre-processing. More than the linguistic situation of Bambara itself, the socio-cultural context within which Bambara speakers live poses challenges for automated processing of this language. In this paper, we present the first parallel data set for machine translation of Bambara into and from English and French and the first benchmark results on machine translation to and from Bambara. We discuss challenges in working with low-resource languages and propose strategies to cope with data scarcity in low-resource machine translation (MT).

1 Introduction

Underresourced languages, from a natural language processing (NLP) perspective, are those lacking the resources (large volumes of parallel bitexts) needed to support state-of-the-art performance on NLP problems like machine translation, automated speech recognition, or named entity recognition. Yet the vast majority of the world's languages—representing billions of native speakers worldwide—are underresourced. And the lack of available training data in such languages usually reflects a broader paucity of electronic information resources accessible to their speakers.

For instance, there are over six million Wikipedia articles in English but fewer than sixty thousand in Swahili and fewer than seven hundred in Bambara, the vehicular and most widely-spoken native language of Mali that is the subject of this paper.¹ Consequently, only 53% of the worlds

population have access to “encyclopedic knowledge” in their primary language, according to a 2014 study by Facebook.² MT technologies could help bridge this gap, and there is enormous interest in such applications, ironically enough, from speakers of the languages on which MT has thus far had the least success. There is also great potential for humanitarian response applications (Öktem et al., 2020).

Fueled by data, advances in hardware technology, and deep neural models, machine translation (NMT) has advanced rapidly over the last ten years. Researchers are beginning to investigate the effectiveness of (NMT) low-resource languages, as in recent WMT 2019 and WMT 2020 tasks (Barrault et al., 2019), and in underresourced African languages. Most prominently, the *Masakhane* (V et al., 2020) community³, a grassroots initiative, has developed open-source NMT models for over 30 African languages on the base of the JW300 corpus (Agić and Vulić, 2019), a parallel corpus of religious texts.

Since African languages cover a wide spectrum of linguistic phenomena and language families (Heine and Nurse, 2000), individual development of translations and resources for selected languages or language families are vital to drive the overall progress. Just within the last year, a number of dedicated studies have significantly improved the state of African NMT: van Biljon et al. (2020) analyzed the depth of Transformers specifically for low-resource translation of South-African languages, based on prior studies by Martinus and Abbott (2019) on the Autshumato corpus (Groenewald and du Plooy, 2010). Dossou and Emezue (2020) developed an MT model and compiled resources for translations between Fon and French,

¹meta.wikimedia.org/wiki/List_of_Wikipedias.

²fbnewsroomus.files.wordpress.com/2015/02/state-of-connectivity1.pdf

³masakhane.io

Akinfaderin (2020) modeled translations between English and Hausa, Orife (2020) for four languages of the Edoid language family, and Ahia and Ogueji (2020) investigated supervised vs. unsupervised NMT for Nigerian Pidgin.

In this paper, we present the first parallel data set for machine translation of Bambara into and from English and French and the first benchmark results on machine translation to and from Bambara. We discuss challenges in working with low-resource languages and propose strategies to cope with data scarcity in low-resource MT. We discuss the socio-cultural context of Bambara translation and its implications for model and data development. Finally, we analyze our best-performing neural models with a small-scale human evaluation study and give recommendations for future development. We find that the translation quality on our in-domain data set is acceptable, which gives hope for other languages that have previously fallen under the radar of MT development.

We released our models and data upon publication⁴. Our evaluation setup may serve as benchmark for an extremely challenging translation task.

2 The Bambara Language

Bambara is the first language of five million people and the second language of approximately ten million more. Most of its speakers are members of Bambara ethnic groups, who live throughout the African continent. Approximately 30–40 million people speak some language in the Mande family of languages, to which Bambara belongs (Lewis et al., 2014).

Bambara is a tonal language with a rich morphology. Over the years, several competing writing systems have developed, however, as an historically predominately oral language, a majority of Bambara speakers have never been taught to read or write the standard form of the language. Many are incapable of reading or writing the language at all. The standardization of words and the coinage of new ones are still works in progress; this poses challenges to automated text processing.

During Muslim expansion and French colonization, Arabic and French mixed with local languages, resulting in a lingua franca, e.g., Urban Bambara. Most of the existing Bambara resources are cultural (folk stories or news/topical) or come

⁴https://github.com/israaar/mt_bambara_data_models

from social media or text messages, and these are written in a melange of French, Bambara and Arabic. Consequently, corpora based on common Bambara usage must account for the code switching found in these mixtures.

Most of these characteristics are shared with related languages, e.g., a subset of the Mande family of languages, where many languages are mutually intelligible. Thus, our hope is that our approach will be transferable to the other twelve official local languages of Mali, or to other African languages with a comparable socio-cultural and linguistic embedding, for example Wolof (non-Mande), which is comparable in terms of number of speakers, borrowings from Arabic and French influence, and oral traditions.

The next section will provide more details on digital resources and describe the process of exploring and collecting data and choosing parallel corpora for the training of the NMT model.

3 Data Collection

3.1 Bambara Corpora

We discovered that there has been no prior development of automatic translation of Bambara, despite a relatively large volume of research on the language (Culy, 1985; Aplonova and Tyers, 2017; Aplonova, 2018). As a pilot study for assessing the potential for automatic translation of Bambara, Leventhal et al. (2020) crowdsourced a small set of written or oral translations from French to Bambara. Additional work was carried out exploring novel crowdsourcing strategies for data collection in Mali Luger et al. (2020).

The *Corpus Bambara de Référence* (Vydrin et al., 2011) is the largest collection of electronic texts in Bambara. It includes scanned and text-based electronic formats. A number of parallel texts based on this data exist. For example, Vydrin (2018) analyzed Bambara’s separable adjectives using this data.

To survey the known available sources of parallel texts with Bambara, we consulted with a number of authorities on Bambara, including the *Academie Malienne des Langues* (AMALAN) in Mali and the *Institut National des Langues et Civilisations Orientales* (INALCO) in France, as well as a number of individual linguists and machine translation experts throughout the world. These two organisations play key roles in the definition and the promotion of a standard form of written Bambara

through the collection and annotation of corpora, the publishing of dictionaries, and, in formulating recommendations for language policy in Mali.

Our efforts uncovered several sources of parallel texts between Bambara and French and/or English that are listed in Table 5 in Appendix A. The table provides a rating of each of the identified resources, and the rationale why they were in- or excluded from our translation study. Ultimately, most of these resources proved either of very little or no practical use as sources of training data. Many did not actually contain aligned texts and some not even suitable monolingual text.

A systematic problem was lack of adherence to the standardized Bambara orthography, due to it being a predominately oral language. This is also one of the reasons why our search for parallel data on the web generally did not yield many finds—commonly being used in written form, Bambara is used even less on the web. For example, the Bambara Wikipedia contains currently 667 articles (compared to 6M for English), of which a large percentage are only stubs. Of the small number of full articles, most do not consistently employ the standard orthography of Bambara. A selection of those, however, was prepared to be used as monolingual data for MT data augmentation.

Most African NMT studies have been based on the *JW300* corpus (Agić and Vulić, 2019), e.g. most of the Masakhane benchmarks (V et al., 2020). *JW300* only contains less than 200 sentences of Dyula, a closely related language to Bambara that it is mutually intelligible with. It might be useful for future cross-lingual studies (Wu et al., 2019; Goyal et al., 2020), but in order to avoid interference between languages, we focus on Bambara data exclusively in this first study.

The most promising for our NMT approach was a dictionary data set from *SIL Mali*⁵ with examples of sentences used to demonstrate word usage in Spanish, French, English, and Bambara; and a tri-lingual health guide titled “Where there is no doctor.”⁶ Detailed corpus statistics are listed in Table 1.

3.2 Sentence Alignment

The part of the dictionary that we are focusing on in this study, are the dictionary entries that consist of examples of Bambara expressions followed by

⁵<https://www.sil-mali.org/en/content/introducing-sil-mali>

⁶<https://gafe.dokotoro.org/>

		Bambara	French	English
Dict.	glosses	3,548	4,847	4,855
	examples	2,023	2,021	2,021
	aligned	2,158	2,146	2,158
Medical	chapters		27	
	files		336	
	paragraphs	9,336	9,367	9,356
	unigrams	8,209	9,893	6,935
	bigrams	26,430	25,746	31,412
	trigrams	5,816	11,312	21,398
	stopwords	147	123	69

Table 1: The dictionary data set from *SIL Mali* and the medical health guide “Where there is no doctor” with examples in French, English, and Bambara.

their translations in French and in English. Most of these are single sentences, so there is sentence-to-sentence alignment in the majority of cases. However, there remains a sufficient number of exceptions to render automated pairing impossible. Part of the problem lies in the unique linguistic and cultural elements of the bambaraphone environment; it is often not possible to meaningfully translate an expression in Bambara without giving an explanation of the context.

The medical health guide is aligned by chapters, each of which is roughly aligned by paragraphs. But at the paragraph level there are too many exceptions for automated pairing to be feasible. Many of the bambaraphone-specific problems found in the dictionary dataset are present at the sentence level as well, particularly in explanations of concepts that can be succinctly expressed in English or French but for which Bambara lacks terminology and the bambaraphone environment lacks an equivalent physical or cultural context.

Both datasets therefore required manual alignment by individuals fluent in written Bambara and either French or English. The annotators need to be able to exercise expert-level judgment on linguistic and, occasionally, medical questions. Access to such human resources was a major factor limiting the quantity of data we were able to align. Because of this, and since the dictionary data was more closely aligned at the sentence level and did not require as much domain knowledge as the medical dataset, we have thus far only used the dictionary dataset in our machine learning experiments.

In order to facilitate this alignment, we imple-



Figure 1: The custom aligner we developed to manually align the dictionary data set. The controls are as follows: for each language, “>” goes to the next item and “<” goes to the previous item; for all languages, “>>>” goes to the next items and “<<<” goes to the previous items; “Aligned B-F-E” saves to memory the alignment of all 3 languages; “Aligned B-F” saves to memory the alignment of Bambara and French items; “Aligned B-E” saves to memory the alignment of Bambara and English items; “Save” saves to a new file; “Continue Saving” continues saving the file created.

mented an alignment interface, as shown in Figure 1. It allows annotators to manually align sentences and to save those sentence pairs that another annotator considered properly aligned. In separate tasks, four annotators with a secondary school level understanding of Bambara performed alignment on French-Bambara and English-Bambara sentence pairs using the tool.

3.3 Preprocessing

Before we could align these sentences, we needed to clean the retrieved dictionary entries. Below we give examples of cases we had to handle manually, going through the entire corpus line by line.

1. Only one language is represented: Discarded.

2. Ambiguous pronouns in Bambara:⁷

Before:

fr: “Il/elle est né à Bamako en 1938.”

bam: “A bangera Bamako san 1938.”

After:

fr: “Il est né à Bamako en 1938.”

bam: “A bangera Bamako san 1938.”

and

fr: “Elle est né à Bamako en 1938.”

bam: “A bangera Bamako san 1938.”

3. Additional explanations in the other languages while those are absent in Bambara:

Before:

fr: “Un doigt ne peut pas prendre un caillou (C’est important d’aider les uns les autres).”

bam: “Bolokoni kelen te se ka bele ta.”

⁷One can imagine that translating these into French or English is difficult since there is no indicator of the correct choice. (Johnson, 2018)

After:

fr: “Un doigt ne peut pas prendre un caillou.”

bam: “Bolokoni kelen te se ka bele ta.”

4. Proverbs:

Before:

fr: “Proverbe: Une longue absence vaut mieux qu’un communiqué (d’un décès).”

bam: “Fama ka fisa ni komunike ye.”

After:

fr: “Une longue absence vaut mieux qu’un communiqué.”

bam: “Fama ka fisa ni komunike ye.”

Data preparation, including alignment, proved to be about 60% of the overall time spent in person-hours on the experiment and required on-the-ground organisation and recruitment of skilled volunteers in Mali.

3.4 Parallel Data

The final data set contains 2,146 parallel sentences of Bambara-French and 2,158 parallel sentences of Bambara-English—a very small data set for NMT compared to the massive state-of-the-art models that are trained on billions of sentences (Arivazhagan et al., 2019). We split the data randomly into training, validation, and test sets of 75%, 12.5% and 12.5% respectively. The training set is composed of 1611 sentences, the validation set of 268 sentences, the test set of 267 sentences for Bambara-French. The training set is composed of 1620 sentences, the validation set of 270 sentences, the test set of 268 sentences for Bambara-French.

3.5 Monolingual Data

In addition to the translations, we obtained a dataset of 488 monolingual Bambara sentences, sampled from all articles in the Bambara Wikipedia and covering a range of topics, but with preponderance of articles related to Mali. We used this monolingual dataset for experiments in data augmentation through back-translation, described in Section 5.1.

4 NMT Development

4.1 Hyperparameters

Our NMT is a transformer (Vaswani et al., 2017) of appropriate size for a relatively smaller training dataset (van Biljon et al., 2020). It has six layers with four attention heads for encoder and decoder, the transformer layer has a size of 1024, and the hidden layer size 256, the embeddings have 256 units. Embeddings and vocabularies are not shared across languages, but the softmax layer weights are tied to the output embedding weights. The model is implemented with the Joey NMT framework (Kreutzer et al., 2019) based on PyTorch (Paszke et al., 2019).

Training runs for 120 epochs in batches of 1024 tokens each. The ADAM optimizer (Kingma and Ba, 2014) is used with a constant learning rate of 0.0004 to update model weights. This setting was found to be best to tune for highest BLEU, compared to decaying or warmup-cooldown learning rate scheduling. For regularization, we experimented with dropout and label smoothing (Szegedy et al., 2016). The best values were 0.1 for dropout and 0.2 for label smoothing across the board. For inference, beam search with width of 5 is used. The remaining hyperparameters are documented in the Joey NMT configuration files that we will provide with the code.

4.2 Segmentation

There is no standard tokenizer for Bambara. Therefore, we simply apply whitespace tokenization for word-based NMT models and compute BLEU with "international" tokenization.⁸

Of the 542 distinct word types in the Bambara dev set, 166 are not contained in the vocabulary (seen during training), 174 of 590 (29.5%) for the test split. For the French portion it is 243 of 713 (34.1%) for the dev split and 274 of 756 (36.2%) for the test split. Because of this large proportion of unknown words, we segment the data for

⁸Tokenizing on punctuation and special symbols, except when surrounding a digit.

both language pairs into subword units (byte pair encodings, BPE) (500 or 1000, separately) using subword-nmt⁹ (Sennrich et al., 2016), and apply BPE dropout to the training sets of both languages (Provilkov et al., 2019). We also experiment with character-level translation for French.

5 Results

5.1 Automatic Evaluation

Segmentation. We evaluate the models' translations against reference translations on our held-out sets with corpus BLEU (Papineni et al., 2002) and ChrF (Popović, 2015) computed with SacreBLEU (Post, 2018).¹⁰ Tables 2 and 3 show the results for French and English translations respectively. We find that word- and character-level modeling performs sub par compared to subword-level segmentation, which is in line with previous work on low-resource MT. The word-based model cannot resolve out-of-vocabulary words, and the character-level model struggled with word composition. With BPE, smaller subwords seem to perform slightly better than larger ones. BPE dropout (Provilkov et al., 2019), which was previously reported to be helpful for low-resource MT (Richburg et al., 2020), did not increase the quality of the results. We observe a trend towards higher scores for translations into Bambara than in the reverse direction, but this cross-lingual comparison has to be taken with a grain of salt, since it is influenced by source and target complexity (Bugliarello et al., 2020). Ambiguities on the Bambara side, such as the gender of pronouns illustrated in the example in Section 3.2, might make translation into English and French particularly difficult.

Back-translation. In addition, we experimented with back-translated Wikipedia data: fine-tuning the original model on a combination of the back-translated and original data, or training it from scratch on a combination of both, as e.g. in (Przytupa and Abdul-Mageed, 2019). However, this did not yield improvements over the original BPE model.¹¹ We speculate that the mismatch between domains hindered improvement. Indeed, we discovered that when we selected only short sentences

⁹<https://github.com/rsennrich/subword-nmt>

¹⁰BLEU+case.mixed+numrefs.1+smooth.exp+tok.intl+version.1.4.9, chrF2+case.mixed+numchars.6+numrefs.1+space.False+version.1.4.9

¹¹Not included in tables.

NMT Model	Configuration	bam→fr		fr→bam	
		BLEU	ChrF	BLEU	ChrF
(1) Word	2464 (fr) / 1724 (bam) words	18.9	0.3	20.6	0.3
(2) Char	97 (fr) / 89 (bam) chars	11.7	0.2	10.6	0.2
(2) BPE	500 subword merges each	19.1	0.3	21.1	0.3
(3) BPE	1000 subword merges each	19.2	0.3	20.4	0.3
(4) (2) + BPE dropout	dropout=0.1	17.8	0.3	16.6	0.2
(5) (3) + BPE dropout	dropout=0.1	16.0	0.3	17.7	0.2
Test scores for the best models from above		20.9	0.3	21.4	0.3

Table 2: NMT results for French and Bambara translations in corpus BLEU and ChrF on the dev set. The best model (in bold) is evaluated on the test set, results are reported in the last row.

NMT Model	Configuration	bam→en		en→bam	
		BLEU	ChrF	BLEU	ChrF
(1) Word	2364 (en) / 1745 (bam) words	19.1	0.3	17.7	0.3
(2) Char	81 (en) / 86 (bam) chars	13.2	0.2	12.5	0.2
(3) BPE	500 subword merges each	18.1	0.3	21.3	0.3
(4) BPE	1000 subword merges each	19.1	0.3	20.0	0.3
(5) (3) + BPE dropout	dropout=0.1	17.0	0.2	19.3	0.3
(6) (4) + BPE dropout	dropout=0.1	15.2	0.2	18.2	0.2
Test scores for the best models from above		14.8	0.3	20.9	0.3

Table 3: NMT results for English and Bambara translations in corpus BLEU and ChrF on the dev set. The best model (in bold) is evaluated on the test set, results are reported in the last row.

from the Wikipedia data, we observed slightly better results, but they still did not outperform the baseline. This highlights the importance of general domain evaluation sets for future work, so that the effectiveness of leveraging additional out-of-domain data can be measured.

Multilingual modeling. Another promising approach for the improvement of extremely low-resourced languages is multilingual modeling (Johnson et al., 2017). In our case, we combined the tasks of translating from English and French into Bambara to strengthen the Bambara decoding abilities of the translation model, by concatenating the training data and learning joint BPE models. The training data is filtered so that it does not contain sentences from the evaluation sets of the respective other language. However, we do not find improvements over the bilingual model.¹² We would have expected improvements on translating

¹²Initially, with cross-lingual overlap between training data of one language and evaluation data of the other language we found large improvements, which shows that the model can transfer from English to French sources very well (and vice versa).

into Bambara because of larger variation on the source side. However, one reason for not seeing this improvement might be that the sentences are relatively short, and fluency is not as much of an issue as in larger scale studies from previous works.

5.2 Human evaluation

Two native Bambara speakers from Mali, co-authors of this paper, with both college-level French and English reading skills, evaluated a random sample of 21 of the test set translations from Bambara into French and 21 different test set translations from Bambara into English produced by the highest scoring models. Both native speakers received their basic education in Bambara and can read and write the language with fluency.

Evaluation Schema. For translations that had only a limited correspondence to the source text, the evaluators were given a number of questions specific to the quality of the translations. For translations that rose to the level of being qualified as conveying most of the sense of the source text we asked for two numerical ratings: First, whether

“most people” would be able to understand what the meaning of the source sentence from its translation. This was intended to cover translations where the technical accuracy might be low, as a BLEU score might measure, but that substantially conveyed the meaning of the source text. Second, whether the translation was a “good translation,” meaning that exact word choices and structure hewed closely to the style and meaning of the source text. We chose to use relatively inexact terminology to describe the ranking criteria as we felt that, as non-professional translators, our evaluators would have difficulty using more technical guidance.

Quantitative results. The evaluation schema and the results obtained are presented in Table 4. The proportion of adequate words appears similar for English and French (rows 1 and 3). However, more English translations are judged as being adequate (rows 2 and 4). The overall percentage of translations that might be said to be useful (row 7), in that they convey at least the gist of the source sentence, is low at 36%, similar to the results obtained by automated methods.

Examples. The following translation excited our admiration because the sentence is relatively complex and the translation is flawless: “Farafinna tilebinyanfan tun tilalen bε Angilew ni Tubabuw cε.”, which gets translated to “L’Afrique de l’Ouest a été divisée entre les Anglais et les Français.” (“West Africa was divided between the English and the French”.)

We also observe that the MT system often translated verb tense correctly, perhaps helped by the fact that verb tenses are extremely simple in Bambara. Some of the sentences that did not qualify as adequate translations nonetheless were instructive and demonstrated specific pattern recognition capabilities.

For another example with the Bambara source “I bε gojogojo wa?”, the model translates “Have you ever eaten you is your wife?”. The word “gojogojo” is slang Bambara mainly used by youth playfully employing reduplication, onomatopoeia, and inspiration from a foreign language. While the term produced a nonsense sentence, the translation seems to carry some of the playfulness in the source sentence. We also notice that it does pick up the subject and uses interrogative word order. In Bambara the word order is the same word order used for a declarative sentence, “You” - “are” -

“athletic”, the sentence is made interrogative by the interrogative marker “wa”.

Looking at the following example, with Bambara source “Araba, ɔkutɔburukalo tile 8 san 2003” and translation “Après la mort de 25 ans.”, the failure of translating time expressions is surprising because we would have expected the system to have been trained on this pattern - it says, Wednesday, the month of October, the 8th day, the year 2003. The translation (“After the death of 25 years”) is not close. Still, there are markers of time in both the original and translated sentences.

Finally, for Bambara source “A nalolen don i n’a fɔ suruku.” the translation says “Il met les mains dans ses poches.” (“He put the hands in his pockets.”), even though the correct translation would say “He is as crude as a hyena” (the word for crude does not translate very exactly into English). While the translation seems to have nothing to do with the source, it has the right subject and somehow seems a reasonable guess if you did not know the key words and gives a bit of the spirit of the source sentence.

6 Conclusion

Our study constitutes the first attempt of modeling automatic translation for the extremely low-resource language of Bambara. We identified challenges for future work, such as the development of alignment tools for small-scale datasets, and the need for a general domain evaluation set. The current limitation of processing written text as input might furthermore benefit from the integration of spoken resources through speech recognition or speech translation, since Bambara is primarily spoken and the lack of standardization in writing complicates the creation of clean reference sets and consistent evaluation.

References

- Željko Agić and Ivan Vulić. 2019. [JW300: A wide-coverage parallel corpus for low-resource languages](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3204–3210, Florence, Italy. Association for Computational Linguistics.
- Orevaoghene Ahia and Kelechi Ogueji. 2020. Towards supervised and unsupervised neural machine translation baselines for nigerian pidgin. “AfricaNLP” *Workshop at the 8th International Conference on Learning Representations*.

Evaluation Criterion	English	French	All
(1) Percentage of words in the translated sentence that are <i>related</i> to the subject of the source.	43%	51%	47%
(2) Translated sentences containing words <i>related</i> to the subject of the source.	71%	52%	62%
(3) Percentage of words that are plausible <i>direct translations</i> of words in the source.	37%	40%	39%
(4) Translated sentences containing words that are plausible <i>direct translations</i> of words in the source.	71%	52%	62%
(5) Translated sentences that say <i>something intelligible</i> whether or not related to the source.	71%	66%	69%
(6) Translated sentences that <i>convey most</i> of the information in the source.	33%	38%	36%
(7) Translated sentences where most people would be able to get, at a minimum, <i>the gist</i> of the source.	33%	38%	36%
(8) <i>Rating</i> of translation quality for sentences of (7) with “ <i>understandable by most people</i> ” as the criterion.	4.1/5	4.8/5	4.5/5
(9) <i>Rating</i> of translation quality for sentences of (7) with “ <i>good translation</i> ” as the criterion.	4.4/5	4.8/5	4.6/5

Table 4: Evaluation results for translations of the test set from Bambara into English and French. Ratings in the last two lines are on a scale from 1 (poor quality) to 5 (good quality).

- Adewale Akinfaderin. 2020. [HausaMT v1.0: Towards English–Hausa neural machine translation](#). In *Proceedings of the The Fourth Widening Natural Language Processing Workshop*, pages 144–147, Seattle, USA. Association for Computational Linguistics.
- Ekaterina Aplonova. 2018. Development of a bambara treebank. *ARANEA 2018*, page 7.
- Ekaterina Aplonova and Francis Tyers. 2017. Towards a dependency-annotated treebank for bambara. In *Proceedings of the 16th International Workshop on Treebanks and Linguistic Theories*, pages 138–145.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. [Massively multilingual neural machine translation in the wild: Findings and challenges](#).
- Loïc Barrault, Ondřej Bojar, Marta R Costa-Jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, et al. 2019. Findings of the 2019 conference on machine translation (wmt19). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61.
- Elan van Biljon, Arnu Pretorius, and Julia Kreutzer. 2020. [On optimal transformer depth for low-resource language translation](#). “AfricaNLP” Workshop at the 8th International Conference on Learning Representations.
- Emanuele Bugliarello, Sabrina J. Mielke, Antonios Anastasopoulos, Ryan Cotterell, and Naoaki Okazaki. 2020. [It’s easier to translate out of English than into it: Measuring neural translation difficulty by cross-mutual information](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1640–1649, Online. Association for Computational Linguistics.
- Christopher Culy. 1985. The complexity of the vocabulary of bambara. *Linguistics and philosophy*, 8(3):345–351.
- Bonaventure F. P. Dossou and Chris C. Emezue. 2020. [Ffr v1.0: Fon-french neural machine translation](#).
- ∇, Iroro Orife, Julia Kreutzer, Blessing Sibanda, Daniel Whitenack, Kathleen Siminyu, Laura Martinus, Jamiil Toure Ali, Jade Abbott, Vukosi Marivate, Salomon Kabongo, Musie Meressa, Espoir Murhabazi, Orevaoghene Ahia, Elan van Biljon, Arshath Ramkilowan, Adewale Akinfaderin, Alp Öktem, Wole Akin, Ghollah Kioko, Kevin Degila, Herman Kamper, Bonaventure Dossou, Chris Emezue, Kelechi Ogueji, and Abdallah Bashir. 2020. [Masakhane – machine translation for africa](#).
- Vikrant Goyal, Sourav Kumar, and Dipti Misra Sharma. 2020. [Efficient neural machine translation for low-resource languages via exploiting related languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 162–168, Online. Association for Computational Linguistics.
- J. Hendrik Groenewald and Liza du Plooy. 2010. Processing parallel text corpora for three south african language pairs in the autshumato project. In *Proceedings of the Second Workshop on African Language Technology*, Valletta, Malta.
- Bernd Heine and Derek Nurse. 2000. *African languages: An introduction*. Cambridge University Press.
- Melvin Johnson. 2018. [Providing gender-specific translations in google translate](#).
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Julia Kreutzer, Jasmijn Bastings, and Stefan Riezler. 2019. [Joey NMT: A minimalist NMT toolkit for](#)

- novices. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, pages 109–114, Hong Kong, China. Association for Computational Linguistics.
- Michael Leventhal, Allahsera Tapo, Sarah Luger, Marcos Zampieri, and Christopher M. Homan. 2020. [Assessing human translations from french to bambara for machine learning: a pilot study](#).
- M Paul Lewis, Gary F Simons, and Charles D Fennig. 2014. *Ethnologue: Languages of Africa and Europe*. SIL international.
- Sarah Luger, Allashera Auguste Tapo, Christopher M. Homan, Marcos Zampieri, and Michael Leventhal. 2020. Towards a crowdsourcing platform for low resource languages – a semi-supervised approach. In *Proceedings of the Eighth Conference on Human Computations*. AAAI.
- Laura Martinus and Jade Z. Abbott. 2019. [A focus on neural machine translation for african languages](#). *CoRR*, abs/1906.05685.
- Iroro Orife. 2020. Towards neural machine translation for edoid languages. “AfricaNLP” *Workshop at the 8th International Conference on Learning Representations*.
- 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 on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8026–8037. Curran Associates, Inc.
- Maja Popović. 2015. chrF: character n-gram f-score for automatic mt evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Ivan Provilkov, Dmitrii Emelianenko, and Elena Voita. 2019. [Bpe-dropout: Simple and effective subword regularization](#).
- Michael Przystupa and Muhammad Abdul-Mageed. 2019. [Neural machine translation of low-resource and similar languages with backtranslation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 224–235, Florence, Italy. Association for Computational Linguistics.
- Aquia Richburg, Ramy Eskander, Smaranda Muresan, and Marine Carpuat. 2020. [An evaluation of subword segmentation strategies for neural machine translation of morphologically rich languages](#). In *Proceedings of the The Fourth Widening Natural Language Processing Workshop*, pages 151–155, Seattle, USA. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, Long Beach, CA, USA.
- Valentin Vydrin. 2018. [Where corpus methods hit their limits: the case of separable adjectives in bambara](#). *Rhema*, (4).
- Valentin Vydrin, Kirill Maslinsky, Jean-Jacques Méric, and A Rovenchak. 2011. Corpus bambara de référence.
- Nianheng Wu, Eric DeMattos, Kwok Him So, Pin-zhen Chen, and Çağrı Çöltekin. 2019. Language discrimination and transfer learning for similar languages: experiments with feature combinations and adaptation. In *Proceedings of the Sixth Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 54–63.
- Alp Öktem, Mirko Plitt, and Grace Tang. 2020. Tigrinya neural machine translation with transfer learning for humanitarian response. “AfricaNLP” *Workshop at the 8th International Conference on Learning Representations*.

A Appendices

URL	Description	Rate	Pros	Cons
http://www.alanwood.net/unicode/n%27ko.html	Test for Unicode support in web browsers	2	character(n'ko, decimal), decimal, character(n'ko, hex), hex, name.	no audio present, no usable data.
http://www.fakoli.net/index.html	N'ko Institute of America, dedicated to the exploration of n'ko	2	accessible to English speakers.	confusing, writes n'ko using latin alphabets, no audio, and no usable data.
http://www.personal.psu.edu/ejp10/symbolcodes/bylanguage/nko.html	N'ko computing information (penn state)	2	facilitate installation of n'ko fonts and gives general information of it as a font	no usable data.
https://ipfs.io/ipfs/QmXoypiZjW3WknFiJnKlWHCnL7zvedxjQkDDP1mXwo6ucc/wiki/N%27Ko_alpha.html	N'ko alphabet	2	provides information about the alphabet from all angle possible	no audio present, no usable data.
http://www.bmanuel.org/clr/clr2_mp.html	Multilingual and parallel corpora	1	n/a	Unaccessible when we recently visited. no usable data.
https://podcasts.apple.com/us/podcast/learn-bambara-with-jeff-fraze/id433868610	podcast	2	short speeches.	not accurate, no usable data.
http://www.rfi.fr/emission/kan-jum-be-yen-mandenkan	rss/podcast	2	audio available	not transcribed, no usable data.
https://www.voabambara.com/podcasts	rss/podcast	2	audio/video available	not transcribed, no usable data.
https://www.voafrique.com/z/3551	podcast	2	audio/video available	not transcribed, no usable data.
https://www.voabambara.com/z/5098	podcast	2	audio/video available	not transcribed, no usable data.
http://ma.rfi.fr/	streaming station	2	available	not transcribed, no usable data.
http://dictionary.ankataa.com	dictionary	2	available; easily understandable	no phonetics, no usable data.
http://www.mali-pense.net/bm/lexicon/index.htm	dictionary	2	french/bambara – bambara/french	no supporting audio, no usable data.
http://cormand.huma-num.fr/	project of bambara de reference	3	Dictionary like	Not publicly available.

Table 5: Datasets for Bambara including: a URL, linking to where we found the source; a description, describing what it is about (as we understood it); a rating from 1 to 5, one is “related to our topic with no usable data”, and five is “related to our topic with usable data”; pros, describing its pluses; and cons, describing its minuses.