

# The first neural machine translation system for the Erzya language

David Dale\*

dale.david@yandex.ru

## Abstract

We present the first neural machine translation system for translation between the endangered Erzya language and Russian and the dataset collected by us to train and evaluate it. The BLEU scores are 17 and 19 for translation to Erzya and Russian respectively, and more than half of the translations are rated as acceptable by native speakers. We also adapt our model to translate between Erzya and 10 other languages, but without additional parallel data, the quality on these directions remains low. We release the translation models along with the collected text corpus, a new language identification model, and a multilingual sentence encoder adapted for the Erzya language.

## 1 Introduction

Out of the 7 thousand languages spoken around the world, only a minor fraction is covered by machine translation tools. For example, Google Translate<sup>1</sup> supports only 133 languages, and a recent model by NLLB Team et al. (2022) supports 202 languages. Most other languages are often considered “low-resource”, although some of them have millions of native speakers. In the context of machine translation, the resources that are low are, primarily, parallel and monolingual text corpora. In this work, we create a machine translation system for the previously uncovered Erzya language with only publicly available resources, a very small budget, and limited human efforts. We hope that it will inspire researchers and language activists to enlarge the coverage of existing NLP resources, and in particular, translation systems.

Our language of choice is Erzya (myv), which is spoken primarily in the Republic of Mordovia, located in the center of the European part of the

Russian Federation. The language, along with its close relative Moksha (mdf), belongs to the Mordvinic branch of the Uralic language family. These two languages, although not mutually intelligible (Janurik, 2017), are often referred to under the common name “Mordovian”. Erzya has had a written tradition since the beginning of the 19th century (Rueter, 2013). Its most widely used alphabet is Cyrillic, although there is a Latin-based alternative alphabet<sup>2</sup>. Erzya has supposedly 300 thousand speakers<sup>3</sup>, and it is one of the three official languages in Mordovia. According to the UNESCO classification, the Erzya language has a status of “definitely endangered” (UNESCO, 2010). Some researchers (Janurik, 2017) put it between the levels 6b (“threatened”) and 7 (“shifting”) on the EGIDS scale (Lewis and Simons, 2010), as it is widely used and transmitted between generations in rural communities but is being gradually displaced by Russian in urban areas. More details about the use of Erzya are given by Rueter (2013), who is also a major current contributor to Erzya NLP resources.

As far as we know, prior to this work, no neural machine translation (NMT) systems for Erzya have been published. To fill this gap, we create and publicly release<sup>4</sup> the following deliverables:

- A language identification model with enhanced recall for Erzya and Moksha languages;
- A sentence encoder for Erzya compatible with LaBSE (Feng et al., 2022);
- A small parallel Erzya-Russian corpus and a larger monolingual Erzya corpus;
- Two neural models for translation between Erzya and 11 other languages.

For translation between Russian and Erzya, we validate our models both by automatic metrics and with judgments of native speakers. More than half

\*The research was conducted between the author’s employments at Skolkovo Institute of Science and Technology (Skoltech) and at Meta AI.

<sup>1</sup><https://translate.google.com>

<sup>2</sup><http://valks.erzja.info> (currently blocked in Russia)

<sup>3</sup>In the 2010 census, 430 thousand people reported speaking Erzya or Moksha, but their proportions are unclear.

<sup>4</sup>The source code and links to other resources are provided at <https://github.com/slone-nlp/myv-nmt>.

of the translations are rated as acceptable.

## 2 Related Work

Low-resource NLP and, in particular, machine translation, have attracted a lot of attention. Among the recent ambitious projects are [Bapna et al. \(2022\)](#) and [NLLB Team et al. \(2022\)](#) that aim at creating NMT systems for hundreds of languages and rely heavily on collection of large online corpora and transfer learning. Other works, such as [Hämäläinen and Alnajjar \(2019\)](#), focus on efficient use of existing vocabularies and morphosyntactic tools to train machine translation systems for very low-resourced languages.

As far as we know, there are no published large parallel corpora or NMT systems for Erzya. [Rueter and Tyers \(2018\)](#) develop an Erzya treebank with a few hundred translations to English and Finnish. [Архангельский \(2019\)](#) present an Erzya web corpus<sup>5</sup> along with the way it was collected, but the corpus is available only via the web interface. For other published corpora, the situation is similar. There exists a half-finished rule-based machine translation system between Erzya and Finnish<sup>6</sup>, and a grammar parser for Erzya<sup>7</sup>. The software package UralicNLP ([Hämäläinen, 2019](#)) supports Erzya among other languages.

There have been a few attempts to transfer machine learning-based NLP resources to Erzya from high-resource languages. [Alnajjar \(2021\)](#) adapt Finnish, English, and Russian word embeddings to Erzya. [Muller et al. \(2021\)](#), [Ács et al. \(2021\)](#) and [Wang et al. \(2022\)](#) evaluate the performance of multilingual BERT-like models on natural language understanding tasks for new languages, including Erzya.

None of the works known to us train machine learning-based models that are capable of generating Erzya language.

## 3 Methodology and Experiments

### 3.1 Data Collection

As there are no large open-access corpora for Erzya, we compile Erzya and Erzya-Russian data from various sources:

- 12K parallel sentences from the Bible<sup>8</sup>;

<sup>5</sup><http://erzya.web-corpora.net/>

<sup>6</sup><https://github.com/apertium/apertium-myv-fin>

<sup>7</sup><https://github.com/timarkh/uniparser-grammar-erzya>

<sup>8</sup><http://finugorbib.com>

- 3K parallel Wikimedia sentences from OPUS ([Tiedemann, 2012](#));
- 42K parallel words or short phrases collected from various online dictionaries;
- the Erzya Wikipedia and the corresponding articles from the Russian Wikipedia;
- 18 books, including 3 books with Erzya-Russian bitexts<sup>9</sup>;
- Soviet-time books and periodicals<sup>10</sup>;
- The Erzya part of Wikisource<sup>11</sup>;
- Short texts by modern Erzya authors<sup>12</sup>;
- News articles from the Erzya Pravda website<sup>13</sup>;
- Texts found in LiveJournal<sup>14</sup> by searching with the 100 most frequent Erzya words.

A more detailed account of the data sources is given in Appendix A.

After filtering these texts with the language identification model (Section 3.2), we gathered 330K unique Erzya sentences. A bilingual part of the texts was used for mining additional parallel sentences in Section 3.4.

### 3.2 Language Identification

To make sure that the extra collected data is in the Erzya language, we train a FastText ([Joulin et al., 2016](#)) language classifier for the 323 languages present in Wikipedia. The 267 thousand training texts were sampled from Wikipedia with probabilities proportional to  $n_{lang}^{1/5}$ , where  $n_{lang}$  is the size of Wikipedia in that language<sup>15</sup>. To increase the recall for Erzya and Moksha languages, we augment this training dataset with Erzya and Moksha Bible texts. The resulting model has 89% accuracy and 86% macro F1 score on the Wikipedia test set (sampled with the same temperature). For Erzya, it has 97% precision and 82% recall. Hyperparameters for all trained models are listed in Appendix B.

### 3.3 Erzya Sentence Encoder

To compute sentence embeddings, we use an encoder based on LaBSE ([Feng et al., 2022](#)), with an extended vocabulary. First, we use the BPE algorithm ([Sennrich et al., 2016](#)) over a monolingual

<sup>9</sup><http://lib.e-mordovia.ru>

<sup>10</sup><https://fennougrica.kansalliskirjasto.fi>

<sup>11</sup>[https://wikisource.org/wiki/Main\\_Page/?oldid=895127](https://wikisource.org/wiki/Main_Page/?oldid=895127)

<sup>12</sup><https://rus4all.ru/myv/>

<sup>13</sup><http://erziapr.ru>

<sup>14</sup><https://www.livejournal.com>

<sup>15</sup>We adopted the idea of temperature sampling with  $T=5$  from [Tran et al. \(2021\)](#) and several other works.

Erzya corpus to add 19K extra merged tokens to the vocabulary. Then, we fine-tune the model on the limited initial parallel data (the Bible, OPUS, and dictionaries): we update only the token embeddings matrix, using the contrastive loss from [Feng et al. \(2022\)](#) over computed sentence embeddings. Finally, after collecting more parallel sentences, we fine-tune the full model on a mixture of tasks: contrastive loss over sentence embeddings, standard masked language modeling loss, and sentence pair classification to distinguish correct translations from random pairs.

### 3.4 Mining Parallel Sentences

When mining parallel sentences, we strive for high precision. To compensate for the questionable quality of our sentence encoder, we apply the following procedure<sup>16</sup>.

- We perform only local mining, i.e. we compare sentences only across paired documents (for Wikipedia and translated books), or within one document (for the web sources).
- To evaluate similarity of two sentences, we multiply the cosine similarity between their LaBSE embeddings by the ratio of the length of the shortest sentence to that of the longest one.
- We further penalize the similarities by partially subtracting from them the average similarities of each sentence to its closest neighbors, similarly to using distance margin from [Artetxe and Schwenk \(2019\)](#).
- Given two documents in Russian and Erzya, we use dynamic programming to select a sequence of sentence pairs that have the maximal sum of pairwise similarity scores and go in the same order in both documents.
- We accept only the sentence pairs with a score above a threshold, which was manually tuned for each source of texts.

In total, this approach yielded 21K more unique parallel sentence pairs. The manual inspection found that more than 90% of them were matched correctly.

### 3.5 Training Machine Translation Models

To benefit from transfer learning, we base our model on the mBART50 model ([Tang et al., 2020](#)) pretrained on multiple languages, including two

Uralic ones (Finnish and Estonian). We extend its BPE vocabulary with 19K new Erzya tokens, using the same method as in Section 3.3, and add a new myv\_XX language code to it. Embeddings for the new tokens are initialized as the averages of the embeddings of the Russian tokens aligned with them in the parallel corpus<sup>17</sup>, inspired by [Xu and Hong \(2022\)](#).

We make two copies of this model and train them to translate in the myv-ru and ru-myv directions, respectively. The myv-ru model is trained on the joint parallel corpus of sentences and words. The ru-myv model is trained on the union of this corpus and the back-translated corpus generated by the myv-ru model from the monolingual myv data.

After training the models on these two languages, we adapt them to 10 more languages: ar, de, en, es, fi, fr, hi, tr, uk, and zh, resulting in the myv-mul and mul-myv models (below, by mul we denote any of these 10 languages). We fine-tune the two models jointly, using a version of online-back translation and self-training. Specifically, we generate the training pairs in four alternating steps:

1. Sample a ru-mul sentence pair from the CCMatrix ([Schwenk et al., 2021](#)) dataset, translate from ru to myv with the mul-myv model;
2. Sample a ru-mul pair from the CCMatrix, translate from mul to myv with the mul-myv model;
3. Sample a ru-myv pair from our parallel corpus, translate from myv to mul with the myv-mul model;
4. Sample a myv text from the monolingual myv corpus, translate from mul to myv and ru with the myv-mul model.

At each step, we update both models on the myv-mul and myv-ru pairs in both directions. For the self-training updates, we multiply the loss by the coefficient  $\lambda_{ST} = 0.05$  to decrease the impact of self-training relatively to back-translation (the choice of the coefficient is suggested by experiments in [He et al. \(2022\)](#)).

During the initial experiments, we noticed that, when translating from Russian to Erzya, the model often just copied Russian phrases with only word

<sup>17</sup>We compute alignments with a naive formula: the alignment weight between tokens  $t_i$  and  $t_j$  is estimated as  $\frac{n_{ij}^2}{n_i n_j}$ , where  $n_i$  and  $n_j$  are their respective frequencies in the parallel corpus, and  $n_{ij}$  is the number of sentence pairs with  $t_i$  in one sentence and  $t_j$  in another.

<sup>16</sup>For more details on the mining procedure, please read the source code in the repository that we release.

endings sometimes changed. Sometimes this is acceptable because Erzya has multiple Russian loanwords, but often there exist native words that are preferable. To alleviate this problem, in step 1 we generate 5 alternative ru-myv translations using diverse beam search (Vijayakumar et al., 2016), and choose the one with the largest proportion of words recognized as myv by our language identification model. This problem was also the reason why we chose to train two different models from translation to Erzya and from Erzya: this way, the decoder and encoder of a model never work with the same language.

## 4 Evaluation

### 4.1 Data

For model evaluation, we prepare a held-out corpus of 3000 aligned Erzya-Russian sentences from 6 diverse sources: the Bible, Erzya folk tales (Sheyanova, 2017), the Soviet 1938 constitution, descriptions of folk children’s games (Брыжинский, 2009), modern Erzya fiction and poetry, and Wikipedia. To evaluate English and Finnish translation, we use translations from the Erzya universal-dependency treebank (Rueter and Tyers, 2018): 441 sentence pairs for en, and 309 for fi. We split all these sets into development and test parts, and report the results on the test set.

### 4.2 Automatic Metrics

For all evaluated directions (between myv and ru, en, fi) we calculate BLEU (Papineni et al., 2002) and ChrF++ (Popović, 2017). Both these metrics estimate the proportion of common parts in the translation and the reference, but BLEU is calculated as precision over word n-grams, whereas ChrF++ aggregates precision and recall of word and character n-grams (which is more suitable for morphologically rich languages such as Erzya and Russian). The values of these metrics on the test set are given in Table 1. For translation from and to Russian, the BLEU scores are 17 and 19 points, respectively. For English and Finnish, however, BLEU is well below 10. We hypothesize that the low quality may be attributed to the domain mismatch between the Erzya-origin and English- or Finnish-origin training corpora, but without detailed test sets we cannot verify this.

For the Russian test set, the performance varies greatly depending on the domain (Table 2). The constitution has the highest scores because its Erzya

Direction	BLEU	ChrF++
ru-myv	17.71	41.16
myv-ru	19.68	38.63
en-myv	2.77	28.03
myv-en	5.44	25.99
fi-myv	4.79	27.42
myv-fi	3.02	22.34

Table 1: Reference-based scores on the test sets.

Source	ru-myv		myv-ru	
	BLEU	ChrF++	BLEU	ChrF++
bible	10.00	36.92	10.71	33.55
tales	7.00	33.90	7.30	28.42
constitution	27.82	62.96	33.31	60.60
games	10.33	31.19	9.85	26.57
fiction	8.68	30.59	5.95	26.60
wiki	28.39	48.56	32.24	47.55

Table 2: Scores by section on the myv-ru test set.

text is saturated with Russian loanwords and is easy to generate and understand. For Wikipedia, the scores are also high, probably because its Erzya articles are often translated from Russian in a rather literal way. The other domains have a more artistic style, and the translations are on average much less literal.

Some examples of the translations and references are given in Table 3.

### 4.3 Manual Evaluation

We recruit three native speaker volunteers to evaluate some translations manually. The evaluation protocol is similar to XSTS (NLLB Team et al., 2022), but evaluates fluency in addition to semantic similarity. The scores are between 1 (a useless translation) and 5 (a perfect translation), with 3 points standing for an acceptable translation without serious errors. Criteria for each score are given in Appendix C.

Each of the 3 annotators rated a few randomly sampled translations from the dev split of each source: 12 pairs in the ru-myv and 17 pairs in the myv-ru directions, which amounts to 87 sentence pairs annotations in total. The average length of the labelled texts was 97 characters, or 14 words.

It turned out that, despite the specified annotation criteria, the annotators were calibrated very differently: their average ratings were 2.9, 3.5, and 4.1. We chose a pessimistic aggregation strategy: for each of the 29 evaluated sentence pairs, we took the worst of the scores by our 3 volunteers.

For translation to Erzya, the average pessimistic score was 2.75, and 58% translations were rated as acceptable (i.e. all the 3 reviewers rated them

Type	Text
source (ru)	И вот что рассказывают наши русские старики: «Когда здесь не было этого села, в этом овраге были разбойники; у них были землянки.
source (myv)	Вана мезе ёвтнить миненек рузонь аятне: зярдо велесь тесэ арасель, се латксонть эрясть розбойникть, эрясть землянкасо.
translation (ru-myv)	Ды вана мезе ёвтнить минек эрянь аятне: «Зярдояк те велесэнь арасель, тосо оврагасо ульнесь розбойникть; сынст ульнесь землянкаст.
translation (myv-ru)	Вот что нам говорят русские старцы: когда деревня здесь не была, то там жилибойники, жили на земле.
source (myv)	Кода авазо, анокстась лапужа кирькст, истя жо педавтинзе, валаськавтинзе педавтома таркатнень начко кедьсэ.
source (en)	Like his mother, he prepared flat rings, and stuck them onto the patty in the same way, and smoothed out the seams with his wet hands.
source (fi)	Samalla tavalla kuin äitinsä Ketšai valmisti litteitä rinkuloita liitti ne samalla tavalla, ja siloitti liitoksen määrällä kädellä.
translation (myv-en)	Like his mother, he prepared flat circles, and also filled the canvas with a needle.
translation (myv-fi)	Kuten äiti, valmistelee tasa-alaiset kentät, myös venytetään, lyödään venyttäjän käsillä.
translation (en-myv)	Кода аванзо, сон анокстыль валаня суркст, теке ладсо педавтызе сынст пацьказонзо ды вадяшась кедень летькенть марго.
translation (fi-myv)	Истя жо, кода авазо Кетшай анокстыль лантсо кевпанть, сон солодиль сынст теке ладсо ды солодиль эйсэст кедьлапушкасо.

Table 3: A few examples of translations and references.

with at least 3 points). For translation to Russian, the average score was 2.71, with 53% acceptable translations.

An additional comment from the annotators was that some of the source Erzya texts were inadequate. In particular, some games sentences contained grammatical errors<sup>18</sup>, and most constitution sentences contained Russian words with Erzya endings instead of their Erzya equivalents. This suggests that one of the next steps in improving our NMT system might be to filter the training and evaluation data for better language quality.

## 5 Conclusions and Future Work

In this paper, we present the first NMT system for the endangered Erzya language, capable of translating between it and 11 diverse languages, primarily Russian. During its development, we have collected about 30K parallel Russian-Erzya sentences and 300K monolingual Erzya sentences, and trained a language identification model and a BERT-based sentence encoder that support Erzya. All the resources are publicly released. These efforts have occupied about two man-weeks of working time and almost no expenses<sup>19</sup>. We hope that these results will inspire the NLP community to develop resources for other endangered languages.

<sup>18</sup>We are not certain whether these errors are due to the low quality of the source text, or to the natural variations within the Erzya language.

<sup>19</sup>All the expenses incurred totalled \$9.99 for the paid subscription to the Google Colab system (<https://colab.research.google.com/signup>).

The quality of our system may be improved by collecting more texts in Erzya and filtering them better than we did. Another promising direction is a more efficient usage of the vocabularies and parsers that are already available for the language, e.g. for generating synthetic training data. Finally, we hope to attract more native speakers for creating larger and cleaner train and test datasets.

One open research question is that of transfer between languages: whether Erzya translation benefits from knowledge of, for example, Hungarian or Estonian, and whether knowledge of Erzya can bring improvements to other languages, such as Moksha. In further studies, we hope to shed some light on this direction as well.

## 6 Acknowledgements

We gratefully acknowledge support from the volunteers who participated in the manual evaluation of translation quality: Semyon Tumaikin, Zinyoronj Santyai, and Evgenia Chugunova. We are also grateful to the reviewers for their helpful suggestions which substantially improved this work.

## References

- Fenno-ugrica. <https://fennougrica.kansalliskirjasto.fi/>. Accessed: 2022-08-01.
- Livejournal. <https://www.livejournal.com>. Accessed: 2022-08-01.
- a. The wikimedia dump service. <https://dumps.wikimedia.org/myvwiki/20220601/>. Accessed: 2022-08-01.

- b. Wikisource. [https://wikisource.org/wiki/Main\\_Page/?oldid=895127](https://wikisource.org/wiki/Main_Page/?oldid=895127). Accessed: 2022-08-01.
- Портал национальных литератур. Эрзянский язык. <https://rus4all.ru/myv/>. Accessed: 2022-08-01.
- Эрзянь Правда. <http://erziapr.ru/>. Accessed: 2022-08-01.
1938. Советской Социалистической Республикатненъ Союзонъ Конституциясь (Основной Законсь). Мордгиз, Саранск.
- Judit Ács, Dániel Lévai, and Andras Kornai. 2021. Evaluating transferability of BERT models on uralic languages. In *Proceedings of the Seventh International Workshop on Computational Linguistics of Uralic Languages*, pages 8–17, Syktyvkar, Russia (Online). Association for Computational Linguistics.
- Khalid Alnajjar. 2021. When word embeddings become endangered. *arXiv preprint arXiv:2103.13275*.
- Mikel Artetxe and Holger Schwenk. 2019. Margin-based parallel corpus mining with multilingual sentence embeddings. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3197–3203, Florence, Italy. Association for Computational Linguistics.
- Ankur Bapna, Isaac Caswell, Julia Kreutzer, Orhan Firat, Daan van Esch, Aditya Siddhant, Mengmeng Niu, Pallavi Baljekar, Xavier Garcia, Wolfgang Macherey, et al. 2022. Building machine translation systems for the next thousand languages. *arXiv preprint arXiv:2205.03983*.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2022. Language-agnostic BERT sentence embedding. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.
- finugorbib.com. <http://finugorbib.com/>. Accessed: 2022-08-01.
- Mika Härmäläinen and Khalid Alnajjar. 2019. A template based approach for training nmt for low-resource uralic languages - a pilot with finnish. In *Proceedings of the 2019 2nd International Conference on Algorithms, Computing and Artificial Intelligence*, ACAI 2019, page 520–525, New York, NY, USA. Association for Computing Machinery.
- Zhiwei He, Xing Wang, Rui Wang, Shuming Shi, and Zhaopeng Tu. 2022. Bridging the data gap between training and inference for unsupervised neural machine translation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6611–6623, Dublin, Ireland. Association for Computational Linguistics.
- Mika Härmäläinen. 2019. UralicNLP: An NLP library for Uralic languages. *Journal of Open Source Software*, 4(37):1345.
- Boglárka Janurik. 2017. Erzya-russian bilingual discourse: A structural analysis of intrasentential code-switching patterns. *Szeged: Szegedi Tudományegyetem PhD dissertation. Manuscript*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2016. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759*.
- M Paul Lewis and Gary F Simons. 2010. Assessing endangerment: expanding fishman’s gids.
- lib.e mordovia.ru. <http://lib.e-mordovia.ru>. Accessed: 2022-08-01.
- marlamuter.com. Эрзя-Руш мутер. <http://marlamuter.com/muter/Эрзя>. Accessed: 2022-08-01.
- mordovians.ru. Русско-Эрзянский словарь. [https://www.mordovians.ru/erzyanskiy\\_yazyk](https://www.mordovians.ru/erzyanskiy_yazyk). Accessed: 2022-08-01.
- mordvarf.com. Русско-Эрзянский словарь. <https://mordvarf.com/русско-эрзянский-словарь/>. Accessed: 2022-08-01.
- Benjamin Muller, Antonios Anastasopoulos, Benoît Sagot, and Djamé Seddah. 2021. When being unseen from mBERT is just the beginning: Handling new languages with multilingual language models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 448–462, Online. Association for Computational Linguistics.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia-Gonzalez, Prangthip Hansanti, John Hoffman, Semaire Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. No language left behind: Scaling human-centered machine translation.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

- Maja Popović. 2017. [chrF++: words helping character n-grams](#). In *Proceedings of the Second Conference on Machine Translation*, pages 612–618, Copenhagen, Denmark. Association for Computational Linguistics.
- Jack Rueter. 2013. [The erzya language. where is it spoken?](#) *Études finno-ougriennes*, (45).
- Jack Rueter and Francis Tyers. 2018. [Towards an open-source universal-dependency treebank for Erzya](#). In *Proceedings of the Fourth International Workshop on Computational Linguistics of Uralic Languages*, pages 106–118, Helsinki, Finland. Association for Computational Linguistics.
- Holger Schwenk, Guillaume Wenzek, Sergey Edunov, Edouard Grave, Armand Joulin, and Angela Fan. 2021. [CCMatrix: Mining billions of high-quality parallel sentences on the web](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6490–6500, Online. 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.
- Irina Ivanovna Sheyanova. 2017. *Skaz forms of Moravian literature*. The Research Institute of the Humanities by the Government of the Republic of Moravia, Saransk.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. Multilingual translation with extensible multilingual pretraining and finetuning. *arXiv preprint arXiv:2008.00401*.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- Chau Tran, Shruti Bhosale, James Cross, Philipp Koehn, Sergey Edunov, and Angela Fan. 2021. [Facebook AI’s WMT21 news translation task submission](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 205–215, Online. Association for Computational Linguistics.
- UNESCO. 2010. [Unesco atlas of the world’s languages in danger \(pdf\)](#).
- Ashwin K Vijayakumar, Michael Cogswell, Ramprasath R Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2016. Diverse beam search: Decoding diverse solutions from neural sequence models. *arXiv preprint arXiv:1610.02424*.
- Xinyi Wang, Sebastian Ruder, and Graham Neubig. 2022. [Expanding pretrained models to thousands more languages via lexicon-based adaptation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 863–877, Dublin, Ireland. Association for Computational Linguistics.
- Minhan Xu and Yu Hong. 2022. [Sub-word alignment is still useful: A vest-pocket method for enhancing low-resource machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 613–619, Dublin, Ireland. Association for Computational Linguistics.
- Тимофей Александрович Архангельский. 2019. [Интернет-корпуса финно-угорских языков России](#). *Ежегодник финно-угорских исследований*, (3):528–537.
- Василий Сергеевич Брыжинский. 2009. *Мордовские народные игры*. Мордовское кн. издательство, Saransk.
- М.Е. Евсевьев. 1964. Мордовские народные сказки и загадки. *Евсевьев М.Е. Избранные труды*, 3:412.
- И.Н Рябов, Е.Ф. Клементьева, and Г.В. Рябова. 2011. *Эрзянский язык. Учебное пособие*. Издательство Мордовского университета, Саранск.
- Валентина Ивановна Щанкина. 2011. *Русско-моксианско-эрзянский словарь*. Поволжский центр культур финно-угорских народов, Саранск.
- Борис Эрюшов. Русско-Эрзянский словарь. <http://lazalyk.narod.ru/business.html>. Accessed: 2022-08-01.

## A Data sources

Source	Type	Size
Erzya-Russian dictionaries: <a href="http://marlamuter.com">marlamuter.com</a> , <a href="http://mordovians.ru">mordovians.ru</a> , <a href="http://mord-varf.com">mord-varf.com</a> , Рябов et al. (2011), Щанкина (2011), Эрюшов	phrase pairs	47860
The myv-ru Wikimedia corpus on OPUS (Tiedemann, 2012)	sentence pairs	3202
The Bible ( <a href="http://finugorbib.com">finugorbib.com</a> )	sentence pairs	12483
<a href="#">Sheyanova (2017)</a> (aligned)	sentence pairs	1023
<a href="#">Брыжинский (2009)</a> (aligned)	sentence pairs	4203
Erzya and Russian Wikipedia ( <a href="#">wik</a> , <a href="#">a</a> ) (aligned)	sentence pairs	11479
Livejournal ( <a href="#">lj</a> ) (aligned)	sentence pairs	1799
Modern Erzya fiction and poetry ( <a href="#">rus</a> ) (aligned)	sentence pairs	916
The Soviet 1938 constitution ( <a href="#">con</a> , 1938) (aligned)	sentence pairs	304
Mordovian tales and riddles ( <a href="#">Евсеев, 1964</a> ) (aligned)	sentence pairs	3776
Various Erzya fiction books ( <a href="http://lib.e.mordovia.ru">lib.e.mordovia.ru</a> )	sentences	52870
Various Soviet-time books and periodicals ( <a href="#">fen</a> )	sentences	54798
Erzya Wikisource, filtered by language ( <a href="#">wik</a> , <a href="#">b</a> )	sentences	120470
Articles from the Erzya Pravda website ( <a href="#">pra</a> )	sentences	43772
Livejournal ( <a href="#">lj</a> )	sentences	36584
Erzya Wikipedia ( <a href="#">wik</a> , <a href="#">a</a> )	sentences	59569
<a href="#">Брыжинский (2009)</a>	sentences	5194

Table 4: The sources used to construct the training and evaluation datasets. The “size” column denotes the number of sentences or phrases in the source.

## B Models’ hyperparameters

### B.1 Language identification

For the language identification model, we use the official FastText implementation<sup>20</sup>. We train it with initial learning rate of 0.05 for 100 epochs, using minimum word count of 100, 64-dimensional embeddings and 200K hash buckets for character n-grams with n from 1 to 4. Then we quantize the model with retraining on the same dataset, a cutoff of 50000, and norm pruning.

### B.2 Sentence encoder

For the sentence encoder model, we use a PyTorch port of LaBSE<sup>21</sup>, in which we remove tokens for all languages, except Russian and English, and add Erzya tokens. For vocabulary extension, we set the minimal token count for stopping BPE at 30.

After extending the vocabulary, we fine-tune the model on the initial parallel sentences and phrases using the LaBSE contrastive loss with margin 0.3 and batch size 4 for 500K steps, updating only the embeddings, and passing the gradient only through the encoded myv sentence. We use the Adafactor optimizer with learning rate of  $10^{-5}$  and clipping the gradient norm at 1. Then we update the model for 500K steps with learning rate  $2 \times 10^{-6}$ , updating all the parameters, and alternating batches with the LaBSE loss, MLM loss, and the loss of classifying the correct and incorrect sentence pairs. Incorrect pairs are generated either by sampling one of the sentences randomly, or by randomly inserting, deleting, or swapping words in one of the sentences in a correct parallel pair.

### B.3 Machine translation models

Both myv-ru and ru-myv models were initialized from mBART50<sup>22</sup> with the vocabulary extended with Erzya tokens. They were trained with Adafactor optimizer using batch size of 8 and learning rate of  $10^{-6}$

<sup>20</sup><https://github.com/facebookresearch/fastText>

<sup>21</sup><https://huggingface.co/sentence-transformers/LaBSE>

<sup>22</sup><https://huggingface.co/facebook/mbart-large-50-many-to-many-mmt>



for 4 epochs: on the first epoch, only token embeddings were updated, and on the remaining epochs, all parameters were updated.

The myv-mul and mul-myv models were initialized from myv-ru and ru-myv, respectively. They were jointly trained for 40K updates with batch size of 1.

For inference, we used beam size of 5 and repetition penalty of 5.0.

Both the sentence encoder and the translation models were trained using the PyTorch<sup>23</sup> and Transformers<sup>24</sup> Python packages.

## C Quality annotation guidelines

The following annotation criteria (in Russian) were suggested to the annotators in Section 4.3.

- 5 points: a perfect translation. The meaning and the style are reproduced completely, the grammar and word choice are correct, the text looks natural.
- 4 points: a good translation. The meaning is reproduced completely or almost completely, the style and the word choice are natural for the target language.
- 3 points: an acceptable translation. The general meaning is reproduced; the mistakes in word choice and grammar do not hinder understanding; most of the text is grammatically correct and in the target language.
- 2 points: a bad translation. The text is mainly understandable and mainly in the target language, but there are critical mistakes in meaning, grammar, or word choice.
- 1 point: a useless translation. A large part of the text is in the wrong language, or is incomprehensible, or has little relation to the original text.

---

<sup>23</sup><https://pytorch.org>

<sup>24</sup><https://huggingface.co/docs/transformers/>