

Sentiment Analysis Using Aligned Word Embeddings for Uralic Languages

Khalid Alnajjar

Rootroo Ltd
Helsinki, Finland
name@rootroo.com

Mika Hämäläinen

Rootroo Ltd
Helsinki, Finland
name@rootroo.com

Jack Rueter

University of Helsinki
Finland
first.last@helsinki.fi

Abstract

In this paper, we present an approach for translating word embeddings from a majority language into 4 minority languages: Erzya, Moksha, Udmurt and Komi-Zyrian. Furthermore, we align these word embeddings and present a novel neural network model that is trained on English data to conduct sentiment analysis and then applied on endangered language data through the aligned word embeddings. To test our model, we annotated a small sentiment analysis corpus for the 4 endangered languages and Finnish. Our method reached at least 56% accuracy for each endangered language. The models and the sentiment corpus will be released together with this paper. Our research shows that state-of-the-art neural models can be used with endangered languages with the only requirement being a dictionary between the endangered language and a majority language.

1 Introduction

Most of the languages spoken in the world are endangered to one degree or another. The fact of being endangered sets some limitations on how modern NLP research can be done with such languages given that many endangered languages do not have vast textual resources available online, and even with the resources that are available, there is a question about the quality of the data resulting from a variety of factors such as fluency of the author, soundness of spelling and, on the lowest level, inconsistencies in character encoding (see Hämäläinen 2021).

This paper focuses on the following Uralic languages: Erzya (myv), Moksha (mdf), Komi-Zyrian (kpv) and Udmurt (udm). Unesco classifies

these languages as definitely endangered (Moseley, 2010). In terms of NLP, these languages have FSTs (Rueter et al., 2020, 2021), Universal Dependencies Treebanks (Partanen et al., 2018; Rueter and Tyers, 2018) (excluding Udmurt) and constraint grammars available in Giella repositories (Moshagen et al., 2014). For some of the languages, there have also been efforts in employing neural models in disambiguation (Ens et al., 2019) and morphological tasks (Hämäläinen et al., 2021). Out of these languages, only Erzya has several neural based models available such as machine translation models (Dale, 2022), a wav2vec model and a Stanza model (Qi et al., 2020).

In this paper, we present a method for translating word embeddings models from larger languages into the endangered languages in question. Furthermore, we fine-tune the models with language specific text data, align them and show results in a sentiment analysis task where no training data is provided in any of the endangered languages. We have made our data and models publicly available on Zenodo¹.

2 Related work

Apart from the work described earlier in the context of the endangered languages in question, there has been a lot of previous work on multilingual NLP where a model is trained in one language to sentence classification and then applied in the context of other languages. In this section, we will describe some of those approaches together with sentiment analysis approaches.

A recent paper demonstrates sentiment analysis on 100 languages (Yilmaz et al., 2021). The authors use RoBERTa-XLM to extract feature vectors. These are then used in training a bi-directional LSTM based classifier model. Another line of work (Liu and Chen, 2015) compares

¹<https://zenodo.org/record/7866456>

several different multilabel classification methods on the task of sentiment analysis showing that RAKEL (Tsoumakas et al., 2010) gave the best performance on raw token input. A recent paper (Hämäläinen et al., 2022) demonstrated promising results in French sentiment analysis on a model that was trained in English, Italian, Spanish and German. The approach relied on a multilingual BERT (Devlin et al., 2019). Öhman (2021) suggests that lexicon based approaches, while viable for endangered languages, are not particularly suitable for sentiment analysis.

In the context of cross-lingual NLP, there is work on POS tagging. For instance, Kim et al. 2017 propose a new model that does not require parallel corpora or other resources. The model uses a common BLSTM for knowledge transfer and another BLSTM for language-specific representations. It is trained using language-adversarial training and bidirectional language modeling as auxiliary objectives to capture both language-general and language-specific information.

Another line of work by Xu et al. 2018 focuses on cross-lingual transfer of word embeddings, which aims to create mappings between words in different languages by learning transformation functions over corresponding word embedding spaces. The proposed algorithm simultaneously optimizes transformation functions in both directions by using distributional matching and minimizing back-translation losses. This approach uses a neural network implementation to calculate the Sinkhorn distance, a distributional similarity measure, and optimize objectives through back-propagation.

For machine translation Chen et al. 2022 demonstrate the importance of both multilingual pretraining and fine-tuning for effective cross-lingual transfer in zero-shot translation using a neural machine translation (NMT) model. The paper presents SixT+, a many-to-English NMT model that supports 100 source languages but is trained on a parallel dataset in only six languages. SixT+ initializes the decoder embedding and full encoder with XLM-R large (Conneau et al., 2020) and trains encoder and decoder layers using a two-stage training strategy.

3 Data

We use two books, *Suomi eilen ja nyt* (*Finland yesterday and now*) by Häkkinen (1997) and

Павлик Морозов (Pavlik Morozov) by Gubarev (1948) both of which are available in Finnish, Erzya, Moksha, Komi-Zyrian and Udmurt. The sentences of the books have been aligned across all the languages at the Research Unit for Volgaic Languages in University of Turku. The size of the corpus for each language can be seen in Table 1.

	tokens	sentences
Finnish	43k	3.1k
Erzya	50k	3.6k
Moksha	51k	3.4k
Komi-Zyrian	50k	3.3k
Udmurt	53k	3.6k

Table 1: The corpus size for each language

Out of the entire corpus, we annotate 35 negative sentences and 33 positive sentences for evaluation for Finnish. We use the alignment information to project this annotation for the rest of the languages as well and verify manually that the sentences express the same sentiment in each language. This forms our test corpus for sentiment analysis that consists of altogether 68 sentiment annotated sentences.

Furthermore, we lemmatize all the texts using the FSTs provided in UralicNLP (Hämäläinen, 2019). The corpus is lemmatized because we intend to translate and align a lemmatized word embeddings model. This also makes the overall approach more robust given that covering the entire morphology of a language would require us to have much larger corpora.

4 Word embeddings

Word embeddings capture the semantic and syntactic links between words by constructing vector representations of words. These vectors can be utilized to measure the semantic similarity between words, find analogous concepts, cluster words (Hämäläinen and Alnajjar, 2019; Stekel et al., 2021) and more. In this work, we use English and Finnish as the big languages that facilitate aligning and classifying words and sentences for the endangered languages. English has an overnumerous amount of linguistic resources, whether as raw text or labeled data, while the endangered resources that we are working with have translation dictionaries for Finnish. For this reason, we use Finnish as the intermediate language that bridges these endangered languages with English resources.

The English model that we utilize is trained on the English Wikipedia dump of February 2017 and Gigaword 5th edition² (Fares et al., 2017). For Finnish, we used recent word embeddings trained by Language Bank of Finland (2022). These embeddings have been trained on several Finnish newspapers. Both of these models have been trained on lemmatized text.

The English word vectors have a dimension size of 300, while the Finnish word vectors have a dimension size of 100. In order to make the dimension sizes of the two sets of embeddings compatible, dimensionality reduction is applied to the English embeddings using principal component analysis (PCA) (Tipping and Bishop, 1999). This process reduces the dimensionality of the English embeddings to 100, allowing them to be compared and analyzed alongside the Finnish embeddings.

4.1 Creation of embeddings

We aim to create word embeddings for endangered languages, which currently lack pre-existing embeddings. We use dictionaries from GiellaLT³, which we augment using graph-based methods to predict new translations through the Ve’rdd⁴ platform (Alnajjar et al., 2022, 2021). We present the number of dictionary translations from each endangered language to Finnish that we obtained from the base dictionaries and predictions in Table 2.

	Translations	Predictions	Total
kpv	10983	14421	25404
mdf	36235	3903	40138
myv	18056	5018	23074
udm	36502	6966	43468

Table 2: Number of translations and predictions from the source languages to Finnish

To create embeddings for the endangered languages, we adopt a method of cloning the Finnish embeddings and substituting the Finnish lemma with its corresponding translation in the endangered language. Where translations were absent, we omitted the word vector. The resulting embeddings consist of 7,908, 10,338, 7,535, and 9,505 word vectors for kpv, mdf, myv, and udm, respectively. The lower number of word coverage can be

²<http://vectors.nlp1.eu/repository/20/17.zip>

³<https://github.com/giellalt>

⁴<https://akusanat.com/verdd/>

attributed to multi-word expressions present in the dictionaries but not the embeddings.

In the next step of our study, we fine-tuned the word embeddings for both Finnish and the endangered languages by using two books as additional data sources. This involved expanding the vocabulary of each embeddings model whenever a new word was encountered in the data. We also adjusted the embeddings weights based on the co-occurrences of words in the text, using a window size of 5 and a minimum count of 5 for a word to be considered in the vocabulary. After completing this process, the vocabulary size of the endangered language embeddings were 10,396, 11,877, 9,030, and 11,080, in the same order as mentioned above.

4.2 Alignment of embeddings

Our goal here is to align the Finnish word embeddings with the English ones, followed by aligning the embeddings of endangered languages to the Finnish embeddings, in a supervised manner. This was achieved by creating alignment dictionaries and aligning the embedding spaces together similarly to Alnajjar (2021).

To align Finnish embeddings with English, we used the Fin-Eng dictionary by Ylönen (2022), which is based on the March 2023 English Wiktionary dump. We also used the Finnish-English dictionaries provided by MUSE (Conneau et al., 2017). Regarding the endangered languages, we use the XML dictionaries to align them with Finnish. We set aside 20% of the Wiktionary and XML data for testing the alignments.

One thing that we have noticed is the lack of the words “no” and “not” in the English embeddings due to stopword removal. To address this, we appended a translation from “not” to “nt” in the Finnish-English alignment data used in the training stage. Whenever the text contained these words, they were automatically mapped to “nt” in the following steps of our research.

We followed the approach described by MUSE (Conneau et al., 2017) to align all the embeddings, with 20 iterations of refinement to align Finnish with English and 5 iterations to align all the other languages to Finnish.

5 Sentence embeddings

Word embeddings represent the meaning of a single word, whereas sentence embeddings represent the meaning of an entire sentence or docu-

Language	Label	Precision	Recall	F1-Score	Accuracy
eng	neg	0.77	0.76	0.76	0.76
	pos	0.75	0.76	0.76	
fin	neg	0.77	0.75	0.76	0.75
	pos	0.73	0.75	0.74	
kpv	neg	0.57	0.57	0.57	0.56
	pos	0.55	0.55	0.55	
mdf	neg	0.63	0.65	0.64	0.63
	pos	0.64	0.62	0.63	
myv	neg	0.71	0.69	0.70	0.69
	pos	0.67	0.69	0.68	
udm	neg	0.69	0.63	0.66	0.63
	pos	0.58	0.63	0.60	

Table 3: Precision, recall, f1-score and accuracy for each language and label

ment. Sentence embeddings are capable of capturing more the context and excel at tasks that call for comprehension of the meaning of a whole text, such as sentiment analysis. Hence, we build sentence embeddings for English that are based on the English word embeddings.

The procedure for creating sentence embeddings was conducted by averaging the word embeddings of a given sentence and subsequently feeding them to two fully-connected feed-forward layers, thereby constructing a Deep Averaging Network (DAN). The sentence embeddings are trained on the STS Benchmark (Cer et al., 2017) using SBERT, a method for sentence embeddings that was proposed by (Reimers and Gurevych, 2019).

6 Sentiment analysis

We create a sentiment classifier that takes in the sentence embeddings and predicts a sentiment polarity label. For training the sentiment analysis model, we use the Stanford Sentiment Treebank (Socher et al., 2013), Amazon Reviews Dataset (McAuley and Leskovec, 2013) and Yelp Dataset⁵. These datasets are available in English and we use their sentiment annotations (positive-negative) to train our model.

The sentiment classifier is constructed as a three-layer fully-connected network, wherein the hidden layers are comprised of 300 neurons each. In order to mitigate overfitting, a dropout operation (Srivastava et al., 2014) is performed prior to the final classification layer. The model consists of 121,202 trainable parameters in total, and is trained over the course of three epochs.

⁵<https://www.yelp.com/dataset>

7 Results

In this section, we show the results of the sentiment classification model on the in-domain, English-language train splits of the sentiment corpora we used to train the model. Furthermore, we show the results of the sentiment classification model when applied on our own annotated data for the 4 endangered Uralic languages in question and Finnish. These results can be seen in Table 3.

All in all, our model performs relatively well. The accuracy for Finnish is almost as high as it is for English despite not having any Finnish sentiment annotated training data. This means that our approach can achieve rather good results when there is a lot of translation data available between the two languages. The results drop for the endangered languages, but we do find the 69% accuracy for Erzya to be quite formidable, however, the result for Komi-Zyrian of 56% leaves some room for improvement.

8 Conclusions

In this paper, we outlined a method for translating word embeddings from a majority language, Finnish, to four minority languages - Erzya, Moksha, Udmurt, and Komi-Zyrian. The word embeddings were aligned and a new neural network model was introduced. This model was trained using English data to carry out sentiment analysis and was then applied to data in the endangered languages using the aligned word embeddings.

We built an aligned sentiment analysis corpus for the four endangered languages and Finnish and used it to test our model. The results were promising and our study demonstrated that even the latest neural models can be utilized with endangered languages if a dictionary between the endangered language and a larger language is available.

Acknowledgments

This research is supported by FIN-CLARIN and Academy of Finland (grant 345610 *Kielivarojen ja kieliteknologian tutkimusinfrastruktuuri*).

References

- Khalid Alnajjar. 2021. When word embeddings become endangered. In *Multilingual Facilitation*, pages 275–288. Rootroo Ltd.
- Khalid Alnajjar, Mika Hämäläinen, Niko Tapio Partanen, and Jack Rueter. 2022. Using graph-based methods to augment online dictionaries of endangered languages. In *Proceedings of the Fifth Workshop on the Use of Computational Methods in the Study of Endangered Languages*, pages 139–148.
- Khalid Alnajjar, Jack Rueter, Niko Partanen, and Mika Hämäläinen. 2021. Enhancing the erzya-moksha dictionary automatically with link prediction. *Folia Uralica Debreceniensia*, 28:7–18.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Guanhua Chen, Shuming Ma, Yun Chen, Dongdong Zhang, Jia Pan, Wenping Wang, and Furu Wei. 2022. Towards making the most of cross-lingual transfer for zero-shot neural machine translation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 142–157, Dublin, Ireland. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.
- David Dale. 2022. The first neural machine translation system for the erzya language. In *Proceedings of the first workshop on NLP applications to field linguistics*, pages 45–53.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, pages 4171–4186, Minneapolis, Minnesota.
- Jeff Ens, Mika Hämäläinen, Jack Rueter, and Philippe Pasquier. 2019. Morphosyntactic disambiguation in an endangered language setting. In *Proceedings of the 22nd Nordic Conference on Computational Linguistics*, pages 345–349, Turku, Finland. Linköping University Electronic Press.
- Murhaf Fares, Andrey Kutuzov, Stephan Oepen, and Erik Velldal. 2017. Word vectors, reuse, and replicability: Towards a community repository of large-text resources. In *Proceedings of the 21st Nordic Conference on Computational Linguistics*, pages 271–276, Gothenburg, Sweden. Association for Computational Linguistics.
- Vitali Gubarev. 1948. Павлик Морозов. Книгам лыкшы мары изд-во.
- Mika Hämäläinen. 2019. Uralicnlp: An nlp library for uralic languages. *Journal of open source software*.
- Mika Hämäläinen. 2021. Endangered languages are not low-resourced! *Multilingual Facilitation*.
- Mika Hämäläinen and Khalid Alnajjar. 2019. Let’s face it. finnish poetry generation with aesthetics and framing. In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 290–300.
- Mika Hämäläinen, Khalid Alnajjar, and Thierry Poibeau. 2022. Video games as a corpus: Sentiment analysis using fallout new vegas dialog. In *Proceedings of the 17th International Conference on the Foundations of Digital Games*, pages 1–4.
- Mika Hämäläinen, Niko Partanen, Jack Rueter, and Khalid Alnajjar. 2021. Neural morphology dataset and models for multiple languages, from the large to the endangered. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaL-iDa)*, pages 166–177.
- Kaisa Häkkinen. 1997. Suomi eilen ja tänään. Independently published, Helsinki.
- Joo-Kyung Kim, Young-Bum Kim, Ruhi Sarikaya, and Eric Fosler-Lussier. 2017. Cross-lingual transfer learning for POS tagging without cross-lingual resources. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2832–2838, Copenhagen, Denmark. Association for Computational Linguistics.
- Language Bank of Finland. 2022. Word embeddings trained with word2vec from the finnish text collection. In <http://urn.fi/urn:nbn:fi:lb-2022041406>.

- Shuhua Monica Liu and Jiun-Hung Chen. 2015. A multi-label classification based approach for sentiment classification. *Expert Systems with Applications*, 42(3):1083–1093.
- Julian McAuley and Jure Leskovec. 2013. Hidden factors and hidden topics: understanding rating dimensions with review text. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 165–172.
- Christopher Moseley, editor. 2010. *Atlas of the World's Languages in Danger*, 3rd edition. UNESCO Publishing. Online version: <http://www.unesco.org/languages-atlas/>.
- Sjur Moshagen, Jack Rueter, Tommi Pirinen, Trond Trosterud, and Francis M Tyers. 2014. Open-source infrastructures for collaborative work on under-resourced languages. *Collaboration and Computing for Under-Resourced Languages in the Linked Open Data Era*, pages 71–77.
- Emily Öhman. 2021. The validity of lexicon-based sentiment analysis in interdisciplinary research. In *Proceedings of the Workshop on Natural Language Processing for Digital Humanities*, pages 7–12.
- Niko Partanen, Rogier Blokland, KyungTae Lim, Thierry Poibeau, and Michael Rießler. 2018. The first Komi-Zyrian universal dependencies treebanks. In *Second Workshop on Universal Dependencies (UDW 2018), November 2018, Brussels, Belgium*, pages 126–132.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D Manning. 2020. Stanza: A python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108.
- Nils Reimers and Iryna Gurevych. 2019. Sentencebert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Jack Rueter, Mika Hämmäläinen, and Niko Partanen. 2020. Open-source morphology for endangered mordvinic languages. In *Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS)*. The Association for Computational Linguistics.
- Jack Rueter, Niko Partanen, Mika Hämmäläinen, and Trond Trosterud. 2021. Overview of open-source morphology development for the komi-zyrian language: Past and future. In *Proceedings of the Seventh International Workshop on Computational Linguistics of Uralic Languages*. The Association for Computational Linguistics.
- Jack Michael Rueter and Francis M Tyers. 2018. Towards an open-source universal-dependency treebank for erzya. In *International Workshop for Computational Linguistics of Uralic Languages*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958.
- Moshe Stekel, Amos Azaria, and Shai Gordin. 2021. Word sense induction with attentive context clustering. In *Proceedings of the Workshop on Natural Language Processing for Digital Humanities*, pages 144–151.
- Michael E. Tipping and Christopher M. Bishop. 1999. Mixtures of Probabilistic Principal Component Analyzers. *Neural Computation*, 11(2):443–482.
- Grigorios Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. 2010. Random k-labelsets for multilabel classification. *IEEE transactions on knowledge and data engineering*, 23(7):1079–1089.
- Ruochen Xu, Yiming Yang, Naoki Otani, and Yuxin Wu. 2018. Unsupervised cross-lingual transfer of word embedding spaces. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2465–2474.
- Selim F Yilmaz, E Batuhan Kaynak, Aykut Koç, Hamdi Dibeklioğlu, and Suleyman Serdar Kozat. 2021. Multi-label sentiment analysis on 100 languages with dynamic weighting for label imbalance. *IEEE Transactions on Neural Networks and Learning Systems*.
- Tatu Ylönen. 2022. Wiktextextract: Wiktionary as machine-readable structured data. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 1317–1325.