

# When Word Embeddings Become Endangered

Khalid Alnajjar<sup>[0000–0002–7986–2994]</sup>

Department of Digital Humanities,  
Faculty of Arts, University of Helsinki, Finland  
`khalid.alnajjar@helsinki.fi`

**Abstract.** Big languages such as English and Finnish have many natural language processing (NLP) resources and models, but this is not the case for low-resourced and endangered languages as such resources are so scarce despite the great advantages they would provide for the language communities. The most common types of resources available for low-resourced and endangered languages are translation dictionaries and universal dependencies. In this paper, we present a method for constructing word embeddings for endangered languages using existing word embeddings of different resource-rich languages and the translation dictionaries of resource-poor languages. Thereafter, the embeddings are fine-tuned using the sentences in the universal dependencies and aligned to match the semantic spaces of the big languages; resulting in cross-lingual embeddings. The endangered languages we work with here are Erzya, Moksha, Komi-Zyrian and Skolt Sami. Furthermore, we build a universal sentiment analysis model for all the languages that are part of this study, whether endangered or not, by utilizing cross-lingual word embeddings. The evaluation conducted shows that our word embeddings for endangered languages are well-aligned with the resource-rich languages, and they are suitable for training task-specific models as demonstrated by our sentiment analysis model which achieved a high accuracy. All our cross-lingual word embeddings and the sentiment analysis model have been released openly via an easy-to-use Python library.

**Keywords:** Cross-lingual Word Embeddings · Endangered Languages · Sentiment Analysis.

## 1 Introduction

The interest in building natural language processing (NLP) solutions for low-resourced languages is constantly increasing [1], not only because of the challenges associated with dealing with scarce resources but also because NLP solutions facilitate documenting and analysing languages. Examples of such solutions are applying optical character recognition to scan books [45], normalizing historical variation [7], using speech recognition [19] and more. However, most of the existing research is conducted in a simulated setting [15,24,22] where a reduced portion of the resource-rich language is used to represent a low-resourced language. Other approaches consider Wikipedias of languages having a small number of articles (i.e., < 500,000) such as Latin, Hindi, Thai and Swahili [9,48].

In this paper, we are dealing with languages that are classified as endangered based on UNESCO Atlas<sup>1</sup>. These languages are Erzya<sup>2</sup> (myv), Moksha (mdf), Komi-Zyrian (kpv) and Skolt Sami (sms). The most common methodology for documenting endangered languages is constructing translation dictionaries, whether digitizing physical dictionaries or reaching to native speakers. Universal dependencies (UD) written by dedicated researchers studying such endangered languages might also be available, and, in a fortunate scenario, they would include translations to a language with more speakers. The bigger languages that endangered languages are translated to are very inconsistent and vary depending on the language family, geographically close languages and the languages spoken by the documenter.

English, without a doubt, is currently the most resourced language in the field of NLP. However, English translations are not frequently found for endangered and low-resourced languages. To overcome this and make using existing English resources possible, we leverage the recent advances in the field of NLP for aligning word embeddings of big languages such as Finnish and Russian with English word embeddings.

The contributions in this paper are:

- Proposing a method for constructing word embeddings for low-resourced and endangered languages, which are also aligned with word embeddings of big languages.
- Building a universal sentiment analysis model that achieves high accuracy in both endangered and resource-rich languages covered in this work.
- Releasing an open-source and easy-to-use Python library with all the word embeddings and the sentiment analyzer model to support the community and researchers<sup>3</sup>.

This paper is structured as follows. Section 2 contains a brief description of the related work on building cross-lingual low-resource word embeddings. Thereafter, we describe the linguistic resources used in this work, including the translation dictionaries, universal dependencies and existing word embeddings of resource-rich. The proposed method for constructing cross-lingual word embeddings for endangered languages is elaborated then, followed by the description of the sentiment analysis model. We then present the results and evaluation for word embeddings and sentiment analysis model. Lastly, we discuss and highlight our remarks in the conclusions.

## 2 Related work

The largest scale model for capturing the computational semantics of endangered Uralic languages, Erzya, Moksha, Komi-Zyrian and Skolt Sami, is, perhaps, SemUr [16]. The database consists of words that are connected to each other based

<sup>1</sup> <http://www.unesco.org/languages-atlas/index.php>

<sup>2</sup> See [37] for an insightful description of the situation of the language.

<sup>3</sup> <https://github.com/mokha/semantics>

on their syntactic co-occurrences in a large internet corpus for Finnish. The extracted relations have been automatically translated by using Jack Rueter’s XML dictionaries. In human evaluation, the quality was surprisingly acceptable given that the method was based on word-level translations. This gives hope in using these high-quality dictionaries in building computational semantic models.

Apart from SemUr, there has not been any other attempts in automatically modelling semantics for endangered Uralic languages. Some recent work, however presents interesting work on higher-resourced languages using word embeddings [2,11]. In general, word embeddings based methods such as word2vec [28] and fastText [6] are optimal for the task of applying high-resource language data to endangered languages as they work on word-level.

Several recent approaches such as GPT-2 [34], ELMo [33] and BERT [10] aim to capture richer semantic representations from text. However, they are very data intensive and their representation is no longer on the level of individual words. This makes it more difficult to use them for endangered languages.

Recently, neural networks have been used heavily in the field of NLP due to their great capabilities in learning a generalization, which resulted in high accuracies. However, neural networks demand a large amount of data, which usually is not available for low-resource languages. Despite this, researchers have employed neural networks in a low-resource setting by producing synthetic data. For instance, Hämäläinen and Rueter have built a neural network to detect cognates for between two endangered languages [18], Skolt Sami and North Sami. Their approach reached to a better accuracy when they combined data, synthetically produced by a statistical model, with real data.

### 3 Linguistic resources

Here, we describe the linguistic resources used throughout the research presented in this paper. We will focus on resources related to the endangered languages (i.e., Erzya, Moksha, Komi-Zyrian and Skolt Sami), while still providing a brief introduction to resource-rich resources. The resources for endangered languages that we cover here are: 1) translation dictionaries, 2) universal dependencies and 3) finite-state transducers. This list by no means is inclusive of all available and useful resources for endangered languages, as additional resources might exist such as the work of Jack Rueter on Online dictionaries [41] and making them usable even through click-in-text interfaces [38]. In terms of resource-languages, we describe their word-embeddings.

#### 3.1 Translation dictionaries

Low-resource and endangered languages commonly have translation dictionaries to a bigger language. For our case, such dictionaries are multilingual and are provided in an Extensible Markup Language (XML) format. Fortunately, the target languages of the translations are mostly consistent in all the dictionaries

(which is not the typical case), but each dictionary contains different portions of translations.

Table 1 shows a statistical summary of the translations existing in the dictionaries. The source language represents the endangered languages and the target language indicates the resource-rich language. A meaning group in the dictionaries may contain multiple translations that can be used interchangeably as they share the same meaning. The analysis shows that for Erzya (myv) and Skolt Sami (sms), Finnish (fin) translations are the most common ones, whereas Russian (rus) and English (eng) translations are the most frequent ones for Komi-Zyrian (kpv) and Moksha (mdf).

**Table 1.** An overview of translation in the XML dictionaries of the low-resourced languages. Language codes are given in ISO 639-3.

Source language	Target language	Meaning groups	Translations	Total
myv	fin	<b>8388</b>	<b>14344 (59.89%)</b>	23950
	rus	5631	7608 (31.77%)	
	eng	1917	1998 (8.34%)	
kpv	fin	1352	2046 (14.89%)	13744
	rus	<b>15492</b>	<b>10585 (77.01%)</b>	
	eng	1078	1113 (8.10%)	
sms	fin	<b>20503</b>	<b>27522 (68.02%)</b>	40461
	rus	4872	6026 (14.89%)	
	eng	5824	6913 (17.09%)	
mdf	rus	37	37 (0.27%)	13626
	eng	<b>6587</b>	<b>13589 (99.73%)</b>	

Entries in the dictionaries are in the lemma form, and, typically, their part-of-speech tags are provided. Further metadata information might exist, such as stems and example usages of the word in the source language. We use the Giella [29] dictionaries that have been mainly authored by Jack Rueter through UralicNLP [21]. While Moksha has Finnish translations, the Moksha dictionary in UralicNLP did not contain any of these translations because the data was missing from the repository.

### 3.2 Universal dependencies

Universal dependencies (UD) [47] is a standard framework for annotating the grammar (parts of speech, morphological features, and syntactic dependencies of sentences. Additionally, UD allows annotators to supply their own comments. In the UD we are dealing with, translation sentences might appear in the comments. The UD of the endangered languages can be obtained directly from Universal Dependencies’ website<sup>4</sup>. At the time of writing, 1,690, 167, 104 and 435 sen-

<sup>4</sup> <https://universaldependencies.org/>

tences were in Erzya’s [44], Moksha’s [42], Skolt Sami’s [30] and Komi-Zyrian’s UDs<sup>5</sup> [32], respectively. These numbers highlight the insufficient amount of data present for training machine learning or NLP models for endangered languages. We have used the UralicNLP [21], a Python library, to read the universal dependencies.

### 3.3 Finite-state transducers

The common automatic tools found for endangered languages are finite-state transducers (FSTs), as they are rule-based which allows language experts to define how the finite-state machine should behave depending on the language. As a result, FSTs make it possible to lemmatize words and produce mini- and full-paradigms. In this work, we use Jack Rueter’s FSTs for Skolt Sami [39], Erzya and Moksha [40], and Komi-Zyrian [36]. The FSTs are supplied as part of the UralicNLP [21] Python library.

### 3.4 Word embeddings of resource-rich languages

Word embeddings are a vector representation of words, which are built based on the surrounding context of the word. Semantic similarity between words captured in the word embeddings can be measured using cosine similarity, which can then be utilized to cluster meanings in text [17]. Common usages for word embeddings is to acquire semantically similar words to an input word. For example, the most 5 similar words to “king” are “queen”, “monarch”, “prince”, “sultan”, and “ruler”. The vector nature of these words makes it possible to perform vector operations such as addition, multiplication and subtractions. With such operations, analogies could be predicted such as “king” - “man” + “woman” = “queen”. Simply, this asks what is the equivalent of a king that is not a man but rather a woman in the semantic space, the answer is a queen.

When building word embeddings, there are many preprocessing configurations and hyperparameters that influence the performance of the models, such as lemmatization, part-of-speech tagging, window size, the dimension size of the embeddings, minimum and maximum thresholds for word frequencies and so on. There is no fixed nor optimal configuration that is apt for all applications.

In the translation dictionaries, words and their translations are provided in their lemma form. Due to this reason, the vocabulary in any word embeddings we will be using has to be lemmatized. Ideally, all the hyperparameters and configurations for word embeddings should be the same to capture similar features and semantics, which would yield better results across models once they are aligned. For the scope of this research, we use the most similar models we could get our hands on.

We utilize the Russian and English [12], and Finish [25] word embeddings. The Russian embeddings are trained on a news corpus, while the English is based on Wikipedia and Gigaword 5th Edition corpora [31]. The Finnish word

<sup>5</sup> There is a UD for Komi-Permyak [43] which is close to Komi-Zyrian.


embeddings are trained on Common Crawls. The dimension size of the English and Russian embeddings is 300 but 200 is the size of the Finnish one. The window size is 5 for all embeddings but Finnish, which is 2. These differences, in addition to other reasons, end up affecting the quality of the models we will build of endangered languages. We discuss them more in the Discussion section.

## 4 Cross-lingual word embeddings for endangered languages

Cross-lingual word embeddings are word embeddings where vectors across multiple languages are aligned. For instance, the vector for “dog” in the English embeddings points roughly to the same direction for the same word in other languages (i.e., “koira” and “собака” for Finnish and Russian, respectively). Example applications for employing cross-lingual word embeddings are: headline generation [5], loan word identification [27] and cognate identification [26].

Before we build and align the word embeddings, we apply a dimensionality reduction using the method proposed in [35] to the three pre-trained models (i.e., English, Russian and Finnish). We set the target dimension to 100. This is to ensure that the vectors in all the embeddings share the same size. Subsequently, we process the vocabulary of the Finnish by removing all occurrences of the hashtag symbol “#”, which is there to mark compounds. Regarding the Russian word embeddings, the vocabulary contained part-of-speech information and, hence, each lemma might be present multiple times. To address this, we discard the part-of-speech information and use all vectors matching the target lemma.

To align the main three word embedding models, we employ the state-of-the-art supervised multilingual word embeddings alignment technique introduced in MUSE [8]. Figure 1 illustrates transforming the word embeddings of the source language  $X$  with the target language  $Y$  so that words in both languages are aligned together. In this example the source language is English and the target language is Italian. What supervised means in this context is that the alignment process relies on a bilingual dictionary that guides the transformation process. In our work, we set the target language to English and align both Russian and Finnish models with it using the bilingual dictionaries released as part of MUSE. The models are refined over 20 iterations.



**Fig. 1.** A visualization of the transformation process of aligning word embeddings in  $X$  in accordance to the ones in  $Y$ , taken from [8].

Following the alignment of the resource-rich models, we construct the word embeddings for the endangered languages: Erzya, Moksha, Komi-Zyrian and Skolt Sami. In doing so, we iterate over all the lexemes in the dictionary of a given endangered language. In the case where a lexeme had translations to any of the three resource-rich languages and the translation existed in the word embeddings of the corresponding language, a vector for the lexeme is constructed as the centroid –an average vector– of all translation vectors.

Once the word embeddings for the endangered languages have been constructed, we fine-tune them using the sentences in their universal dependencies. Lastly, we realign each word embeddings model with the resource-rich language having most translations to. In other words, Erzya and Skolt Sami are aligned with Finnish but Komi-Zyrian and Moksha are aligned with Russian and English, respectively. The models are aligned over 5 refinement steps.

## 5 Sentiment analysis

In this section, we describe an experiment with the newly produced word embeddings. We apply them in the task of sentiment analysis. We hand pick all positive and negative sentences from the Erzya treebank [44] based on the translations provided in the treebank in English and Finnish. This constitutes our Erzya test corpus that contains 23 negative sentences and 22 positive sentences, giving us a total of 45 sentences.

We use the Stanford Sentiment Treebank for English [46] to train our sentiment analyzer model. As the Erzya test data is binary – negative and positive sentences – we treat the sentiment information in the treebank as binary as well, ignoring any neutral examples. It is important to note that we do not use any examples written in Erzya during the training, only sentences in English.

We train a neural model that takes in a sentence in English as a source and a sentiment label (positive or negative) as a target. We train the neural model with the aligned embeddings by substituting the words in the input sentences with their vectors. As our models are lemmatized, we need to ensure that all words are lemmatized in the in input as well. We use spaCy [20] for this lemmatization step. The architecture and training of the neural model is inspired by the work presented in [23], where bi-grams are added to the input sentences during the training phase and the neural network is a linear classifier.

Table 3 shows some examples of the input in Erzya, its translation in English and the correctly predicted label. For Erzya, we use the lemmas from the treebank, and get their closest English vectors through the aligned word embeddings. This way, the model treats the Erzya sentences as though they were English and it can predict the sentiment in the language it did not see during the training. The resulting model was trained for 30 epochs and it reached to 53.3% accuracy for Erzya and 75.5% accuracy for English in the treebank sentences and an accuracy of 83.5% in English in the Stanford Sentiment Treebank dataset. We have obtained an accuracy boost for Erzya predictions, reaching 57.8%, when we also considered vectors of other resource-rich languages with the aid of the

**Table 2.** The top 3 semantically similar words to the English vector, in all languages. The score after the colon is the semantic similarity score (the higher, the more similar).

eng (input)	fin	rus	myv	mdf	sms	kpv
dog	koira: 0.7100 kissa: 0.6618 namipalo: 0.6263	поймать: 0.4354 убивать: 0.4310 .родственник: 0.4271	пинелевкс: 0.8172 псарня: 0.6691 киска: 0.6340	пине: 0.9977 кутерь: 0.8220 ката: 0.7547	piännai: 0.7197 piänngaž: 0.7078 kaazzâž: 0.6521	гиджгысь: 0.5691 барсук: 0.5271 вежель: 0.5238
cat	rotta: 0.6631 Fretti: 0.6484 kissa: 0.6461	щенок: 0.5246 бася: 0.4885 детеныша: 0.4794	обизьган: 0.7474 лыркай: 0.7018 пинелевкс: 0.6993	ката: 0.9990 зака: 0.9990 пине: 0.7478	kaazzâž: 0.6800 žee'vet: 0.6672 kue'ttžeevai: 0.6665	черепаха: 0.6121 питон: 0.5996 обезьяна: 0.5937
king	Ahasveros: 0.5551 kuningas: 0.5522 kuninkas: 0.5243	наследник: 0.4433 гордо: 0.4188 исходемик: 0.4122	инеазормастор: 0.6751 озадкс: 0.6643 инеазорнь: 0.6315	кароль: 0.9971 тюштя: 0.8601 оцязор: 0.7768	daam: 0.5301 koongõs: 0.5214 lää'ssnõmm: 0.5035	королева: 0.4869 принц: 0.4736 герцог: 0.4648
queen	kuningatar: 0.5902 prinsessa: 0.5867 kruununprinsessa: 0.5686	паркер-боулз: 0.5578 энистон: 0.5314 чад: 0.5063	инеазорава: 0.9954 инеазорнь: 0.5428 венчакай: 0.5360	оцязорова: 0.9972 kevârij: 0.6191 лемдяз: 0.6191	koongõskaav: 0.7180 prince'ss: 0.5865 kä'll-lä'dd: 0.5457	королева: 0.7227 принцесса: 0.6614 принц: 0.4903
car	auto: 0.7728 kottero: 0.6843 katumaasturi: 0.6627	машина: 0.6621 бмв: 0.6234 bmw: 0.6170	автомобиль: 0.7716 автомашина: 0.7716 уаз: 0.7438	машина: 0.8568 автокрандаз: 0.6957 ардомбяль: 0.6377	autt: 0.6826 mõõnnämneävv: 0.6572 luâdastvuejjamautt: 0.6438	мотик: 0.6299 водитель: 0.5915 автобуса: 0.5914
man	spolle: 0.5062 pedofiliä: 0.5029 puukottaja: 0.4986	пожизненно: 0.4450 остин::пауэрс: 0.4377 кривенко: 0.4291	мююкт: 0.4911 нарт: 0.4911 гурямка: 0.4869	алья: 0.7974 ава: 0.5362 сакал: 0.5212	so'rmmejei: 0.4548 upsee'r: 0.4522 nuõtrooumaž: 0.4468	айулов: 0.4970 катаржик: 0.4538 допроситны: 0.4063
woman	romaninainen: 0.5813 somalinainen: 0.5713 maahanmuuttajanainen: 0.5436	юлия::печерская: 0.5349 столбова&mdash: 0.5157 воспитатель: 0.5079	аваломань: 0.5539 авасыме: 0.5428 пекиязь: 0.4938	ава: 0.9988 авакань: 0.6255 ни: 0.5704	neežzan: 0.6134 ääumai: 0.4610 äälm: 0.4610	айулов: 0.4585 мам: 0.4035 колготки: 0.3830
France	Ranska: 0.6330 Belgia: 0.6097 Iso-Britannia: 0.5757	франция: 0.5325 деша: 0.4916 арно: 0.4801	французонь: 0.4922 француз: 0.4922 австриец: 0.4586	Кранцмастор: 0.9964 кранц: 0.7155 кранцава: 0.7155	Franskjännam: 0.6357 Jõnn-Britann: 0.5778 Itaal: 0.5331	забастовка: 0.4077 кõрень: 0.3972 японец: 0.3698
Finland	Tanska: 0.5735 Norja: 0.5732 Viro: 0.5612	тудегепев: 0.4599 инсбрук: 0.4462 либерец: 0.4398	Финляндия: 0.4457 Суоми: 0.4457 Россия: 0.4384	шведонь: 0.6399 шведава: 0.6399 швед: 0.6399	Lä'dd: 0.6780 Lää'ddjännam: 0.6780 Taarr: 0.6424	ненеч: 0.4165 вужкыв: 0.3451 подувкыв: 0.3451
see	Muuttu: 0.4886 tämä: 0.4860 ainavain: 0.4824	видеть: 0.5243 тренд: 0.5057 тенденция: 0.4935	покш: 0.5228 те: 0.5200 но: 0.4920	няемс: 0.9982 няема: 0.5860 ила-крда: 0.5248	õinn: 0.5315 o'ddjõõttâd: 0.4936 tiett-aa: 0.4881	дзик: 0.4642 эся: 0.4590 шензьõдлыны: 0.4540
want	haluta: 0.5709 siksi: 0.5219 molempi: 0.5146	жить: 0.5960 хотеть: 0.5225 актриса::юлия::михалков: 0.5222	мирямс: 0.5852 секс: 0.5654 одемс: 0.5654	пиштемс: 0.6032 мезевок: 0.5864 мезе-мезе: 0.5861	soovšed: 0.5562 haa'leed: 0.5319 jee'res: 0.5257	гажавны: 0.8506 желайтны: 0.6250 ньõтчыдысь: 0.5561
day	lomaamu: 0.5336 reissupäivä: 0.5230 lähtõa: 0.5088	утра: 0.4900 7days.ru: 0.4613 выплакать: 0.4506	поздаямс: 0.4964 покшнэ: 0.4948 час: 0.4730	цяс: 0.5621 ой: 0.5621 шиньгучка: 0.5351	minut: 0.4695 jâdda: 0.4684 kõskkpei'vv: 0.4673	мõдасув: 0.4533 вежонпом: 0.4377 салют: 0.3980



**Table 3.** Example sentences in Erzya and their translations in English, along with the predicted sentiment by our method for each sentence.

Erzya	English	Sentiment
Зярошкаль цёрыньгъ кенярксозо!	You can imagine the boy’s delight!	Positive
Чизэ лембе.	It is a warm day.	
Сехге паро шка.	The best time of all.	
Цёрыньгъ аламодо визькс теевсь.	The boy felt a little ashamed.	Negative
Баагинень ёмавомась — пек берянь тежксесь.	Losing a bell was a really bad sign.	
Весе те — апаро вийтнень таңдавтнемс.	This is all meant to scare away the evil spirits.	

translation dictionary (Finnish in this case, as Erzya has many translations to Finnish).

The resulting accuracy is respectable given that the test data is fundamentally different from the training data. First of all, the testing and training are in different languages. Second of all, they represent very different genres: the training data is based on movie reviews, whereas the testing data has sentences from novels.

## 6 Discussion and Conclusions

The work conducted in this paper has been a first step for using machine learning in modelling the semantics of some of the endangered Uralic languages. It is evident that these aligning based approaches embraced before in the literature cannot get us too far in truly representing the semantics do to socio-cultural mismatches in concepts. For instance, we saw that *Finland*, which is a very important concept for a Finnish model was completely misaligned with geographically close countries such as *Denmark*, *Norway* and *Estonia*. Alignment can only get us so far and using models trained on larger languages has its inherent problems when applied to completely new domains in a completely different language.

Even the starting quality for the pretrained embeddings was low. The Russian model was unacceptably bad and the Finnish model has too many words that are not lemmatized at all, or are lemmatized to a wrong lemma. When the quality of the models available for a high-resourced languages is substandard, one cannot expect any sophisticated machine learning method to come to the rescue. Unfortunately in our field, too little attention is paid to the quality of resources and more attention is paid into single values representing overall accuracies and overall performance.

As there is no shortcut to happiness, we should look into the data available in the endangered languages themselves. For instance, FU-Lab has a plethora of resources for Komi languages [14,13] that are just waiting for lemmatization. Once lemmatized, these resources could be used to build word embeddings directly in that language. Of course, this requires collaboration between many parties and willingness to make data openly available. While this might not be an issue with FU-Lab, it might be with some other instances holding onto their immaterial rights too tight.

At the current, stage our dictionary editing system, Ve'rdd [4,3], contains words for multiple endangered languages and their translations in a graph structure. This data could be extended by predicting new relations into the graph with semantic models such as word embeddings. This could help at least in resolving meaning groups and polysemy of the lexical entries. However, the word embeddings available for the endangered languages in question has not yet reached to a stage mature enough for their incorporation as a part of the lexicon.

## 7 Acknowledgement

I would like to dedicate the acknowledgement section to Jack Rueter, for all his work on endangered languages and his brilliant ideas on improving the current technological state of endangered languages. Jack's enthusiasm and dedication to endangered languages is clearly shown in all the various dictionaries and FSTs built and maintained by him. He supervised my work on building the dictionary editing system, Ve'rdd [4,3]. He was always available for discussing and supporting my work, without him Ve'rdd would not be in the great level it is at at the moment. He truly is a pioneer in the field, and the entire community appreciates all of his work.

## References

1. Proceedings of the 3rd Workshop on the Use of Computational Methods in the Study of Endangered Languages Volume 1 (Papers). Association for Computational Linguistics, Honolulu (Feb 2019), <https://www.aclweb.org/anthology/W19-6000>
2. Adams, O., Makarucha, A., Neubig, G., Bird, S., Cohn, T.: Cross-lingual word embeddings for low-resource language modeling. In: Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers. pp. 937–947. Association for Computational Linguistics, Valencia, Spain (Apr 2017), <https://www.aclweb.org/anthology/E17-1088>
3. Alnajjar, K., Hämäläinen, M., Rueter, J.: On editing dictionaries for uralic languages in an online environment. In: Proceedings of the Sixth International Workshop on Computational Linguistics of Uralic Languages. p. 26–30. The Association for Computational Linguistics (2020)
4. Alnajjar, K., Hämäläinen, M., Rueter, J., Partanen, N.: Ve'rdd. narrowing the gap between paper dictionaries, low-resource nlp and community involvement. In: Proceedings of the 28th International Conference on Computational Linguistics: System Demonstrations. International Committee on Computational Linguistics (2020)
5. Alnajjar, K., Leppänen, L., Toivonen, H.: No time like the present: Methods for generating colourful and factual multilingual news headlines. In: Proceedings of the 10th International Conference on Computational Creativity. pp. 258–265. Association for Computational Creativity (Jun 2019)
6. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics* **5**, 135–146 (2017)

7. Bollmann, M.: A large-scale comparison of historical text normalization systems. arXiv preprint arXiv:1904.02036 (2019)
8. Conneau, A., Lample, G., Ranzato, M., Denoyer, L., Jégou, H.: Word translation without parallel data. arXiv preprint arXiv:1710.04087 (2017)
9. Das, A., Ganguly, D., Garain, U.: Named entity recognition with word embeddings and wikipedia categories for a low-resource language. *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)* **16**(3), 1–19 (2017)
10. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: 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: Human Language Technologies, Volume 1 (Long and Short Papers)*. pp. 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019). <https://doi.org/10.18653/v1/N19-1423>, <https://www.aclweb.org/anthology/N19-1423>
11. Duong, L., Kanayama, H., Ma, T., Bird, S., Cohn, T.: Learning crosslingual word embeddings without bilingual corpora. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)*. Association for Computational Linguistics, Texas, USA (November 2016)
12. Fares, M., Kutuzov, A., Oepen, S., Velldal, E.: Word vectors, reuse, and replicability: Towards a community repository of large-text resources. In: *Proceedings of the 21st Nordic Conference on Computational Linguistics, NoDaLiDa, 22-24 May 2017, Gothenburg, Sweden. vol. 0*, pp. 271–276. Linköping University Electronic Press, Linköpings universitet (2017)
13. Федина, МС: Создание официально-делового подкорпуса национального корпуса коми языка. In: *УПРАВЛЕНИЕ СОЦИАЛЬНО-ЭКОНОМИЧЕСКИМ РАЗВИТИЕМ СУБЪЕКТА РОССИЙСКОЙ ФЕДЕРАЦИИ*, pp. 205–216 (2015)
14. Федина, Марина Серафимовна: Корпус Коми Языка Как База Для Научных Исследований. In: *II Международная научная конференция «Электронная письменность народов Российской Федерации: опыт, проблемы и перспективы»* проводится в рамках реализации Государственной программы «Сохранение и развитие государственных языков Республики Башкортостан и языков народов Республики Башкортостан» на 2019–2024 гг. Ответственный редактор: Ахмадеева АУ. p. 45 (2019)
15. Gu, J., Hassan, H., Devlin, J., Li, V.O.: Universal neural machine translation for extremely low resource languages. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. pp. 344–354 (2018)
16. Härmäläinen, M.: Extracting a semantic database with syntactic relations for finnish to boost resources for endangered uralic languages. *The Proceedings of Logic and Engineering of Natural Language Semantics 15 (LENLS15)* (2018)
17. Härmäläinen, M., Alnajjar, K.: Let’s FACE it: Finnish poetry generation with aesthetics and framing. In: *12th International Conference on Natural Language Generation*. pp. 290–300. The Association for Computational Linguistics (2019). <https://doi.org/10.18653/v1/w19-8637>
18. Härmäläinen, M., Rueter, J.: Finding Sami cognates with a character-based NMT approach. In: *Proceedings of the 3rd Workshop on the Use of Computational Methods in the Study of Endangered Languages Volume 1 (Papers)*. pp. 39–45. Association for Computational Linguistics, Honolulu (Feb 2019), <https://www.aclweb.org/anthology/W19-6006>

19. Hjortnaes, N., Partanen, N., Rießler, M., Tyers, F.M.: Towards a speech recognizer for komi, an endangered and low-resource uralic language. In: Proceedings of the Sixth International Workshop on Computational Linguistics of Uralic Languages. pp. 31–37 (2020)
20. Honnibal, M., Montani, I., Van Landeghem, S., Boyd, A.: spaCy: Industrial-strength Natural Language Processing in Python (2020). <https://doi.org/10.5281/zenodo.1212303>
21. Hämäläinen, M.: UralicNLP: An NLP library for Uralic languages. *Journal of Open Source Software* 4(37), 1345 (2019). <https://doi.org/10.21105/joss.01345>
22. Hämäläinen, M., Alnajjar, K.: A template based approach for training nmt for low-resource uralic languages - a pilot with finnish. In: ACAI 2019: Proceedings of the 2019 2nd International Conference on Algorithms, Computing and Artificial Intelligence. pp. 520–525. ACM, United States (Dec 2019). <https://doi.org/10.1145/3377713.3377801>
23. Joulin, A., Grave, E., Bojanowski, P., Mikolov, T.: Bag of tricks for efficient text classification. In: Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers. pp. 427–431. Association for Computational Linguistics, Valencia, Spain (Apr 2017), <https://www.aclweb.org/anthology/E17-2068>
24. Khayrallah, H., Thompson, B., Post, M., Koehn, P.: Simulated multiple reference training improves low-resource machine translation. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 82–89. Association for Computational Linguistics, Online (Nov 2020). <https://doi.org/10.18653/v1/2020.emnlp-main.7>, <https://www.aclweb.org/anthology/2020.emnlp-main.7>
25. Laippala, V., Ginter, F.: Syntactic n-gram collection from a large-scale corpus of internet finnish. In: Human Language Technologies-The Baltic Perspective: Proceedings of the Sixth International Conference Baltic HLT. vol. 268, p. 184 (2014)
26. Lefever, E., Labat, S., Singh, P.: Identifying cognates in English-Dutch and French-Dutch by means of orthographic information and cross-lingual word embeddings. In: Proceedings of the 12th Language Resources and Evaluation Conference. pp. 4096–4101. European Language Resources Association, Marseille, France (May 2020), <https://www.aclweb.org/anthology/2020.lrec-1.504>
27. Mi, C., Yang, Y., Wang, L., Zhou, X., Jiang, T.: Toward better loanword identification in Uyghur using cross-lingual word embeddings. In: Proceedings of the 27th International Conference on Computational Linguistics. pp. 3027–3037. Association for Computational Linguistics, Santa Fe, New Mexico, USA (Aug 2018), <https://www.aclweb.org/anthology/C18-1256>
28. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)
29. Moshagen, S., Rueter, J., Pirinen, T., Trosterud, T., Tyers, F.M.: Open-Source Infrastructures for Collaborative Work on Under-Resourced Languages (2014), the LREC 2014 Workshop “CCURL 2014 - Collaboration and Computing for Under-Resourced Languages in the Linked Open Data Era”
30. Nivre, J., Zeman, D., Rueter, J., Juutinen, M., Hämäläinen, M.: Ud\_skolt\_sami-giellagas 2.7 (Nov 2020)
31. Parker, R., Graff, D., Kong, J., Chen, K., Maeda, K.: English gigaword fifth edition, 2011. Linguistic Data Consortium, Philadelphia, PA, USA (2011)
32. Partanen, N., Blokland, R., Lim, K., Poibeau, T., Rießler, M.: The first Komi-Zyrian Universal Dependencies treebanks. In: Proceedings of the Second Workshop on Universal Dependencies (UDW

- 2018). pp. 126–132. Association for Computational Linguistics, Brussels, Belgium (Nov 2018). <https://doi.org/10.18653/v1/W18-6015>, <https://www.aclweb.org/anthology/W18-6015>
33. Peters, M.E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., Zettlemoyer, L.: Deep contextualized word representations. In: Proc. of NAACL (2018)
  34. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners. OpenAI blog **1**(8), 9 (2019)
  35. Raunak, V., Gupta, V., Metze, F.: Effective dimensionality reduction for word embeddings. In: Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019). pp. 235–243. Association for Computational Linguistics, Florence, Italy (Aug 2019). <https://doi.org/10.18653/v1/W19-4328>, <https://www.aclweb.org/anthology/W19-4328>
  36. Rueter, J.: Хельсинкиса университетын кыв туялысь Ижкарын перымса кы-  
въяс симпозиум вылын лыддьомтор. Permistika pp. 154–158 (2000)
  37. Rueter, J.: The erzya language. where is it spoken? *Études finno-ougriennes* (45) (2013)
  38. Rueter, J.: Giellatekno open-source click-in-text dictionaries for bringing closely related languages into contact. In: Proceedings of the Third Workshop on Computational Linguistics for Uralic Languages. pp. 8–9 (2017)
  39. Rueter, J., Härmäläinen, M.: FST morphology for the endangered Skolt Sami language. In: Proceedings of the 1st Joint Workshop on Spoken Language Technologies for Under-resourced languages (SLTU) and Collaboration and Computing for Under-Resourced Languages (CCURL). pp. 250–257. European Language Resources association, Marseille, France (May 2020), <https://www.aclweb.org/anthology/2020.sltu-1.35>
  40. Rueter, J., Härmäläinen, M., Partanen, N.: Open-source morphology for endangered mordvinic languages. In: Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS). pp. 94–100. Association for Computational Linguistics, Online (Nov 2020). <https://doi.org/10.18653/v1/2020.nlposs-1.13>, <https://www.aclweb.org/anthology/2020.nlposs-1.13>
  41. Rueter, J., Härmäläinen, M.: Synchronized mediawiki based analyzer dictionary development. In: Proceedings of the Third Workshop on Computational Linguistics for Uralic Languages. pp. 1–7 (2017)
  42. Rueter, J., Nivre, J., Zeman, D., Kabaeva, N., Levina, M.: Ud\_moksha-jr 2.7 (Nov 2020)
  43. Rueter, J., Partanen, N., Ponomareva, L.: On the questions in developing computational infrastructure for Komi-permyak. In: Proceedings of the Sixth International Workshop on Computational Linguistics of Uralic Languages. pp. 15–25. Association for Computational Linguistics, Wien, Austria (10–11 Jan 2020), <https://www.aclweb.org/anthology/2020.iwclul-1.3>
  44. Rueter, J., Tyers, F.: Towards an open-source universal-dependency treebank for Erzya. In: Proceedings of the Fourth International Workshop on Computational Linguistics of Uralic Languages. pp. 106–118. Association for Computational Linguistics, Helsinki, Finland (Jan 2018). <https://doi.org/10.18653/v1/W18-0210>, <https://www.aclweb.org/anthology/W18-0210>
  45. Silfverberg, M., Rueter, J.: Can morphological analyzers improve the quality of optical character recognition? In: Septentrio Conference Series. pp. 45–56. No. 2 (2015)
  46. Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C.D., Ng, A., Potts, C.: Recursive deep models for semantic compositionality over a sentiment treebank.

- In: Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, pp. 1631–1642. Association for Computational Linguistics, Seattle, Washington, USA (Oct 2013), <https://www.aclweb.org/anthology/D13-1170>
47. Zeman, D., Nivre, J., Abrams, M., Ackermann, E., Aepli, N., Aghaei, H., Agić, Ž., Ahmadi, A., Ahrenberg, et al.: Universal dependencies 2.7 (2020), <http://hdl.handle.net/11234/1-3424>, LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University
  48. Zhou, S., Rihwani, S., Wieting, J., Carbonell, J., Neubig, G.: Improving candidate generation for low-resource cross-lingual entity linking. *Transactions of the Association for Computational Linguistics* **8**, 109–124 (2020), <https://www.aclweb.org/anthology/2020.tacl-1.8>