

Inferring Translation Candidates for Multilingual Dictionary Generation with Multi-Way Neural Machine Translation

Mihael Arcan, Daniel Torregrosa, Sina Ahmadi and John P. McCrae

Insight Centre for Data Analytics
Data Science Institute
National University of Ireland Galway
`firstname.lastname@insight-centre.org`

Abstract. In the widely-connected digital world, multilingual lexical resources are one of the most important resources, for natural language processing applications, including information retrieval, question answering or knowledge management. These applications benefit from the multilingual knowledge as well as from the semantic relation between the words documented in these resources. Since multilingual dictionary creation and curation is a time-consuming task, we explored the use of multi-way neural machine translation trained on corpora of languages from the same family and trained additionally with a relatively small human-validated dictionary to infer new translation candidates. Our results showed not only that new dictionary entries can be identified and extracted from the translation model, but also that the expected precision and recall of the resulting dictionary can be adjusted by using different thresholds.

Keywords: Neural machine translation · Dictionary generation · Automatic inference

1 Introduction

The growing amount of semantically structured monolingual, as well as multilingual resources, such as dictionaries or knowledge graphs (KGs), offers an excellent opportunity to explore, link and to enrich them with missing multilingual knowledge. Since a manual translation of such resources is very time-consuming and expensive, this work focuses on the translation and evaluation of dictionary entries between French and Spanish, extracted from the Wiktionary dictionary.¹ Furthermore, we focused on a less-resourced scenario, where a sufficient amount of parallel data to train a neural machine translation (NMT) system is not available.

In this scenario, we targeted bilingual dictionary generation using no or a minimal set of parallel sentences of the targeted language pair, and trained a multi-way NMT system on well-resourced language pairs. Our parallel corpus consisted of parallel sentences between English and Spanish, French and Romanian, and

¹ <https://www.wiktionary.org/>, dump version 20190201

Italian and Portuguese, where the targeted translation direction, i.e. French-Spanish, was not explicitly linked to the parallel training corpus. To improve the translation quality of the Wiktionary dictionary entries, we continued training the existing multi-way NMT system with a limited set of parallel entries of the targeted language pair. For this goal, we used two different datasets: the French-Spanish dictionary from Apertium [6], a rule-based machine translation platform, (Figure 1) and a small parallel French-Spanish corpus.²




Fig. 1. Apertium RDF graph, with the dictionary used in this work highlighted.

2 Related work

In this section, we give an overview on methods to acquire multilingual lexicons using similarity measures and graph based approaches.

[24] generated a bilingual dictionary using the structure of the source dictionaries. They introduced the Inverse Consultation (IC) approach which measures the semantic distance between two words based on the number of their common words in the pivot language. Using this method, [18] created an English-Gujarati lexicon using Hindi as the pivot. Similarly, [26] used English as an intermediate language to create a Chinese-Japanese lexicon. The IC method was extended by taking more lexical and semantic information into account [10]. For instance, [3] used part-of-speech information and semantic classes to produce a Japanese-Malay dictionary with English as the pivot. [21] created a Japanese-Swedish dictionary by linking words based on the sense definitions, whereby [11] constructed a Japanese-Chinese dictionary using a pivot English lexicon and co-occurrences information for more accurate alignment.

The high dependency of the IC method on one language as a pivot has been shown to create limited translations with ambiguity and low recall [20, 17]. One way to remedy this is to use multiple pivot languages with additional resources. [15] generated a Korean-Japanese dictionary using English and Chinese pivot languages and an English thesaurus. [22] described the automatic generation of a

² The used datasets and the trained NMT models are available at <http://server1.nlp.insight-centre.org/tiad2019/>

multilingual resource, called PanDictionary. In this work, the authors used probabilistic inference over the translation graph. The construction of the dictionary consisted of extracting knowledge from existing dictionaries and combining the obtained knowledge into a single resource. [9] took advantage of the semantic structure of WordNet as the pivot language for creating a new lexicon for less-resourced languages. [14] used string distance to create bilingual lexicons based on transduction models of cognates, as languages belonging to a specific language family usually share many cognates.

In recent years, there has been an increasing usage of graph-based algorithms such as random walk and graph sampling techniques for multilingual dictionary generation [2, 27]. [16] proposed a system for generating translation candidates using a graph-based approach with a weighting scheme and a collocation-based model based on the parallel corpus Europarl. In contrast, [1] focused on finding cycles of translations in the graph. By finding cycles of translations in the graph of all lexical entries with translations treated as undirected edges, the proposed approach was able to infer translations with reasonable accuracy. [5] used supervised machine learning to predict high-quality candidate translation pairs. They train a Support Vector Machine for classifying valid or invalid translation candidates. For this, they used several features, e.g. frequency of source word in a dictionary or minimum/maximum path length. Furthermore, string similarity leveraging, i.e. Levenshtein distance, was also taken into consideration.

3 Methodology

Multi-way neural machine translation To perform experiments on NMT models with a minimal set of parallel data, i.e. for less-resourced languages, we trained a multi-source and multi-target NMT model [8] with well-resourced language pairs. In our work, we have chosen parallel corpora part in the Romance language family, i.e. Spanish, Italian, French, Portuguese, Romanian, as well as English. To train the multi-way NMT system, we restricted the language pairs to English-Spanish, French-Romanian and Italian-Portuguese. Within this setting, the NMT system learns to translate text between languages mentioned above, but not between the French-Spanish language pair, which is the target of this work.

Continuous training with a limited set of parallel data To allow the NMT model to align words in the embedding space between French and Spanish, we used the trained multi-way model³ and continued the training of the network based on a minimal set of the French-Spanish parallel data. Without this procedure, the default multi-way system could not generate a translation of the targeted French-Spanish language pair and would instead generate translations into a non-Spanish language. As an example, when translating the French sentence:

il est donc conclu que l'établissement d'un préjudice ne dépend pas de l'utilisation de 2008 comme année de départ.

³ Trained on English-Spanish, French-Romanian and Italian-Portuguese parallel data

the default multi-way model generates a sentence in Romanian:

prin urmare, se concluzionează că instituirea prejudiciului nu depinde de utilizarea din 2008 ca anul de plec are.

The system further trained on the minimal French-Spanish dataset properly generates the Spanish translation:

por consiguiente, se concluye que el establecimiento de un perjuicio no depende de la utilización del año 2008 como año de partida.

For the continuous training of the multi-way model, we experimented with different datasets:

- Continuous training of the default multi-way model with 2,000 French-Spanish parallel sentences (ii_x in Figure 2)
- Continuous training of the default multi-way model with around 14,000 French-Spanish entries extracted from the Apertium dataset (iii_x in Figure 2)
- First, continuous training of the multi-way model with 2,000 French-Spanish parallel sentences. Second, the sentence adjusted multi-way NMT system is further trained on the Apertium dataset (iv_x in Figure 2)

4 Experimental Setup

In this section, we give an overview on the datasets and the translation toolkit used in our experiment.

4.1 Neural Machine Translation Toolkit

For our experiments, we used **OpenNMT** [12], a generic deep learning framework mainly specialised in sequence-to-sequence models covering a variety of tasks such as machine translation, summarisation, speech processing and question answering. We used the following neural network training parameters: 2 hidden layers, 500 hidden LSTM (long short term memory) units, input feeding enabled, a batch size of 64, 0.3 dropout probability and a dynamic learning rate decay. We train the network for 13 epochs and report the results in Section 5.

4.2 Byte Pair Encoding

A common problem in training a neural network is the computational complexity, which causes the vocabulary to be limited to a specific threshold. Because of this, the neural network cannot learn expressions of rare words. Therefore, if the training method does not see a specific word or phrase multiple times during training, it will not learn the interpretation of the word. This challenge is even more evident in sequence-to-sequence models used for machine translation. To overcome this limitation, different methods were suggested, i.e. character based neural model [4, 13] or the usage of subword units, e.g. Byte Pair Encoding (BPE). The latter one was successfully adapted for word segmentation specifically for Neural Machine Translation [19]. BPE [7] is a form of data compression that

Monolingual	# Tokens	# Types	# Subwords	# Uniq. Subwords	# Lines
English	16,116,345	103,893	16,178,313	15,829	768,344
Spanish	17,760,791	131,384	18,212,147	18,937	768,344
French	19,147,907	119,195	19,641,812	17,491	768,344
Romanian	16,567,238	142,105	17,547,273	19,266	768,344
Italian	17,266,710	128,918	17,819,166	18,164	768,344
Portuguese	17,558,350	125,657	17,914,861	19,210	768,344

	Source		Target		
Multi-Way	# Subwords	# Uniq. Subwords	# Subwords	# Uniq. Subwords	# Lines
train	111,058,273	32,244	102,326,852	32,230	4,366,309
validation	330,736	26,152	306,747	26,096	12,000

Continuous training	Source		Target		
	# Subwords	# Uniq. Subwords	# Subwords	# Uniq. Subwords	# Lines
Apertium	70,486	25,022	70,428	24,468	28,150
ParSent	54,199	1,998	54,487	1,999	2,000

Table 1. Dataset statistics for the DGT corpus for the monolingual resources, as well as the combined multi-way dataset used to train the translation system. Additionally, statistics on the Apertium and parallel sentences (ParSent) used to continue training the default multi-way translation model.

iteratively replaces the most frequent pair of bytes in a sequence with a single, unused byte. Instead of merging frequent pairs of bytes as shown in the original algorithm, characters or character sequences are merged for the purposes of natural language generation. To achieve this, the symbol vocabulary is initialised with the character vocabulary, and each word is represented as a sequence of characters—plus a special end-of-word symbol, which allows restoring the original tokenisation after the generating the answer based on the given question. This process is repeated as many times as new symbols are created.

4.3 Datasets

To train the multi-way model described in Section 3, we used the DGT (*Directorate-General for Translation*) corpus [23], a publicly accessible resource provided by the European Commission to support multilingualism and the re-use of European Commission information available in 24 different European languages. The English, Spanish, French, Romanian, Italian and Portuguese languages were selected to train the multi-way NMT system, from which we extracted 768,344 translated sentences present in all six languages within the DGT corpus (Table 1).

Additionally, we extracted two different Spanish-French dictionaries: one containing all the canonical forms stored in the Spanish-French Apertium bilingual dictionary file,⁴ with a total of 22,229 entries; and another one by accessing the

⁴ Version `cfc5995c5369ddb158cd266fcb8d4b74ed8dbdd0`.

Wiktionary dictionary, with a total of 58,659 entries. While the former dictionary was used to refine the output of the NMT system, the later was used as a gold standard for evaluation.

4.4 Evaluation and metrics

In order to evaluate the system, we used the trained multi-way NMT models applying the default, 1-best translation settings and a 10-best with a beam width of 10 to translate the corresponding side of the Wiktionary dictionary.

The main evaluation metrics that are used are precision and recall, using the dictionary entries extracted from Wiktionary as a reference. As NMT does not only provide the translations, but also the confidence score of each one, we can perform a receiver operating characteristic analysis; by using different thresholds of the confidence score, we are able to obtain a smaller, more precise dictionary, or a larger, less precise one.

5 Results

In this section, we present the results of the evaluation of the dictionaries generated using four different approaches: (i) by dumping the Apertium French-Spanish dictionary, (ii) the multi-way NMT (English-Spanish, French-Romanian and Italian-Portuguese) model further trained on French-Spanish sentences, (iii) the multi-way NMT model further trained on the French-Spanish Apertium dictionary entries and (iv) the multi-way NMT model further trained first on the French-Spanish sentences and then on the French-Spanish Apertium dictionary entries. Further experiments without explicit parallel data are demonstrated in [25].

5.1 Quantitative evaluation

The results for the quantitative evaluation can be found in Figure 2. As expected, adjusting the threshold to be more strict leads to higher precision at the expense of having a lower recall; using a looser threshold leads to the opposite behaviour.

The low recall can be explained by the nature of our reference dictionary; Wiktionary includes many different idioms, proper nouns, colloquial or dialectal variants. For example, in the English dictionary, we found entries such as *Adyghe Autonomous Oblast* (classified as a proper noun), *birds of the feather flock together* (classified as a phrase) or *black over Bill's mother's* (classified as an adjective). This affects the precision of the translations, as the NMT system will have a hard time translating proper nouns, terminological expressions or idioms, due to their infrequent appearance in the parallel corpora used to train the translation models.

While the performance of the model tuned on both sentences and dictionary entries (iv) is slightly higher than the one of the model tuned only on dictionary entries (iii), the second approach might be much more interesting, as it requires




Fig. 2. Results for the receiver operating characteristic curve using all possible score thresholds for the systems tuned on sentences (ii_x), on the generated dictionary (iii_x) and both (iv_x) with 1-best and default beam of 5 ($x=1$) and 10-best and beam of 10 ($x=10$), and the dictionary extracted from Apertium (i). While all the possible thresholds are shown as a line, only a subset of them are shown as points to ease the visualisation.

no parallel corpora. The performance of the system trained on sentences only (ii) is lower than the other two models, since the amount of unique words in the corpus is smaller than the one present in the Apertium dictionary ($\sim 2,000$ vs $\sim 25,000$). Surprisingly, the precision of the Apertium extracted dictionary (labelled as i in Figure 2) is quite low. In the following section, we perform a qualitative evaluation of the results in order to better explain this behaviour.

5.2 Qualitative evaluation

One of the major issues on automatically evaluating newly generated dictionaries is that there is no such thing as an ideal dictionary; all dictionaries used for

evaluation will be incomplete due to various reasons, such as new meanings arising, old meanings falling in disuse, or infrequent words that get excluded for the sake of clarity and brevity.

For example, in the Spanish-French dictionary, Apertium contains Spanish-French entries such as *bienal-biennale* or *inhumanamente-inhumanément* that might have not been included in Wiktionary due to the use of prefixes, or *polinizador-polinisateur* due to not being frequent enough. Apertium also contains the translation for many proper nouns and idioms that are not covered in Wiktionary (and vice-versa).

When analysing the translations provided by the multi-way NMT system, we observed that the system is not able to pick up some idioms and expressions. For example, the Spanish word *bajo*, part of the dictionary entry *bajo sajón-sous saxon*, has been translated into French as *sous* (en. *under*) instead of *bas* (en. *lower*). Similarly, *syndrome de la page blanche* in French got translated as *síndrome de la página blanca* (en. *writer's block*), that, while being an accurate word-by-word translation, does not have a specific meaning in Spanish. It could have been translated into the more appropriate Spanish idiom *bloqueo del escritor*, but the current idiom was not found in the training set, thus cannot be generated by the NMT model. This further highlights the difficulty of the task.

Additionally, the NMT system was initially not trained to translate between Spanish and French. After concluding the initial training, it was further trained on a few thousand sentences and/or dictionary entries (as described in Section 3). Nevertheless, the system sometimes still produced English translations, since it was initially trained with English paired to Spanish. As an example, when translating the Spanish term *tren de equipajes*, the system generated *train de luggage*, being *bagage* the correct French translation. Finally, as a BPE model was used, the NMT system produced incorrect translations for some infrequent words as well. For example, *cogote* in Spanish (en. *nape*) was translated into *coup*, because the first BPE (*co*) leads the NMT system to produce the most likely continuation.

6 Conclusion

In this work, we have shown that we can enrich an existing dictionary with multilingual knowledge by using multi-way neural machine translation trained on data that does not include parallel data between the targeted languages in combination with a small set of dictionary entries. The results can be further slightly improved by using a small parallel corpus.

One of the current limitations of the proposed system is that it is not possible to obtain the part-of-speech for the dictionary entries. This limitation can be overcome by obtaining the part-of-speech of each word by using monolingual taggers and annotating each word with the corresponding tag for training the NMT model. Therefore, we will focus on the usage of Universal Dependencies⁵ to uniformly annotate the corpora used to train the multi-way models.

⁵ <https://universaldependencies.org/>

Another possible improvement is the use of a lemmatizer to preprocess the corpus used to train the NMT system: dictionaries usually only list the canonical form of each word, therefore predicting any translation that is not in a canonical form will reduce the accuracy of the generated dictionary.

Finally, we can also improve the evaluation by devising criteria to exclude certain kind of entries such as proper nouns, colloquialisms and idioms from both the reference and the hypothesis, as long as they do not appear in the training set.

Acknowledgment

This publication has emanated from research supported in part by a research grant from Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289, co-funded by the European Regional Development Fund, and the European Union’s Horizon 2020 research and innovation programme under grant agreement No 731015, ELEXIS - European Lexical Infrastructure.

References

1. Morris Alper. Auto-generating bilingual dictionaries: Results of the TIAD-2017 shared task baseline algorithm. In *Proceedings of the LDK 2017 Workshops: 1st Workshop on the OntoLex Model (OntoLex-2017), Shared Task on Translation Inference Across Dictionaries & Challenges for Wordnets co-located with 1st Conference on Language, Data and Knowledge (LDK 2017), Galway, Ireland, June 18, 2017.*, pages 85–93, 2017.
2. Christophe Andrieu, Nando De Freitas, Arnaud Doucet, and Michael I Jordan. An introduction to MCMC for machine learning. *Machine learning*, 50(1-2):5–43, 2003.
3. Francis Bond and Kentaro Ogura. Combining linguistic resources to create a machine-tractable Japanese-Malay dictionary. *Language Resources and Evaluation*, 42(2):127–136, 2008.
4. Marta R. Costa-jussà and José A. R. Fonollosa. Character-based neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 357–361, Berlin, Germany, August 2016. Association for Computational Linguistics.
5. Kathrin Donandt, Christian Chiarcos, and Maxim Ionov. Using machine learning for translation inference across dictionaries. In John P. McCrae, Francis Bond, Paul Buitelaar, Philipp Cimiano, Thierry Declerck, Jorge Gracia, Ilan Kernerman, Elena Montiel-Ponsoda, Noam Ordan, and Maciej Piasecki, editors, *Proceedings of the LDK 2017 Workshops: 1st Workshop on the OntoLex Model (OntoLex-2017), Shared Task on Translation Inference Across Dictionaries & Challenges for Wordnets co-located with 1st Conference on Language, Data and Knowledge (LDK 2017), Galway, Ireland, June 18, 2017.*, volume 1899 of *CEUR Workshop Proceedings*, pages 103–112. CEUR-WS.org, 2017.
6. Mikel L Forcada, Mireia Ginestí-Rosell, Jacob Nordfalk, Jim O’Regan, Sergio Ortiz-Rojas, Juan Antonio Pérez-Ortiz, Felipe Sánchez-Martínez, Gema Ramírez-Sánchez, and Francis M Tyers. Apertium: a free/open-source platform for rule-based machine translation. *Machine translation*, 25(2):127–144, 2011.

7. Philip Gage. A new algorithm for data compression. *C Users J.*, 12(2):23–38, February 1994.
8. Thanh-Le Ha, Jan Niehues, and Alexander Waibel. Toward multilingual neural machine translation with universal encoder and decoder. In *Proceedings of the 13th International Workshop on Spoken Language Translation (IWSLT)*, Seattle, USA, December 8-9 2016.
9. Varga István and Yokoyama Shoichi. Bilingual dictionary generation for low-resourced language pairs. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 2*, pages 862–870. Association for Computational Linguistics, 2009.
10. Hiroyuki Kaji and Toshiko Aizono. Extracting word correspondences from bilingual corpora based on word co-occurrences information. In *Proceedings of the 16th conference on Computational linguistics-Volume 1*, pages 23–28. Association for Computational Linguistics, 1996.
11. Hiroyuki Kaji, Shin’ichi Tamamura, and Dashtseren Erdenebat. Automatic Construction of a Japanese-Chinese Dictionary via English. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC’08)*, Marrakech, Morocco, 2008. European Language Resources Association (ELRA).
12. Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. Opennmt: Open-source toolkit for neural machine translation. In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada, July 2017. Association for Computational Linguistics.
13. Wang Ling, Isabel Trancoso, Chris Dyer, and Alan W. Black. Character-based neural machine translation. *CoRR*, abs/1511.04586, 2015.
14. Gideon S Mann and David Yarowsky. Multipath translation lexicon induction via bridge languages. In *Proceedings of the Second Meeting of the North American Chapter of the Association for Computational Linguistics on Language Technologies*, pages 1–8. Association for Computational Linguistics, 2001.
15. Kyonghee Paik, Francis Bond, and Satoshi Shirai. Using multiple pivots to align Korean and Japanese lexical resources. In *Proc. of the NLPRS-2001 Workshop on Language Resources in Asia*, pages 63–70, 2001.
16. Thomas Proisl, Philipp Heinrich, Stefan Evert, and Besim Kabashi. Translation inference across dictionaries via a combination of graph-based methods and co-occurrence statistics. In *Proceedings of the LDK 2017 Workshops: 1st Workshop on the OntoLex Model (OntoLex-2017), Shared Task on Translation Inference Across Dictionaries & Challenges for Wordnets co-located with 1st Conference on Language, Data and Knowledge (LDK 2017), Galway, Ireland, June 18, 2017.*, 2017.
17. Xabier Saralegi, Iker Manterola, and Inaki San Vicente. Analyzing methods for improving precision of pivot based bilingual dictionaries. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 846–856. Association for Computational Linguistics, 2011.
18. Charles Schafer and David Yarowsky. Inducing translation lexicons via diverse similarity measures and bridge languages. In *proceedings of the 6th conference on Natural language learning-Volume 20*, pages 1–7. Association for Computational Linguistics, 2002.
19. Rico Sennrich, Barry Haddow, and Alexandra Birch. 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, August 2016. Association for Computational Linguistics.

20. Daphna Shezaf and Ari Rappoport. Bilingual lexicon generation using non-aligned signatures. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 98–107. Association for Computational Linguistics, 2010.
21. Jonas Sjöbergh. Creating a free digital Japanese-Swedish lexicon. In *Proceedings of PACLING*, pages 296–300. Citeseer, 2005.
22. Stephen Soderland, Oren Etzioni, Daniel S Weld, Michael Skinner, Jeff Bilmes, et al. Compiling a massive, multilingual dictionary via probabilistic inference. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 1-Volume 1*, pages 262–270. Association for Computational Linguistics, 2009.
23. Ralf Steinberger, Andreas Eisele, Szymon Kloczek, Spyridon Pilos, and Patrick Schlüter. DGT-TM: A freely available Translation Memory in 22 languages. In *Proceedings of the 8th International Conference on Language Resources and Evaluation (LREC12)*, pages 454–459, Istanbul, Turkey, May 2012.
24. Kumiko Tanaka and Kyoji Umemura. Construction of a bilingual dictionary intermediated by a third language. In *Proceedings of the 15th conference on Computational linguistics-Volume 1*, pages 297–303. Association for Computational Linguistics, 1994.
25. Daniel Torregrosa, Mihael Arcan, Sina Ahmadi, and John P. McCrae. TIAD 2019 Shared Task: Leveraging Knowledge Graphs with Neural Machine Translation for Automatic Multilingual Dictionary Generation. In *Proceedings of the Translation Inference Across Dictionaries (TIAD 2019) - Shared Task*, Leipzig, Germany, 2019.
26. Takashi Tsunakawa, Naoaki Okazaki, Xiao Liu, and Jun’ichi Tsujii. A Chinese-Japanese lexical machine translation through a pivot language. *ACM Transactions on Asian Language Information Processing (TALIP)*, 8(2):9, 2009.
27. Marta Villegas, Maite Melero, Núria Bel, and Jorge Gracia. Leveraging RDF Graphs for Crossing Multiple Bilingual Dictionaries. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Thierry Declerck, Marko Grobelnik, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France, 2016. European Language Resources Association (ELRA).