

Coursera Corpus Mining and Multistage Fine-Tuning for Improving Lectures Translation

Haiyue Song¹, Raj Dabre², Atsushi Fujita³, Sadao Kurohashi⁴

^{1,4}Kyoto University

^{2,3}National Institute of Information and Communications Technology

¹song@nlp.ist.i.kyoto-u.ac.jp, ⁴kuro@i.kyoto-u.ac.jp

^{2,3}{raj.dabre, atsushi.fujita}@nict.go.jp

Abstract

Lectures translation is a case of spoken language translation and there is a lack of publicly available parallel corpora for this purpose. To address this, we examine a framework for parallel corpus mining which is a quick and effective way to mine a parallel corpus from publicly available lectures at Coursera. Our approach determines sentence alignments, relying on machine translation and cosine similarity over continuous-space sentence representations. We also show how to use the resulting corpora in a multistage fine-tuning based domain adaptation for high-quality lectures translation. For Japanese–English lectures translation, we extracted parallel data of approximately 40,000 lines and created development and test sets through manual filtering for benchmarking translation performance. We demonstrate that the mined corpus greatly enhances the quality of translation when used in conjunction with out-of-domain parallel corpora via multistage training. This paper also suggests some guidelines to gather and clean corpora, mine parallel sentences, address noise in the mined data, and create high-quality evaluation splits. For the sake of reproducibility, we have released our code for parallel data creation.

Keywords: parallel corpus, machine translation, educational corpus, alignment, lecture translation, fine-tuning

1. Introduction

In recent years, massive open online courses (MOOCs) have proliferated and have enabled people to attend lectures regardless of their geographical location. Typically, such lectures are taught by professors from certain universities and are made available through video recordings. It is common for these lectures to be taught in a particular language and the video is accompanied by subtitles. These subtitles are then translated into other languages so that speakers of those languages also benefit from the lectures. As manual translation is a time-consuming task and there are a large number of lectures, having a high-quality machine translation (MT) system can help ease the distribution of knowledge to a large number of people across the world.

Given the fact that most lectures are spoken in English, translating English subtitles to other languages is an urgent and important task. On the other hand, there are also lectures taught in other languages than English. For instance, several universities in Japan offer online lecture courses, mostly taught in Japanese. Enabling non-Japanese speakers to participate in these courses by translating Japanese lecture subtitles into other languages, including English, is also an important challenge. The TraMOOC project (Kordoni et al., 2015) aims at improving the accessibility of European languages through MT. They focus on collecting translations of lecture subtitles and constructing MT systems for eleven European and BRIC languages. However, the amount of parallel resources involving other languages, such as Chinese and Japanese, are still quite low.

Subtitle translation falls under spoken language translation. Past studies in spoken language translation mainly focused on subtitles for TED talks (Cettolo et al., 2012). Even though the parallel data in this domain should be exploitable for lectures translation to some degree, university lectures are devoted mainly for educational purposes, and

the subtle differences in domains may hinder translation quality. To obtain high-quality parallel data, professional translators are typically employed to translate. However, the cost is often very high and thus using this way to produce large quantities of parallel data is economically infeasible, especially for universities and non-profit organizations. In the case of online lectures and talks, subtitles are often translated by crowdsourcing (Behnke et al., 2018) which involves non-professional translators. The resulting translation can thus be often inaccurate and quality control is indispensable. There are many automatic ways to find parallel sentences from roughly parallel documents (Tiedemann, 2016).¹ In particular, MT-based approaches are quite desirable because of their simplicity and it is possible to use existing translation models to extract additional parallel data. However, using an MT system trained on data from another domain can give unreliable translations which can lead to parallel data of low quality.

In this paper, we propose a new method which combines machine translation and similarities of sentence vector representations to automatically align sentences between roughly aligned document pairs. As we are interested in educational lectures translation, we focus on extracting parallel data from Coursera lectures. Using our method, we have compiled a Japanese–English parallel corpus of approximately 40,000 lines. We have also created test and development sets, consisting of 2,000 and 500 sentences, respectively, through manual filtering. All our data splits are at a document level and thus can be used to evaluate techniques that exploit context. To show the value of our extracted parallel data, we conducted experiments using it

¹cf. Comparable corpora, such as Wikipedia, i.e., pairs of documents containing the contents in same topic but their parallelism is not necessarily guaranteed and the corresponding sentences are not necessarily in the same order.

in conjunction with out-of-domain corpora for Japanese–English lectures translation. We show that although the small in-domain corpus is ineffective by itself, it is very useful when combined with out-of-domain corpora using domain adaptation techniques.

The contributions of this paper are as follows.

- A simple but accurate technique to extract parallel sentences from noisy parallel document pairs. We have made the code publicly available.²
- A Japanese–English parallel corpus usable for benchmarking educational lectures translation: high quality development and test sets are guaranteed through manual verification, whereas the training data is automatically extracted and is potentially noisy.
- An extensive evaluation of robust domain adaptation techniques leveraging out-of-domain data to improve lectures translation for the Japanese–English pair.

2. Related Work

Our work focuses on three main topics: spoken language corpora in the educational domain, parallel corpora alignment, and domain adaptation for low-resource MT.

2.1. Educational Spoken Language Corpora

Most spoken language corpora come from subtitles of online videos and a sizeable portion of subtitles are available for educational videos. Such videos are recorded lectures that form a part of an online course provided by an organization which is usually non-profit. Nowadays, many MOOCs³ have become available and help people to conveniently acquire knowledge regardless of their location. The TraMOOC project (Kordoni et al., 2015; Kordoni et al., 2016) aims at providing access to multilingual subtitles of online courses by using MT. Coursera⁴ is an extremely popular platform for MOOCs and a large number of lectures have multilingual subtitles which are created by professional and non-professional translators alike. A similar MOOC site is Iversity.⁵

Another existing spoken language corpora is for TED talks (Cettolo et al., 2012).⁶ Most talks are for the purpose of educating people, even though they do not belong to the educational lectures domain. On a related note, Opensubtitles (Tiedemann, 2016)⁷ is a collection of subtitles in multiple languages but mixes several domains.

2.2. Parallel Corpus Alignment

Extracting parallel data usable for MT involves crawling documents and aligning translation pairs in the corpora. To align translations, one can use crowdsourcing services (Behnke et al., 2018). However, this can be extremely time-consuming if not expensive. Previous research (Abdelali et

al., 2014) focused on collecting data from AMARA platform (Jansen et al., 2014). They usually aim at European and BRIC languages, such as German, Polish, and Russian. Using automatic alignment methods are more desirable, because they can help extract parallel sentences that are orders of magnitude larger than those that can be obtained by manual translation, including crowdsourcing. Although the quality of the extracted parallel sentences might be low, relying on comparable corpora can help address quality issues (Wolk, 2015) where one can use time-stamp to roughly align corresponding documents (Abdelali et al., 2014; Tiedemann, 2016). In order to obtain high-quality parallel data from these documents, MT-based methods (Sennrich and Volk, 2010; Sennrich and Volk, 2011; Liu et al., 2018) and similarity-based methods (Bouamor and Sajjad, 2018; Wang and Neubig, 2019) can be combined with dynamic programming (Utsuro et al., 1994) for fast and accurate sentence alignment. The LASER tool (Chaudhary et al., 2019)⁸ offers another way to align sentence pairs automatically in an unsupervised fashion.

2.3. Domain Adaptation for Neural Machine Translation

At present, neural machine translation (NMT) is known to give higher quality of translation. To train a sequence-to-sequence model (Sutskever et al., 2014), attention-based model (Bahdanau et al., 2015) or self-attention based model (Vaswani et al., 2017), we need a large parallel corpus for high-quality translation (Zoph et al., 2016; Koehn and Knowles, 2017). In the case of the news domain, there are many corpora, e.g., News Commentary (Tiedemann, 2012), containing large number of parallel sentences that enable high-quality translation. In contrast, for educational lectures translation, only relatively small datasets are available. Transfer learning through fine-tuning an out-of-domain model on the in-domain data (Luong and Manning, 2015; Sennrich et al., 2016a; Zoph et al., 2016; Chu et al., 2017) is the most common way to overcome the lack of data. However, approaches based on fine-tuning suffer from the problem of over-fitting which can be addressed by strong regularization techniques (Hinton and Salakhutdinov, 2006; Chelba and Acero, 2006; Miceli Barone et al., 2017; Thompson et al., 2019). Furthermore, the domain divergence between the out-of- and in-domain corpora is another issue.

3. Our Framework for Mining Coursera Parallel Corpus

This section describes our general framework to compile a parallel corpus in the educational lectures domain, relying on Coursera. Figure 1 gives an overview of our framework, where we assume the availability of in-domain parallel documents (top-left), such as those available from Coursera, and out-of-domain parallel sentences (bottom-right). We give details about the way we prepare the source document pairs, align the sentence pairs in the documents, and create evaluation splits for benchmarking.

²<https://github.com/shyyhs/CourseraParallelCorpusMining>

³<http://mooc.org>

⁴<https://www.coursera.org>

⁵<https://iversity.org>

⁶<https://wit3.fbk.eu/mt.php?release=2017-01-ted-test>

⁷<https://www.opensubtitles.org>

⁸<https://github.com/facebookresearch/LASER>




Figure 1: Overview of our framework.

3.1. Crawling and Cleaning Parallel Documents

Our framework exploits a set of in-domain parallel documents, i.e., translated lectures subtitles, available at Coursera. First, the list of available courses at Coursera is obtained, for instance by scraping. Then, all the subtitles in all available languages are downloaded from each course in the list, for instance by using Coursera-dl.⁹ Note that this results in a multilingual document-aligned subtitle corpus. From the extracted document pairs, we retain those in which the order of the sentences are roughly in the same order of the time-stamps of the lecture.

To obtain high-quality translations, the crawled parallel documents must be intensively cleaned. We consider the following 5-step procedure.

Step 1. Normalizing Text Encoding: First, all the documents are converted into UTF-8 and variants of character encodings were normalized (NFKC).

Step 2. Detecting Language Mismatch: The content of a document is sometimes of a different language than mentioned on the website. Thus, we have to detect and exclude such mismatches. Language detection tools, such as langdetect (Raffel et al., 2019),¹⁰ and/or handwritten rules can be used.

Step 3. Splitting Lines into Sentences: Since not all lines within a document are segmented into sentences, sentence splitting is necessary. Punctuation marks can be regarded as the clue. Files containing no punctuation marks are discarded, because we currently have no reliable way to deal with them.

Step 4. Removing Meta Tokens: Some tokens indicating meta-information, such as “[Music]” and “<<,” in each file are removed.

Step 5. Eliminating Imbalanced Document Pairs: Some document pairs are imbalanced in the sizes: one side has twice or more sentences than the other. Such pairs are eliminated.

3.2. Sentence Alignment

Given crawled and cleaned document pairs, we identify sentence alignments using dynamic programming (DP) as

in Utsuro et al. (1994), assuming the monotonicity of subtitles: the corresponding sentences in each pair of documents are roughly in the same order. Our assumption is based on the fact that the sentences in subtitle corpora are often constructed in accordance with the time-stamps of the sentences they correspond with. Consequently, comparing all pairs of sentences between document pairs is unnecessary (Abdelali et al., 2014). Our DP algorithm relies on MT system, sentence similarity measure, and some constraints based on the nature of lectures subtitles.

3.2.1. Training an Initial MT System

To compute the similarity of an arbitrary pair of sentences in two different languages, we first need to represent them in a common space. One option is to translate one side into the other language using an MT system (Sennrich and Volk, 2010). To train such a system, we can leverage any existing parallel data in related or even distant domains. The MT system should generate translations as accurately as possible. In practice, domain adaptation techniques (Chu et al., 2017) are most useful in training an accurate MT system.

3.2.2. Similarity Measure

The key component in the DP algorithm is the matching function, i.e., similarity measure in our context. Existing methods, such as that in Sennrich and Volk (2010), used sentence-level BLEU scores (Papineni et al., 2002) of machine-translated source sentence against the actual target language sentence as their similarity score: formally,

$$Sim_{BLEU}(f_i, e_j) = BLEU(MT(f_i), e_j), \quad (1)$$

where f_i and e_j are the i -th sentence in the source document and the j -th sentence in the target document, respectively. However, due to the lack of in-domain data, MT system can give only translations of low quality and thus the BLEU scores can be misleading, especially for distant language pairs, such as Japanese and English.

An alternative way is to directly compute cosine similarity of a given sentence pair (Bouamor and Sajjad, 2018), relying on pre-trained multilingual word embeddings to represent sentences in different languages with the same vector space through element-wise addition of word embeddings (Mikolov et al., 2013). However, cross-lingual embeddings are often not accurate for distant language pairs, especially if they have been pre-trained on data from another domain. Taking inspiration from both these approaches, we employ MT combined with cosine similarity of sentence embeddings to measure the similarity of two sentences in different languages, formulated as follows.

$$Sim_{EMB}(f_i, e_j) = \cos(emb(MT(f_i)), emb(e_j)). \quad (2)$$

As in Eq. (1), we first translate each sentence in the source language document into the target language. In practice, we prefer to have English as the target language, because this eliminates the need for cross-lingual vectors and this also enables us to use an abundance of high-quality English pre-trained vectors for several domains. $emb(\cdot)$ represents the embedding of the given sentence, which can be computed by averaging the embeddings of words in that sentence, as in Mikolov et al. (2013).

⁹<https://github.com/coursera-dl/coursera-dl>

¹⁰<https://github.com/Mimino666/langdetect>

Algorithm 1: Document-aware sentence filtering

Input : *DocPairs*, *volume*, *ratio***Output:** *DocPairs*, *SentencePairs*

```
1 SentencePairs ← {};  
2 while |SentencePairs| < volume do  
3   | DocPair ← pickBestDocPair(DocPairs);  
4   | DocPairs ← DocPairs \ {DocPair};  
5   | CandidatePairs ← getAlignments(DocPair);  
6   | Correct ← {};  
7   | foreach Pair ∈ CandidatePairs do  
8     | Judge ← manualEvaluation(Pair);  
9     | if Judge == good then  
10    |   | Correct ← Correct ∪ {Pair};  
11  | if |Correct| > |CandidatePairs| * ratio then  
12  |   | SentencePairs ← SentencePairs ∪ Correct;
```

3.2.3. Constraints

To control the alignment quality, we should introduce the following three types of constraints in our DP algorithm.

- A pair of sentences will not be a match if their similarity is lower than a pre-determined threshold, *th*.
- A pair of sentences will not be a match if one of them is *k* times longer than the other.
- Only 1-1, 0-1, and 1-0 matching are allowed.

3.3. Creating High-quality Evaluation Sets

To benchmark the performance of educational lectures translation, a high-quality test set is indispensable. If we also have another set of high-quality translations, it can be a useful development set for tuning MT systems. We resort to manual cleaning of the scored and aligned sentence pairs obtained using the previous step.

We first sort all document pairs in the descending order of the average similarity of all aligned sentence pairs within each document pair. We then subject these sorted and sentence aligned pairs to human evaluation using Algorithm 1 in order to obtain high-quality test and development sets, where the target volume of each set (*volume*) and document-level comparability (*ratio*) are the two parameters. We use the remaining sentence aligned document pairs for training. Our test, development, and training sets are all constructed at the document level and thus our corpora can be used to evaluate document-level translation (Voita et al., 2019; Wang et al., 2019b; Tan et al., 2019).

4. Creating Japanese–English Parallel Data

Although we have actually extracted document pairs for all available courses on Coursera, henceforth, we report on an application of our framework to create a Japanese–English Coursera dataset.

4.1. Cleaning Documents

Our framework is mostly language independent. The only language specific processes are tokenization and language mismatch detection. We first segmented the both English and Japanese paragraphs with full-stop (“.”), exclamation (“!”), and question marks (“?”) in Latin encoding and their

Algorithm 2: Language detection procedure

Input : *Doc*, *n*, *m*, *EnglishChar*, *JapaneseChar***Output:** *Label*

```
1 Sentences ← pickRandomSentences(Doc, n);  
2 EnSentence ← 0;  
3 JaSentence ← 0;  
4 foreach Sentence ∈ Sentences do  
5   | EnChar ← 0;  
6   | JaChar ← 0;  
7   | OtherChar ← 0;  
8   | foreach Char ∈ Sentence do  
9     | if Char ∈ EnglishChar then  
10    |   | EnChar ← EnChar + 1;  
11    | else if Char ∈ JapaneseChar then  
12    |   | JaChar ← JaChar + 1;  
13    | else  
14    |   | OtherChar ← OtherChar + 1;  
15  | if EnChar > JaChar then  
16  |   | EnSentence ← EnSentence + 1;  
17  | else  
18  |   | JaSentence ← JaSentence + 1;  
19  | if EnSentence ≥ m then  
20  |   | Label ← En;  
21  | else if JaSentence ≥ m then  
22  |   | Label ← Ja;  
23  | else  
24  |   | Label ← Noise;
```

full-width counterparts in UTF-8 followed by a space or the end of line. Then, we tokenized Japanese and English sentences, using Juman++ (Tolmachev et al., 2018)¹¹ and NLTK,¹² respectively.

Algorithm 2 shows our rule-based language detection procedure for the Japanese–English setting. It judges whether the given document is in Japanese or English according to the number of sentences within the document that belong to each language, where the language of each sentence is determined on the basis of the number of English and Japanese characters. More specifically, we define a set of characters, *EnglishChar*, with “a” to “z” and “A” to “Z,” and another set of characters, *JapaneseChar*, with *hiragana* and *katakana*. We empirically set the two thresholds: *n* = 10 and *m* = 8.

We evaluated the performance of the langdetect tool (Rafel et al., 2019) and our algorithm on 100 sample documents, and found that the langdetect has one misclassification whereas ours worked perfectly with 1.00 precision and recall, presumably thanks to the cleanness of the Coursera data. Considering that our simple method worked reasonably accurately, we chose the results of our method for the following steps.

4.2. Creating Initial MT System

As mentioned in Section 1., the TED parallel corpus (Cetolo et al., 2012) is from the spoken language domain and thus it is most similar to the spoken educational lectures

¹¹<https://github.com/ku-nlp/jumanpp>¹²<https://www.nltk.org>

Dataset	Train	Dev	Test
ASPEC	1.0M	1,790	1,812
TED	223k	1,354	1,194

Table 1: Number of sentence pairs in each corpus.

Training schedule	BLEU
A	4.1
T	12.2
AT	14.6
A → T	13.9
A → AT	15.0

Table 2: BLEU score for Ja→En on TED test set. A, T, and AT respectively stand for ASPEC training set, TED training set, and their balanced mixture. “→” indicates that the model trained on the left-hand side data is fine-tuned on the right-hand side data.

domain. However, given its small size, it can lead to only an unreliable MT system. Therefore, we decided to use a larger out-of-domain ASPEC corpus (Nakazawa et al., 2016)¹³ to build a better MT system. Table 1 gives the statistics of the ASPEC and TED corpora that we used to train our initial MT system. We compared fine-tuning and mixed fine-tuning approaches proposed by Chu et al. (2017). When performing mixed fine-tuning on the concatenation of both two corpora, the TED corpus was over-sampled to match the size of the ASPEC corpus. We trained our NMT models using tensor2tensor with its default hyper-parameters. Refer to Section 5.3. for further details on training configurations.

So far, we do not have a test set for the target domain, i.e., Coursera. We therefore evaluated the performance of the MT systems with BLEU score (Papineni et al., 2002) on the TED test set. Table 2 gives the results, where “A,” “T,” and “AT” stand for ASPEC, TED, and their balanced mixture, respectively and “→” means fine-tuning on the right-hand side data. The model first pre-trained and then mixed fine-tuned, i.e., A→AT, gave the best result on the TED test set. Thus, we used this model for sentence alignment.

4.3. Creating Benchmarking Splits

Finally, we extracted parallel sentences using the initial Japanese–English MT system to translate the Japanese sentences into English and the English embeddings available at the NLPL word embeddings repository¹⁴ to compute sentence similarity. With the two parameters for constraining the DP algorithm, i.e., $th = 0.92$ and $k = 2$, we obtained a total of 43,549 pairs of sentences from 884 document pairs. Then, following the procedure in Section 3.3., we manually¹⁵ created the test and development sets, taking the most reliable document pairs. We set 2,000 and 500 sentences as the target *volume* for the test and development sets, respec-

¹³We selected the best 1.0 million sentence pairs.

¹⁴<http://vectors.nlpl.eu/repository>, ID 40: Word2Vec Continuous Skipgram trained on English CoNLL2017 corpus. Download

¹⁵The checker is not a native English or Japanese speaker, but has the N1 certification (highest level) of the Japanese Language Proficiency Test and 99 points in TOEFL iBT.

	# of document pairs	# of aligned lines	# of deleted lines
Test	50	2,005	177
Dev	16	541	56
Train	818	40,770	-

Table 3: Our Japanese–English Coursera parallel data.

Dataset	English	Japanese
	Mean / Median / s.d.	Mean / Median / s.d.
ASPEC	25.4 / 23 / 11.4	27.5 / 20 / 12.0
TED	20.4 / 17 / 13.9	19.8 / 16 / 14.1
Coursera	21.1 / 19 / 11.1	22.2 / 20 / 11.8

Table 4: Statistics on the sentence length.

LM	Corpus	ASPEC	TED	Coursera
	ASPEC	-1.147	-3.013	-2.926
TED	-2.962	-1.097	-2.255	
Coursera	-2.658	-2.335	-0.760	

Table 5: Average per-token log-likelihood.

tively, and set $ratio = 0.50$. As shown in Table 3, a total of 2,779 sentence pairs drawn from 66 documents were manually judged in approximately 4 hours and about 8.4% of them ((177+56)/2,779) were filtered out. All the aligned lines in the remaining document pairs were used to compose the training data. Since we have not performed any manual filtering, this set is presumably noisier than the development and test sets.

4.4. Analysis

We compared our Coursera dataset with the ASPEC and TED datasets regarding average sentence length and domain similarity. Table 4 gives a summary of sentence length: the number of tokens segmented by Juman++ and NLTK for Japanese and English, respectively. Coursera dataset is in between ASPEC and TED in its average sentence length, but relatively closer to TED than to ASPEC. We also computed the similarity between datasets using language model (LM). First, we trained a 4-gram LM on the lower-cased version of English side of each training set. We then computed the per-token log-likelihood of these training sets with each of these LMs. As shown in Table 5, three datasets are visibly distant to each other. Nevertheless, TED seems relatively more exploitable than ASPEC for helping to translate Coursera datasets, presumably because they comprise spoken language unlike ASPEC.

5. Japanese–English Lectures Translation

We now describe how we can utilize the parallel corpus compiled as mentioned in the previous section for Japanese–English educational lectures translation.

5.1. Multistage Fine-Tuning

Although NMT needs a large amount of parallel data to work well, its performance is very sensitive to the domain of the dataset and to the order in which datasets are included in the training. As such, it is common to divide training

ID	Training schedule	Ja→En	En→Ja	ID	Training schedule	Ja→En	En→Ja
A1	A	13.6	10.4	B2	AT	24.5	13.3
A2	A AT	25.6	13.5	B3	AT ATC	26.8	17.0
A3	A AT ATC	27.5	18.0	B4	AT ATC TC	25.1	17.0
A4	A AT ATC TC	25.9	17.6	B5	AT ATC TC C	23.8	17.7
A5	A AT ATC TC C	24.4	17.7	B6	AT ATC C	24.1	17.8
A6	A AT ATC C	24.7	18.5	B7	AT TC	26.4	17.2
A7	A AT TC	26.9	17.5	B8	AT TC C	23.9	17.5
A8	A AT TC C	24.3	17.6	B9	AT C	22.9	17.7
A9	A AT C	23.8	17.2	B10	ATC	22.2	15.8
A10	A ATC	25.7	17.9	B11	ATC TC	22.0	15.4
A11	A ATC TC	25.2	17.4	B12	ATC TC C	21.2	16.6
A12	A ATC TC C	24.3	17.5	B13	ATC C	21.2	16.5
A13	A ATC C	24.3	17.8	B14	TC	15.3	11.2
A14	A TC	25.4	17.6	B15	TC C	16.1	12.2
A15	A TC C	23.8	17.1	B16	C	6.2	6.4
A16	A C	21.6	16.9	D3	AT T	23.2	12.2
C3	A AT T	24.0	12.2	D4	AT T TC	24.6	16.6
C4	A AT T TC	25.8	16.9	D5	AT T TC C	22.3	17.0
C5	A AT T TC C	23.8	17.6	D6	AT T C	22.5	17.0
C6	A AT T C	23.4	17.3	D10	T	17.5	8.9
C10	A T	23.9	12.2	D11	T TC	20.6	13.8
C11	A T TC	25.3	16.3	D12	T TC C	19.8	14.4
C12	A T TC C	23.6	16.6	D13	T C	19.5	14.6
C13	A T C	22.7	16.9	F14	AC	16.2	13.6
E14	A AC	23.2	17.9	F15	AC C	16.3	13.9
E15	A AC C	22.1	16.5				

Table 6: BLEU scores for all the multistage training options examined in our experiment. Models A1–A16 and B2–B16 represent all the 31 ($= 2^5 - 1$) sub-paths of the A→AT→ATC→TC→C flow. Bold indicates the **initial training**, and red-, blue-, and grey-colored cells mean **inflation**, **deflation**, and **replacement** of training data, respectively.

into multiple stages where each stage uses data from different domains to maximize the impact of the domain-specific training data. As we have larger parallel corpora from other domains, such as TED (0.2M pairs; non-educational spoken domain) and ASPEC (3.0M pairs; scientific domain), we can leverage domain adaptation techniques, such as fine-tuning and mixed fine-tuning (Chu et al., 2017). Furthermore, Imankulova et al. (2019) and Dabre et al. (2019) showed that training in multiple stages where each stage contains different proportions of various types of training data leads to the best results. Following them, we decided to conduct an extensive experiment with multistage training with different proportions of training data from different domains at each stage.

5.2. Datasets

As in the previous section, we performed Juman++ and NLTK tokenization for Japanese and English, respectively. Henceforth, we refer to the ASPEC training data of 1.0 million lines as “A,” the TED training data of 0.2 million lines as “T,” and the Coursera training data of 40k lines as “C.” When combining more than one dataset, we oversampled the smaller ones to match the size of the largest one. We denote the concatenated corpus by a concatenation of the letters representing them: e.g., AT for the mixture of ASPEC data with 5 times oversampled TED data, and ATC for the concatenation of ASPEC with 5 times oversampled TED data and 25 times oversampled Coursera data.

Following the observations in Section 4.4., we decided to focus on the training schedule A→AT→ATC→TC→C,

where in each stage we use datasets more similar with in-domain dataset (Wang et al., 2019a), and thoroughly evaluated all of its sub-paths. We also used T and AC for some contrastive experiments.

5.3. Settings for MT

We created a shared sub-word vocabulary for Japanese and English from ASPEC and TED training set using BPE (Sennrich et al., 2016b) with roughly 32k merge operations. This vocabulary was used for all experiments, even when a model is trained only on C.

We trained NMT models using the tensor2tensor framework (Vaswani et al., 2018)¹⁶ with its default “transformer.big” setting, such as dropout=0.2, attention dropout=0.1, optimizer=adam with beta1=0.9, beta2=0.997. We used eight Tesla V100 32GB GPUs with batch size of 4,096 sub-word tokens, and early-stopping on approximate BLEU score computed on the development set: the training process stops when the score shows no gain larger than 0.1 for 10,000 steps. When fine-tuning the model on a different dataset, we resumed the training process from the last checkpoint in the previous stage.

In the decoding step, we used the average of the last 10 checkpoints, and decoded the test sets with a beam size of 4 and a length penalty, α , of 0.6 consistently across all the models. The trained systems were evaluated with BLEU scores computed by sacreBLEU.¹⁷

¹⁶<https://github.com/tensorflow/tensor2tensor>, version 1.14.0.

¹⁷<https://github.com/mjpost/sacreBLEU>

Test set	Model	Ja→En	En→Ja
ASPEC	A (A1 in Table 6)	29.8	41.5
TED	T (D10 in Table 6)	14.7	12.2
Coursera	C (B16 in Table 6)	6.2	6.4

Table 7: BLEU scores for different test sets.

5.4. Results

Table 6 summarizes the BLEU scores of all the MT systems trained up to five training stages. The training schedule with all the stages, i.e., $A \rightarrow AT \rightarrow ATC \rightarrow TC \rightarrow C$ (A5) did not achieve the best results for both translation directions. For the Ja→En task, one of the intermediate models, $A \rightarrow AT \rightarrow ATC$ (A3), gave the best BLEU score with more than 20 points gain over the model trained only on the in-domain parallel data (B16). For the reverse direction, i.e., the En→Ja task, the schedule $A \rightarrow AT \rightarrow ATC \rightarrow C$ (A6) achieved the best BLEU score with a 12.1 point BLEU gain. In contrast, training a model directly in one stage on ATC (B10) gave significantly lower results than the multi-stage results.

Whenever new training data were introduced (marked red in Table 6), the BLEU scores were improved.¹⁸ This is mostly in line with the observations of Chu et al. (2017); starting from out-of-domain and ending with a mixture of out-of- and in-domain data gives the best results for in-domain translation. As shown in Table 5, A is most dissimilar to C, and T is most similar to C. As such, it seems reasonable to gradually introduce the in-domain data by relying on data from related domain for intermediate training steps. According to Dabre et al. (2019), the final stage of fine-tuning on C should give the best translation quality. However, in our setting, this holds true only for some cases in the En→Ja task, suggesting the necessity of hyperparameter tuning for fine-tuning on deflated training data (marked blue in Table 6).

This shows the importance of exhaustively exploring all settings, which confirmed and/or revealed the followings.

- Leveraging out-of-domain data through multistage training is invaluable.
- Gradually inflating the data starting from out-of-domain corpus and adding the in-domain corpus at the end should give the best possible translation quality.

One peculiarity of our results is that the BLEU scores for the Ja→En task were significantly higher than the En→Ja task, which is a reversal of a general tendency for this language pair (Imamura and Sumita, 2018; Nakazawa et al., 2019), even though the BLEU scores in different languages are not directly comparable. Table 7 gives a comparison of three different translation tasks. Upon manual investigation, we identified that the En→Ja translations in TED and Coursera tasks tend to be much shorter than the reference translation, receiving around 0.7 brevity penalty. When we tuned the length penalty for decoding on the development set, we observed 0.5 to 1.0 point BLEU gains on the test set

¹⁸Compare the pairs (A1, A2), (A2, A3), (A1, A10), (B2, B3), (C3, C4), (C10, C11), (D3, D4), (D10, D11), and (A1, E14).

$A \rightarrow AT \rightarrow ATC$	Ja→En	En→Ja
Iteration 1 (A3 in Table 6)	27.5	18.0
Iteration 2	27.2	17.9

Table 8: BLEU scores in different iterations.

for the En→Ja task, but this is not enough to flip the BLEU score tendencies. Another possible reason is the nature of translationese (Rubino et al., 2016). Whereas ASPEC contains mainly Japanese-to-English translations, most talks in TED and Coursera datasets are English-to-Japanese translations. Yet another reason is the difference between written (ASPEC) and spoken (TED and Coursera) languages. We leave deeper exploration for the future.

5.5. Iterative Refinement of Aligned Data

Having obtained a better MT system than the initial one, we can iterate the whole process illustrated in Figure 1, i.e., extracting the best possible parallel sentences using an MT system, and training a new MT system on the new parallel corpus, in order to maximize the quality of both the parallel corpus and the MT system.

To verify the impact of repeating an iteration, we took the best-performing Ja→En MT system, i.e., $A \rightarrow AT \rightarrow ATC$ (A3), and performed sentence alignment for the document pairs used as source for the training set, retaining test and development sets. The re-aligned training data for C were used in the $A \rightarrow AT \rightarrow ATC$ training schedule, where the models until $A \rightarrow AT$ were identical to those obtained in the first iteration, since they did not see C at all.

Table 8 compares the BLEU scores achieved by the $A \rightarrow AT \rightarrow ATC$ models in the first two iterations. Unfortunately, we do not see any improvement in translation quality. We can speculate the following two reasons.

- The dataset of approximately 40,000–45,000 sentences is too small to have any visible impact on translation quality.
- The best possible sentence alignments for Coursera data were already found, owing to our algorithm, similarity measure, and/or the initial MT system trained only on ASPEC and TED.

Nevertheless, our observation does not necessarily hold for every language pair. We thus encourage researchers to try iterative refinement of training data in their own experimental settings.

5.6. Indirect Assessment of the Created In-Domain Data

In this section, we evaluate the superiority of our sentence alignment method, presented in Section 3.2. (henceforth, MT+CS), over other methods, extrinsically, through MT performance. The following two similarity measures were additionally implemented and tested.

Unsupervised: Cosine similarity over the cross-lingual sentence embeddings, learned by an unsupervised method, called VecMap (Artetxe et al., 2017).¹⁹

¹⁹<https://github.com/artetxem/vecmap>

	# of aligned lines	BLEU	
		Ja→En	En→Ja
Unsupervised	40,452	4.0	4.3
MT+BLEU	42,672	2.8	3.4
MT+CS (B16 in Table 6)	40,770	6.2	6.4

Table 9: BLEU scores achieved with only Coursera parallel data extracted by different similarity measures.

MT+BLEU: The metric in Eq. (1), where the MT system was identical to the one used for our similarity metric, i.e., Eq. (2).

These two methods also rely on thresholding of the similarity scores of sentence pairs to get rid of potential noise. We set the threshold to a value such that the number of resulting sentence pairs is roughly the same as the number of pairs produced by our proposed method.

Table 9 compares the number of extracted parallel sentences and the BLEU scores obtained by NMT systems trained only on the automatically aligned in-domain training data. Whereas using BLEU as a measure of sentence similarity for alignment was bad, unsupervised cross-lingual embeddings gave more reliable similarity scores leading to better aligned sentences. Our MT+CS method which combines these two methods was able to give a parallel corpus, giving the highest BLEU score.

However, this does not completely justify the superiority of our similarity measure, because we have created the test and development sets relying on the MT+CS method. This could introduce a bias in the resulting set toward this particular alignment method. Another concern is the difficulty of the translations. Even though we have obtained reasonably high BLEU scores in our experiment (Table 6), due to heavy reliance on the word embeddings, the test and development sets may contain relatively easy sentence pairs in the sense that the sentence-level correspondences are easy to detect with such a simple method. We plan to investigate these aspects in our future work.

6. Conclusion and Future work

In this paper, we proposed a framework to create a dataset for translation of lectures from the educational domain. Specifically, we proposed a novel sentence similarity measure that combines machine translation and cosine similarity over sentence embeddings. Taking Japanese–English translation as a case study, we created a dataset of approximately 40,000, 500 and 2,000 lines of training, development, and test sets, with manual cleaning of the latter two sets to ensure that they can be used to reliably benchmark translation performance. We then utilized the automatically extracted parallel sentences to train an NMT system for Japanese–English lectures translation and show that multistage training in a domain adaptation framework leads to better translation models.

We have released our code used in our experiments for the sake of reproducibility. Given that the data crawled from Coursera is multilingually aligned at the document level, we plan to compile and provide a multilingual parallel corpus for lectures translation in the near future.

Acknowledgments

This work was carried out when Haiyue Song was taking up an internship at NICT, Japan. A part of this work was conducted under the program “Research and Development of Enhanced Multilingual and Multipurpose Speech Translation System” of the Ministry of Internal Affairs and Communications (MIC), Japan.

7. References

- Abdelali, A., Guzman, F., Sajjad, H., and Vogel, S. (2014). [The AMARA Corpus: Building Parallel Language Resources for the Educational Domain](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC)*, pages 1856–1862, Reykjavik, Iceland, May. European Language Resources Association.
- Artetxe, M., Labaka, G., and Agirre, E. (2017). [Learning Bilingual Word Embeddings with \(Almost\) no Bilingual Data](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Vancouver, Canada, July. Association for Computational Linguistics.
- Bahdanau, D., Cho, K., and Bengio, Y. (2015). [Neural Machine Translation by Jointly Learning to Align and Translate](#). In *Proceedings of the Third International Conference on Learning Representations (ICLR)*, San Diego, USA, May.
- Behnke, M., Miceli Barone, A. V., Sennrich, R., Sосoni, V., Naskos, T., Takoulidou, E., Stasimioti, M., van Zaanen, M., Castilho, S., Gaspari, F., Georgakopoulou, P., Kordoni, V., Egg, M., and Kermanidis, K. L. (2018). [Improving Machine Translation of Educational Content via Crowdsourcing](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC)*, Miyazaki, Japan, May. European Language Resources Association.
- Bouamor, H. and Sajjad, H. (2018). [H2@BUCC18: Parallel Sentence Extraction from Comparable Corpora Using Multilingual Sentence Embeddings](#). In *Proceedings of the Eleventh Workshop on Building and Using Comparable Corpora*, pages 43–47, Miyazaki Japan, May. European Language Resources Association.
- Cettolo, M., Girardi, C., and Federico, M. (2012). [Wit³: Web Inventory of Transcribed and Translated Talks](#). In *Proceedings of the 16th Conference of European Association for Machine Translation*, pages 261–268, Trento, Italy, May. European Association for Machine Translation.
- Chaudhary, V., Tang, Y., Guzmán, F., Schwenk, H., and Koehn, P. (2019). [Low-Resource Corpus Filtering using Multilingual Sentence Embeddings](#). *CoRR*, abs/1906.08885.
- Chelba, C. and Acero, A. (2006). [Adaptation of Maximum Entropy Capitalizer: Little Data can Help a Lot](#). *Computer Speech & Language*, 20(4):382–399.
- Chu, C., Dabre, R., and Kurohashi, S. (2017). [An Empirical Comparison of Domain Adaptation Methods for Neural Machine Translation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational*

- Linguistics (Volume 2: Short Papers)*, pages 385–391, Vancouver, Canada, July. Association for Computational Linguistics.
- Dabre, R., Fujita, A., and Chu, C. (2019). [Exploiting Multilingualism through Multistage Fine-Tuning for Low-Resource Neural Machine Translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1410–1416, Hong Kong, China, November. Association for Computational Linguistics.
- Hinton, G. and Salakhutdinov, R. (2006). [Reducing the Dimensionality of Data with Neural Networks](#). *Science*, 313(5786):504 – 507.
- Imamura, K. and Sumita, E. (2018). [Multilingual Parallel Corpus for Global Communication Plan](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC)*, Miyazaki, Japan, May. European Language Resources Association.
- Imankulova, A., Dabre, R., Fujita, A., and Imamura, K. (2019). [Exploiting Out-of-Domain Parallel Data through Multilingual Transfer Learning for Low-Resource Neural Machine Translation](#). In *Proceedings of Machine Translation Summit XVII Volume 1: Research Track*, pages 128–139, Dublin, Ireland, August. European Association for Machine Translation.
- Jansen, D., Alcalá, A., and Guzman, F. (2014). [Amara: A Sustainable, Global Solution for Accessibility, Powered by Communities of Volunteers](#). In Constantine Stephanidis et al., editors, *Universal Access in Human-Computer Interaction: Design for All and Accessibility Practice*, pages 401–411. Springer International Publishing, Cham.
- Koehn, P. and Knowles, R. (2017). [Six Challenges for Neural Machine Translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver, August. Association for Computational Linguistics.
- Kordoni, V., Cholakov, K., Egg, M., Way, A., Birch, L., Kermandidis, K., Sosoni, V., Tsoumakos, D., van den Bosch, A., Hendrickx, I., Papadopoulou, M., Georgakopoulou, P., Gialama, M., van Zaanen, M., Buliga, I., Jermol, M., and Orlic, D. (2015). [TraMOOC: Translation for Massive Open Online Courses](#). In *Proceedings of the 18th Annual Conference of the European Association for Machine Translation*, page 217, Antalya, Turkey, May. European Association for Machine Translation.
- Kordoni, V., van den Bosch, A., Kermandidis, K. L., Sosoni, V., Cholakov, K., Hendrickx, I., Huck, M., and Way, A. (2016). [Enhancing Access to Online Education: Quality Machine Translation of MOOC Content](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC)*, pages 16–22, Portorož, Slovenia, May. European Language Resources Association.
- Liu, S., Wang, L., and Liu, C.-H. (2018). [Chinese-Portuguese Machine Translation: A Study on Building Parallel Corpora from Comparable Texts](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC)*, Miyazaki, Japan, May. European Language Resources Association.
- Luong, M.-T. and Manning, C. D. (2015). [Stanford Neural Machine Translation Systems for Spoken Language Domains](#). In *Proceedings of the Twelfth International Workshop on Spoken Language Translation*, pages 76–79, Da Nang, Vietnam, December.
- Miceli Barone, A. V., Haddow, B., Germann, U., and Sennrich, R. (2017). [Regularization Techniques for Fine-tuning in Neural Machine Translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1489–1494, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). [Distributed Representations of Words and Phrases and their Compositionality](#). In C. J. C. Burges, et al., editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.
- Nakazawa, T., Yaguchi, M., Uchimoto, K., Utiyama, M., Sumita, E., Kurohashi, S., and Isahara, H. (2016). [ASPEC: Asian Scientific Paper Excerpt Corpus](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC)*, pages 2204–2208, Portorož, Slovenia, May. European Language Resources Association.
- Nakazawa, T., Doi, N., Higashiyama, S., Ding, C., Dabre, R., Mino, H., Goto, I., Pa, W. P., Kunchukuttan, A., Parida, S., Bojar, O., and Kurohashi, S. (2019). [Overview of the 6th Workshop on Asian Translation](#). In *Proceedings of the 6th Workshop on Asian Translation*, pages 1–35, Hong Kong, China, November. Association for Computational Linguistics.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). [BLEU: A Method for Automatic Evaluation of Machine Translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 311–318, Philadelphia, USA, July. Association for Computational Linguistics.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. (2019). [Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer](#). *CoRR*, abs/1910.10683.
- Rubino, R., Lapshinova-Koltunski, E., and van Genabith, J. (2016). [Information Density and Quality Estimation Features as Translationese Indicators for Human Translation Classification](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 960–970, San Diego, USA, June. Association for Computational Linguistics.
- Sennrich, R. and Volk, M. (2010). [MT-based Sentence Alignment for OCR-Generated Parallel Texts](#). In *Proceedings of the Ninth Conference of the Association for Machine Translation in the Americas (AMTA)*, Denver, USA, October–November. Association for Machine Translation in the Americas.
- Sennrich, R. and Volk, M. (2011). [Iterative, MT-based](#)

- Sentence Alignment of Parallel Texts.** In *Proceedings of the 18th Nordic Conference of Computational Linguistics (NODALIDA 2011)*, pages 175–182, Riga, Latvia, May. Northern European Association for Language Technology.
- Sennrich, R., Haddow, B., and Birch, A. (2016a). **Improving Neural Machine Translation Models with Monolingual Data.** In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany, August. Association for Computational Linguistics.
- Sennrich, R., Haddow, B., and Birch, A. (2016b). **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. Association for Computational Linguistics.
- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). **Sequence to Sequence Learning with Neural Networks.** In Z. Ghahramani, et al., editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- Tan, X., Zhang, L., Xiong, D., and Zhou, G. (2019). **Hierarchical Modeling of Global Context for Document-Level Neural Machine Translation.** In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1576–1585, Hong Kong, China, November. Association for Computational Linguistics.
- Thompson, B., Gwinnup, J., Khayrallah, H., Duh, K., and Koehn, P. (2019). **Overcoming Catastrophic Forgetting During Domain Adaptation of Neural Machine Translation.** 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)*, pages 2062–2068, Minneapolis, USA, June. Association for Computational Linguistics.
- Tiedemann, J. (2012). **Parallel Data, Tools and Interfaces in OPUS.** In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC)*, pages 2214–2218, Istanbul, Turkey, May. European Language Resources Association.
- Tiedemann, J. (2016). **Finding Alternative Translations in a Large Corpus of Movie Subtitle.** In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC)*, pages 3518–3522, Portorož, Slovenia, May. European Language Resources Association.
- Tolmachev, A., Kawahara, D., and Kurohashi, S. (2018). **Juman++: A Morphological Analysis Toolkit for Scriptio Continua.** In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 54–59, Brussels, Belgium, November. Association for Computational Linguistics.
- Utsuro, T., Ikeda, H., Yamane, M., Matsumoto, Y., and Nagao, M. (1994). **Bilingual Text, Matching using Bilingual Dictionary and Statistics.** In *Proceedings of the 15th International Conference on Computational Linguistics*, volume 2, pages 1076–1082, Kyoto, Japan, August.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). **Attention is All you Need.** In I. Guyon, et al., editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Vaswani, A., Bengio, S., Brevdo, E., Chollet, F., Gomez, A., Gouws, S., Jones, L., Kaiser, Ł., Kalchbrenner, N., Parmar, N., Sepassi, R., Shazeer, N., and Uszkoreit, J. (2018). **Tensor2Tensor for Neural Machine Translation.** In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Papers)*, pages 193–199, Boston, USA, March. Association for Machine Translation in the Americas.
- Voita, E., Sennrich, R., and Titov, I. (2019). **When a Good Translation is Wrong in Context: Context-Aware Machine Translation Improves on Deixis, Ellipsis, and Lexical Cohesion.** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1198–1212, Florence, Italy, July. Association for Computational Linguistics.
- Wang, X. and Neubig, G. (2019). **Target Conditioned Sampling: Optimizing Data Selection for Multilingual Neural Machine Translation.** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5823–5828, Florence, Italy, July. Association for Computational Linguistics.
- Wang, W., Caswell, I., and Chelba, C. (2019a). **Dynamically Composing Domain-Data Selection with Clean-Data Selection by “Co-Curricular Learning” for Neural Machine Translation.** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1282–1292, Florence, Italy, July. Association for Computational Linguistics.
- Wang, X., Tu, Z., Wang, L., and Shi, S. (2019b). **Exploiting Sentential Context for Neural Machine Translation.** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6197–6203, Florence, Italy, July. Association for Computational Linguistics.
- Wołk, K. (2015). **Noisy-Parallel and Comparable Corpora Filtering Methodology for the Extraction of Bi-lingual Equivalent Data at Sentence Level.** *Computer Science (AGH)*, 16(2):169–184.
- Zoph, B., Yuret, D., May, J., and Knight, K. (2016). **Transfer Learning for Low-Resource Neural Machine Translation.** In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1575, Austin, USA, November. Association for Computational Linguistics.